CHAPTER

11

Imperfect Competition

In previous chapters, we studied the two ends of the market power spectrum: perfect competition and monopoly. In perfect competition, a firm has no market power because it is only one of many producers in the market, the price is driven down to marginal cost, and output is relatively high. In a monopoly, one firm has market power because it is the only producer of a good in the market, price is greater than marginal cost, and output is lower. We also learned about the many pricing strategies that firms with market power can use to earn greater economic profit.

Between these two ends of the spectrum are lots of industries that are neither perfectly competitive nor monopolistic. Coke and Pepsi dominate the cola market together. Nintendo, Sony, and Microsoft dominate console video games. These companies compete but are hardly perfectly competitive. Yet they aren't stand-alone monopolies either. The *industry structure between perfect competition and monopoly* is known as **imperfect competition**.

This chapter introduces this important but sometimes complicated market structure. We begin by looking at several types of **oligopoly**, a *market structure characterized by competition among a small number of firms*. Because there are many possible ways in which oligopolistic firms compete, no single model of oligopoly exists that is applicable to every situation. Having a few competitors in an industry—rather than many or only one—can lead to many possible price and output outcomes. It's not as simple as just picking the one where price equals marginal cost or the like. We might observe many outcomes, depending on the market circumstances.

With oligopolies, firms have some market power, but not necessarily monopoly power, and there is some competition, but not perfect competition. We need to be a little more specific about aspects of the market before we can figure out what prices they will charge, how much each company will produce, and how much profit each firm will earn. Just knowing how many companies are in the market is not enough to know what will happen in an oligopolistic market. Industries with, say, four major firms can look extremely different from each other. Other factors that have an effect on price and quantity decisions in an oligopoly include: (1) whether the companies make identical products (as in an oil oligopoly) or products that are slightly different from one another (like Coke and Pepsi); (2) how intensely the companies compete; and (3) whether they compete with one another by choosing the prices they charge or the quantities they produce.

In this chapter, we present five of the most common models of how oligopolies behave, plus one additional model called **monopolistic competition**, a *type of imperfect competition where a large number of firms have some market power, but each makes zero economic profit in the long run*. Whenever you have this many models as possible explanations for market behavior, it's important to determine which one is appropriate for a specific case. This decision isn't always obvious in practice, so we discuss some ideas for determining which model is most appropriate for various real-world situations.

imperfect competition The industry structure between perfect competition and monopoly.

oligopoly A market structure characterized by competition among a small number of firms.

monopolistic competition A type of imperfect competition where a large number of firms have some market power, but each makes zero economic profit in the long run.

11.1 What Does Equilibrium Mean in an Oligopoly?

Before we introduce the different models of oligopoly, we need to lay some groundwork. Specifically, we have to expand on the idea of what an equilibrium is for these industries. The concept of equilibrium in perfect competition and in monopoly is easy. It means a price at which the quantity of the good demanded by consumers equals the quantity of the good supplied by producers. That is, the market "clears." The market is stable with no excess supply or demand, and the consumers and producers do not want to change their decisions.

The problem with applying this idea of equilibrium to an oligopolistic industry is that each company's action influences what the other companies want to do. To achieve an outcome in which no firm wants to change its decision means determining more than just a price and quantity for the industry as a whole. It has to apply to each firm individually, too.

An equilibrium in an oligopoly starts with the same idea as in perfect competition or monopoly: The market clears. However, it adds a requirement that no company wants to change its behavior (its own price or quantity) once it knows what other companies are doing. In other words, each company must be doing as well as it can *conditional* on what other companies are doing. Oligopoly equilibrium has to be stable not only in equating the total quantities supplied and demanded, but also must be stable among the individual producers in the market.

Nash equilibrium An equilibrium in which each firm is doing its best conditional on the actions taken by other firms.

An *equilibrium in which each firm is doing its best conditional on the actions taken by other firms* is called a **Nash equilibrium**. It is named after Nobel laureate John Nash (who was also the subject of the award-winning book and movie titled *A Beautiful Mind*). The Nash equilibrium concept is even more central in the next chapter when we study game theory, which explores the strategic interaction among firms. For our purposes in this chapter, though, the following example will help clarify what is and what is not a Nash equilibrium in an oligopoly.

Application: An Example of Nash Equilibrium: Marketing Movies

Major superhero action movies like Disney's *Black Panther* or Warner Brothers' *Wonder Woman* are amazingly expensive to make. By the time the studios have paid for the CGI, the actors, and everything else, they're looking at a bill of around $180 million.[1] But on top of these production costs, Disney and Warner Brothers each then had to pay many more millions of dollars on advertising and marketing the films so people would watch their movies.

Let's suppose Disney and Warner Brothers are the only two movie companies that make superhero feature films, and that their advertising influences people's choices of what movie to see. Advertising doesn't increase the overall number of movies people watch, just which movie they do.

Now think of both studios planning to release the next installments in these series, *Black Panther* 2 and *Wonder Woman* 2, on the same summer weekend. Furthermore, let's assume that the cost of production is still $180 million and the cost of advertising is

[1] Industry reports estimate *Black Panther* cost $200 million to make, whereas *Wonder Woman* cost $150 million.

$70 million. If both studios advertise and compete with each other, their marketing efforts will cancel out. As a result, the two will split the market, and each will bring in, let's say, $500 million of revenue. Subtracting the $180 million production cost and the $70 million advertising cost, that leaves $250 million of profit to each studio.

If, on the other hand, the studios could somehow agree not to advertise at all, they would again split the market, but this time each would save the $70 million in advertising costs. In this case, the studio profits would be greater at $320 million each.

Disney and Warner Brothers would prefer the second, higher-profit outcome. The problem is that, due to the nature of advertising's influence on moviegoers, if only one studio advertises and the other doesn't, then the studio that advertises will get a larger share of the audience and the other one will be left with less. Suppose, for example, that the studio engaged in advertising would earn $800 million of revenue, and the other would earn only $100 million. The firm advertising its film therefore earns a profit of $550 million ($800 million of revenue minus the $180 million production cost and the $70 million in advertising). The other studio, the one that doesn't advertise, *loses* $80 million ($100 million of revenue minus the $180 million production cost).

Table 11.1 lays out these scenarios. The table's four cells correspond to the four possible profit outcomes if each firm pursues the strategy described at the top of each column and the start of each row: Both firms advertise (upper left), neither firm advertises (lower right), Warner Brothers advertises and Disney doesn't (upper right), or vice versa (lower left). Profit is measured in millions of dollars. The number before the comma in each cell is Warner Brothers' profit if both studios take the actions that correspond to that cell. The number after the comma is Disney's profit.

Look at the table and contemplate where equilibrium might occur in this industry. At first glance, you may expect that, because they could maximize their joint profits by agreeing not to advertise, the studios should just collaborate and earn $320 million each. This is not a Nash equilibrium, however. Here's why: Suppose a studio applied this reasoning and actually held off from advertising because it believed its profit would be higher. Once the first studio decides not to advertise, however, the other studio has a strong incentive to advertise. The other studio can now earn far more profit by advertising than by going along with the don't advertise plan. Recall that the Nash equilibrium means both companies are doing the best they can, *given what the other is doing*. Because one studio can earn a higher profit by advertising when the other doesn't, agreeing not to advertise is not a Nash equilibrium.

To make this scenario concrete, let's say Disney has decided not to advertise. Looking at the profits in Table 11.1, you can see that if Warner Brothers goes along, it will earn $320 million in profit. If it instead abandons the agreement and chooses to advertise, however, it will earn $550 million. Clearly, Warner Brothers will do the latter. You can also see in the table that it works the other way, too: If Warner Brothers chooses not to advertise, Disney does better by advertising (also earning $550 million instead of $320 million).

Table 11.1 An Advertising Game*

		DISNEY	
		Advertise	Don't Advertise
WARNER BROTHERS	Advertise	250, 250	550, −80
	Don't Advertise	−80, 550	320, 320

*Outcomes are measured in millions of dollars of profit.

Therefore, any agreement to hold off from advertising is not stable because both parties have an incentive to cheat on it. Even if one of them sticks to the agreement, the other will earn more profit by sabotaging it. Because each studio will earn higher profit by advertising when the other does not, an outcome in which neither studio advertises cannot be a Nash equilibrium. Agreeing not to advertise is not a Nash equilibrium.

Our analysis so far has established that if one studio *doesn't* advertise, the other studio wants to advertise. What is a studio's optimal action if the other studio *does* advertise? The answer may be found in Table 11.1. If Disney advertises, Warner Brothers earns $250 million by advertising and loses $80 million by not advertising. A similar situation holds for Disney's best response to Warner Brothers. Therefore, advertising is each studio's best response to the other's choice to advertise.

This means that choosing to advertise is a studio's best course of action regardless of whether the other studio advertises or not. Because this is true for both Disney *and* Warner Brothers, the only Nash equilibrium in this case is for both studios to advertise. It is stable because each company is doing the best it can *given what the other is doing*.

prisoner's dilemma A situation in which the Nash equilibrium is an outcome that is somehow worse for all involved than another (unstable) outcome.

Notice that this is true even though it means the studios' profits in the Nash equilibrium will be $250 million each—lower than the $320 million each would earn if they could both hold off from advertising. *A situation in which the Nash equilibrium is an outcome that is somehow worse for all involved than another (unstable) outcome* is known as a **prisoner's dilemma** in game theory. We will look at such situations in more detail in the next few sections and in the next chapter. ■

11.2 Oligopoly with Identical Goods: Collusion and Cartels

Model Assumptions Collusion and Cartels

- Firms make identical products.
- Industry firms agree to coordinate their quantity and pricing decisions, and no firm deviates from the agreement even if breaking it is in the firm's best self-interest.

collusion Economic behavior in which all the firms in an oligopoly coordinate their production and pricing decisions to collectively act as a monopoly to gain monopoly profits to be split among themselves.

cartel The organization formed when firms collude.

In the next several sections, we examine several different models of imperfect competition. They give very different answers about the way in which firms make decisions, so it's important to know which model is the right one to use. A box at the start of each section lists the conditions an industry must meet for that model to apply. In the first model, *all the firms in an oligopoly coordinate their production and pricing decisions to collectively act as a monopoly to gain monopoly profits to be split among themselves.* This economic behavior is known as **collusion**. The *organization formed when firms collude* is often called a **cartel**.[2]

If the companies in an oligopoly can successfully collude, figuring out the oligopoly equilibrium is easy. The firms act collectively as a single monopolist would, and the industry equilibrium is the monopoly equilibrium (output is the level for which $MR = MC$ and the price is determined by the demand curve, as we saw in Chapter 9).[3] Don't try this at home, though. Cartels and collusion violate the law in most every country of the world, and in the United States, it is a criminal offense that has landed many executives in prison. We discussed in Chapter 9 that governments enforce antitrust laws because of monopolies' potential to harm consumers. That explains why collusion has to be done in secret. Interestingly, the secrecy itself makes it more difficult for cartels to maintain a stable equilibrium.

[2] Sometimes the term "cartel" is reserved for a joint monopoly behavior when the firms involved have a public agreement, whereas "collusion" is used to refer to this behavior when it is done in secret. Both describe the same economic behavior, however.

[3] While determining the market equilibrium price, total quantity, and total profits in a cartel is easy, it's not always easy (either for economists studying cartels or the firms in the cartels themselves) to figure out how the cartel's quantity and profits will be divided among its members. We discuss this problem later in the section.

FREAKONOMICS

Apple Always Wins, or Does It?

In the months leading up to the launch of the iPad, Apple was also preparing to open the iBooks Store, which would allow users to buy and read books on mobile devices. But Apple faced a dilemma. It would be difficult for the iBooks Store to succeed unless it could match the $9.99 price point of Amazon's Kindle, its major competitor. This price, thought by some to be less than marginal cost, had served Amazon well in building a customer base for e-books, but the price made it challenging for a new entrant like Apple to successfully compete.

Amazon's low price point was a concern for book publishers as well. They feared that cheap e-books would hurt sales of their more expensive print copies and, over the longer term, influence the public's expectations regarding book prices. Amazon's growing market power also posed the threat that Amazon might start directly competing with publishers.

The major publishers had already begun engaging in talks, meeting in private dining rooms in New York City restaurants to discuss ways to force Amazon to price above $9.99. Before one of their meetings, David Young, then chairman and CEO of Hachette Book Group, told a fellow publisher, "I hate [Amazon's] bullying behavior and will be happy to support a strategy that restricts their plans for world domination." The publishers started implementing their strategy. They coordinated on raising the wholesale price of e-books and introducing "windowing," which delayed e-book versions of new releases to protect hardcover sales. After one particular correspondence regarding these tactics, Young advised another publishing executive that "it would be prudent for you to double delete this from your email files."

Eddy Cue, Apple's Senior Vice President of Internet Software and Services, learned of the publishers' discontent with Amazon's price point, and he began requesting meetings with the major publishers to encourage them to join the iBooks Store. Cue assured the publishers that Apple would price books higher than Amazon. After the first of his meetings, he reported to Apple's then-CEO Steve Jobs that the publishers were "ecstatic" about the prospect of Apple's entry into the industry.

Over the next several weeks, the publishers and Apple crafted a contract that featured a market-wide transition from a wholesale model to an agency model (whereby the publisher rather than the retailer sets the retail price) and a clause that guaranteed Apple would hold the lowest prices on the market. These changes would effectively drive e-book prices upward while securing Apple's competitive position in the market.

Amazon responded by also moving to an agency model in the following months. Soon thereafter, e-book retail prices increased an average of 14.2% per unit and 42.7% for *New York Times* bestsellers. Publishers sold fewer e-books through Amazon, as might have been expected, though estimates of the magnitude of this decrease varied. On the whole, however, the collusion between Apple and the publishers appeared to be a massive success. Not only had it raised prices, it seriously eroded Amazon's monopoly. Before the price-fixing scheme, Amazon held about 90% of the e-book market. A year-and-a-half after the kickoff of the iBooks Store, when the scheme was in full effect, this share was closer to 60%.

Once again, it seemed as if the old Apple magic had worked: The company revolutionized yet another market by entering it. Maybe that would have been the case had things stayed the way they were. But any cheers of victory among Apple and the publishers quickly faded when the U.S. government sued them for collusion. The presiding judge announced her decision just over a year later: Apple was guilty. The words the judge used were scathing. "To adopt Apple's theory, a fact-finder would be confronted with the herculean task of explaining away reams of documents and blinking at the obvious." Several states and private plaintiffs sought more than $800 million in damages. Apple managed to soften the blow by appealing and reaching a $450 million settlement in case its appeal fell through, but it lost the most important gain the scheme had rendered: its ability to control e-book prices.

The Instability of Collusion and Cartels

The firms in an oligopoly would love to collude. They could earn more profit. Adam Smith, the eighteenth-century philosophy professor and one of the fathers of the discipline of economics, recognized this. He wrote in *The Wealth of Nations*, "People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices."

But colluding is harder than it looks. Each member of a cartel has an incentive not to go along. Although firms in a market might be able to come to some initial agreement over a bargaining table, collusion turns out to be unstable—not an equilibrium.

Think about an industry in which there are two firms, Firm A and Firm B, that want to collude. To keep things simple, say both firms have the same constant marginal cost c. If the two firms act collectively as a monopolist, we can follow the monopoly method from Chapter 9 to determine the market equilibrium. Each firm will operate where marginal revenue equals marginal cost. It's not stable, though, because each will want to increase its output at the other's expense.

Suppose the inverse market demand curve for their product is $P = a - bQ$, where P is the price per unit and Q is the quantity produced. We know from Section 9.2 that the marginal revenue curve corresponding to this linear inverse demand curve is $MR = a - 2bQ$. The firms will produce a quantity that sets their marginal revenue equal to their marginal cost c:

$$MR = MC$$

$$a - 2bQ = c$$

Solving this equation for Q gives $Q = (a - c)/2b$. This is the industry's output when its firms collude to act like a monopolist. If we plug this back into the demand curve equation, we find the market price at this quantity: $P = (a + c)/2$.

This is the industry's *total* production in the collusive monopoly outcome. Any combination of the individual firms' outputs that adds to this total will result in the monopoly price and profit. Of course, the firms have to decide how to split this profit. Because both firms have the same costs, splitting seems reasonable. Each firm produces half of the output, $Q/2 = (a - c)/4b$, and they split the monopoly profit equally. Later in this section, we discuss why collusion is even more unstable when firms have different costs.

Cartel Instability: A Mathematical Analysis To see why collusion is unstable, let's work though an example with specific numbers. Suppose the inverse demand curve is $P = 20 - Q$ and $MC = \$4$. Setting $MR = MC$, as above, the total industry output in a collusive equilibrium will be $Q = 8$ units, and the monopoly price will be $P = \$12$. Assuming that Firms A and B split production evenly, each makes 4 units under collusion. This outcome is shown in **Figure 11.1**.

Collusion and cartels fall apart for the same reason that Disney and Warner Brothers can't agree to stop advertising in our earlier example. It's in each company's interest to expand its output once it knows the other company is restricting output. Each company has the incentive to cheat on the collusive agreement. In other words, collusion is *not a Nash equilibrium.*

Think about either company's output choice in our example. Will Firm A want to stick with the output of 4 (half the monopoly output of 8) if Firm B decides to produce 4? If Firm A decides to increase its output to 5 instead of 4, then the total quantity produced in the industry would increase to 9. This higher output level lowers the price from \$12 to \$11 (the demand curve in Figure 11.1 slopes down so price falls when quantity rises).

Once Firm A cheats and increases its output, the industry is no longer at the monopoly quantity and price level, and total industry profit will fall because of overproduction.

Figure 11.1 Cartel Instability

A cartel would like to operate as a monopoly, restricting output to 8 (where $MR = MC$) and selling each unit at a price of \$12 for an industry profit of $(\$12 - \$4) \times 8 = \$64$. If production and profit are shared equally between two firms, each firm earns a profit of $(\$12 - \$4) \times 4 = \$32$. However, Firm A may earn a greater profit by cheating on the agreement and producing another unit, which raises total output in the market and lowers price to \$11 per unit. At this price and output, Firm A earns a profit of $(\$11 - \$4) \times 5 = \$35$. So, each firm has an individual incentive to cheat, and collusion is not stable.

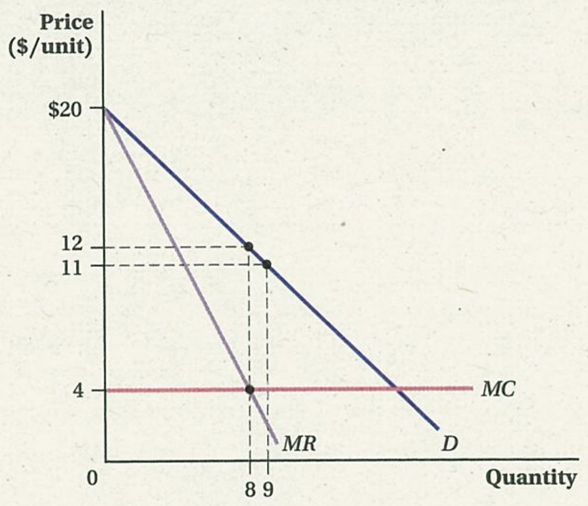

Total profit drops from $Q \times (P - c) = 8 \times (12 - 4) = \$64$ at the monopoly/cartel level to $9 \times (11 - 4) = \$63$ after Firm A increases its output on the sly.

Although the profit of the industry as a whole falls, Firm A, the company that violates the agreement, earns more than with collusion. Its profit under collusion was \$32 (half of the monopoly profits of \$64). But now its profit is $5 \times (11 - 4) = \$35$. The extra sales from increasing production more than make up for the lower prices caused by the increase in production.

A Nash equilibrium requires each firm to be doing the best it can given what the other firm is doing. This example shows that one firm can do better by violating the collusive agreement if the other firm continues to uphold it, so collusion is not a Nash equilibrium. In fact, the cheating firm can do better still by producing more than 5 units. If one firm sticks to the collusive agreement and makes 4 units, the profit-maximizing output for the other firm is 6 units. At this quantity, the price is $20 - (4 + 6) = \$10$, and the cheating firm's profit is $6 \times (10 - 4) = \$36$. (Test this out for yourself. Notice that the cheating firm's profit only starts to fall if it produces 7 or more units.) Both firms face the same incentives to cheat. That's why collusion is difficult.

Increasing the Number of Firms in the Cartel This example was for a two-firm cartel. If more firms are involved, the difficulties of holding an agreement together get even worse. Consider the above example, but now with four firms instead of two. In a collusive agreement, each firm would make 2 units (one-fourth of the total quantity of 8) and earn \$16 (one-fourth of the monopoly profit of \$64). If three of the firms abide by the cartel and make 2 units, but the fourth decides to cheat and make 3, price would again fall to \$11. The company that cheated on the agreement would earn a profit of $3 \times (11 - 4) = \$21$. The \$5 increase in profit from cheating here is even larger than the \$3 increase when there were only two firms. And the cheater would want to expand production even more. If the other firms remain at the collusive output (2 units per firm), the cheater will increase its profit by producing both a 4th and 5th unit. However, producing a 6th unit would reduce the cheater's profit. (To see this, consider the prices that would occur when total output is 8, 9, and 10 units and calculate the cheating firm's profits at those prices.) Because profit falls when the cheating firm produces 6 units, its profit-maximizing output, conditional on the others' production, will be 5 units.

Besides raising the value to cheating, having more firms in a cartel also reduces the damages suffered by any firm that continues to abide by the collusive agreement. This is because the profit losses caused by the cheating will be spread across more firms. This factor further contributes to the difficulty of maintaining collusion when more firms are involved.

This cheating problem is familiar to cartels everywhere. Each firm in the cartel wants every *other* firm to collude, thereby raising the market price, while it steals away business from everyone else by producing more output, thus lowering the market price. Because every firm in a cartel has this same incentive to cheat, it's difficult to persuade anyone to collude in the first place.

figure it out 11.1

Suppose that Squeaky Clean and Biobase are the only two producers of chlorine for swimming pools. The inverse market demand for chlorine is $P = 32 - 2Q$, where Q is measured in tons and P is dollars per ton. Assume that chlorine can be produced by either firm at a constant marginal cost of $16 per ton and there are no fixed costs.

a. If the two firms collude and act like a monopoly, agreeing to evenly split the market, how much will each firm produce and what will the price of a ton of chlorine be? How much profit will each firm earn?

b. Does Squeaky Clean have an incentive to cheat on this agreement by producing an additional ton of chlorine? Explain.

c. Does Squeaky Clean's decision to cheat affect Biobase's profit? Explain.

d. Suppose that both firms agree to each produce 1 ton more than they were producing in part (a). How much profit will each firm earn? Does Squeaky Clean now have an incentive to cheat on this agreement by producing another ton of chlorine? Explain.

Solution:

a. If the firms agree to act like a monopoly, they will set $MR = MC$ to solve for the profit-maximizing output:

$$MR = MC$$
$$32 - 4Q = 16$$
$$4Q = 16$$
$$Q = 4$$

and each firm will produce 2 tons. To find the price, we substitute the market quantity ($Q = 4$) into the inverse demand equation:

$$P = 32 - 2Q = 32 - 2(4) = \$24 \text{ per ton}$$

Each firm will earn a profit of $(\$24 - \$16) \times 2 = \$16$.

b. If Squeaky Clean cheats and produces 3 tons, Q rises to 5 and price falls to $22. Squeaky Clean's profit will be equal to $(\$22 - \$16) \times 3 = \$18$. Therefore, Squeaky Clean does have an incentive to cheat on the agreement because its profit would rise.

c. If Squeaky Clean cheats, the price in the market falls to $22. This reduces Biobase's profit, which is now $(\$22 - \$16) \times 2 = \$12$.

d. If both firms agree to limit production to 3 tons, $Q = 6$ and $P = \$20$. Therefore, each firm earns a profit of $(\$20 - \$16) \times 3 = \$12$. If Squeaky Clean tries to produce 4 tons of chlorine, Q rises to 7 and P falls to $18. Therefore, Squeaky Clean's profit will be $(\$18 - \$16) \times 4 = \$8$. Thus, Squeaky Clean does not have an incentive to cheat on this agreement because its profit would fall.

What Makes Collusion Easier?

Although collusion isn't an especially stable form of oligopoly, there are some conditions that make it more likely to succeed. The first circumstance an aspiring cartel needs is a way to detect and punish cheaters. We just saw that each company in a cartel has the private incentive to produce more output (or charge a lower price) than at the collusive level. If the other firms in a cartel have no way of knowing when a member cheats—and no form of punishment to inflict when they discover someone is cheating—little chance exists that an agreement will hold. That's why collusion is more likely to work when firms can closely observe the transaction quantities and prices of other firms. Such transparency limits the ability of potential cheaters to cut secret, lower-price deals with customers. If a firm cheats, the cartel needs to have some way to enforce the agreement or punish the cheater. Because collusion is generally illegal, the cartel can't really take the cheaters to court, but it might be able to take other actions that reduce the profits of firms that stray from the agreement, such as shutting them out of a share of future cartel profits.

Second, a cartel may find it easier to succeed if there is not much variation in marginal costs across its members. To maximize profit, a monopoly (or a cartel trying to act like a monopoly) wants to use the lowest-cost production method. This desire complicates any scheme to share the monopoly profit among the cartel members and leaves open more opportunities to cheat. Within OPEC, for example, if Saudi Arabia can pump its oil out of the ground for about $4 a barrel, while in Nigeria it costs about $20 per barrel, OPEC would need to explain to Nigeria that the most efficient production strategy would be to sell only Saudi Arabian oil and none from Nigeria.

Third, cartels are more stable when firms take the long view and care more about the future. Think of staying in a cartel (i.e., choosing not to cheat on a collusive agreement) as trading off a short-term opportunity cost to earn a long-term gain. The short-term opportunity cost is giving up the higher profit a firm could obtain by cheating on the agreement. The long-term benefit is that, if the cartel can avoid dissolving into competition, it stands to make monopoly profits. The more the firm values those future monopoly profits relative to the quick hit of additional profit from cheating, the more it will abide by the collusive agreement. Impatient companies, like those in danger of bankruptcy and therefore in desperate need of profit today, are more likely to cheat.

Application: Cartel Bots?

Most people's idea of a cartel is a group of businesspeople who get together to plan out the manipulation of a market. These meetings are often set in either a literal or, more likely these days, figurative "smoke-filled room." Either way, people are involved.

The advent of artificial intelligence (AI) and other algorithmic business practices, however, has recently raised another concern. What if collusion isn't the result of people actually deciding to act in a particular way, but rather because pricing algorithms built into software "decide" to raise prices? Suppose competing companies wrote pricing software that effectively operationalized the command, "If our competitors raise prices, raise our price." It isn't difficult to imagine that the market could fall into a cartel-like outcome without any humans agreeing with one another to actually take the step of increasing prices. How could antitrust authorities prosecute a case where people never made a collusive deal? You can't put an algorithm in jail.

Perhaps we could make even *using* such cartel-spurring algorithms illegal and subject the installers to penalties. But a while a strategy to "price high if they price high" is transparent and easy to track, most pricing algorithms are much more complex. They could plausibly lead to cartel-like outcomes in ways that no one, possibly even their creators, might have imagined. Several AI bots given no other guidance than to set prices to maximize profits might quickly discover for themselves that the best way to do this is to collude. (They would also be fairly good at quickly detecting any cheating behavior, raising the stability of the collusive outcome.)

Policymakers and economists have yet to settle on a recommended course of action. But technological trends make it likely that these sorts of situations will occur with increasing frequency in the future. ■

11.3 Oligopoly with Identical Goods: Bertrand Competition

Model Assumptions Bertrand Competition with Identical Goods

- Firms sell identical products.
- The firms compete by choosing the price at which they sell their products.
- The firms set their prices simultaneously.

In the previous section, we learned that the collusion/cartel model of oligopoly in which firms behave like a monopoly is unlikely to hold in reality because coordination is not an equilibrium and the agreement will likely break down. We need a model in which firms compete against one another. The first such model is as simple as it gets: Firms sell the same product, and consumers compare prices and buy the product with the lowest price. This *oligopoly model in which each firm chooses the price of its product* is called **Bertrand competition**, after Joseph Bertrand, the nineteenth-century French mathematician and economist who first wrote about it. When firms are selling identical products, as we're assuming here, Bertrand oligopoly has a particularly simple equilibrium: $P = MC$, just like perfect competition. In later sections of this chapter, we see how circumstances change when firms sell products that are not identical.

Bertrand competition Oligopoly model in which each firm chooses the price of its product.

Setting Up the Bertrand Model

To set up this model, let's suppose a market with only two companies in it exists. They sell the same product and have the same marginal cost. For example, suppose there are only two stores in a city, a Walmart and a Target, and these stores are located next to each other. They both sell the latest Nintendo Switch and each firm's marginal cost is $300 per console. This includes the wholesale price the firm has to pay Nintendo as well as miscellaneous selling costs, such as stocking the consoles on shelves, checking customers out, and so on.

We need one further assumption: Consumers don't view either store differently in terms of service, atmosphere, or the like. If consumers value these characteristics separately from the consoles, then in a way the products would no longer be identical and we would need to model the firms' behavior using the model of differentiated products discussed later in the chapter.

With only two companies in a market, it might seem as if there would be a lot of market power and high markups over cost. But suppose the customers in this market have a simple demand rule: Buy the console from the store that sells it at the lowest price. If both stores charge the same price, consumers flip a coin to determine where they buy. This rule

means, in effect, that the store charging the lowest price will garner all the demand in the market. If both stores charge the same price, each store gets half the demand.

Suppose the total demand in the market is for Q consoles. Let's denote Walmart's price as P_W and Target's price as P_T. The two stores then face the following demand curves:

Demand for Nintendo Switches at Walmart:

$$\begin{aligned} &Q, && \text{if } P_W < P_T \\ &Q/2, && \text{if } P_W = P_T \\ &0, && \text{if } P_W > P_T \end{aligned}$$

Demand for Nintendo Switches at Target:

$$\begin{aligned} &Q, && \text{if } P_T < P_W \\ &Q/2, && \text{if } P_T = P_W \\ &0, && \text{if } P_T > P_W \end{aligned}$$

Each store chooses its price to maximize its profit, realizing that it will sell the number of units according to the demand curves above. We've assumed the total number of consoles sold, Q, doesn't depend on the price charged. The price only affects which store people buy from. (We could alternatively have allowed Q to depend on the lowest price charged; all the key results discussed below would remain the same.)

Nash Equilibrium of a Bertrand Oligopoly

Remember that in a Nash equilibrium, each firm is doing the best it can given whatever the other firm is doing. So to find the equilibrium of this Bertrand model, let's first think about Target's best response to Walmart's actions. (We could do this in the reverse order if we wanted.) If Target believes Walmart will charge a price P_W for Nintendo Switches, Target will sell nothing if it sets its price above P_W, so we can probably rule that out as a profit-maximizing strategy. Target is left with two options: Match Walmart's price and sell $Q/2$ units, or undercut Walmart and sell Q. Because all it has to do is undercut Walmart by *any* amount, dropping its price just below P_W will only reduce its per-unit margin by a tiny amount, but the store will double its sales because it will take the whole market instead of splitting it.

As an example, suppose $Q = 1{,}000$ and Target thinks Walmart will charge $P_W = \$325$. If Target also charges $P_T = \$325$, it will sell 500 consoles at a profit of $25 each (the $325 price minus the $300 marginal cost). That's a total profit of $12,500. But if Target charges $324.99, it will sell 1,000 Nintendo Switches at a profit of $24.99 each. This is a profit of $24,990—almost double what it was at $175. Target has a strong incentive to undercut Walmart's expected price.

Of course, things are the same from Walmart's perspective: It has the same incentive to undercut whatever price it thinks Target will choose. If it believes Target is going to charge $P_T = \$324.99$ for a Nintendo Switch, Walmart could price its consoles at $324.98 and gain back the entire market. But then Target would have the incentive to undercut *this* expected price, and so on.

This incentive for undercutting would only stop once the price each store expects the other to charge falls to the level of the stores' marginal costs ($300). At that point, cutting prices further would let a store gain the entire market, but that store would be selling every Switch at a loss.

The equilibrium of this Bertrand oligopoly occurs when each store charges a price equal to its marginal cost—$300 in this example. Each obtains half of the market share, and each store earns zero economic profit. The stores would like to charge more, but if either firm raises its price above marginal cost by even the smallest amount, the other firm has a strong incentive to undercut it. And dropping prices below marginal cost would only cause the stores to suffer losses. Thus, the outcome isn't great for the firms, but neither firm can do better by unilaterally changing its price. This is the definition of a Nash equilibrium.

In the identical-good Bertrand oligopoly, one firm cannot increase its profit by raising its price *if* the other firm still charges a price equal to its marginal cost. If the firms could somehow figure out a way to coordinate changes in their actions so that they both raised prices together, they would raise their profits. However, the problem with this strategy, as we saw earlier, is that collusion is unstable. Once the firms are charging prices above marginal cost, a firm can raise its profits by unilaterally changing its action and lowering its price just slightly.

The Bertrand model of oligopoly shows that even with a small number of firms, competition can still be extremely intense under the right conditions. *In fact, the market outcome of Bertrand competition with identical goods is the same as that in a perfectly competitive market: Price equals marginal cost.* This super-competitiveness occurs because either firm can steal the whole market away from the other by dropping price only slightly. The strong incentive to undercut the price leads both firms to drop their prices to marginal cost.

This example had only two firms, but the result would be the same if there were more. The intuition is the same: Every firm's price-cutting motive is so strong that the only equilibrium is for them to all charge a price equal to marginal cost and split the market evenly.[4] The strong assumptions of the Bertrand model with identical products are rare, but some online markets approximate this condition. Where comparing across merchants is really easy, the lowest-priced seller can take the lion's share of the market and these markets often end up with all firms charging the same low price, as the model predicts.

11.4 Oligopoly with Identical Goods: Cournot Competition

Model Assumptions Cournot Competition with Identical Goods

- Firms sell identical products.
- Firms compete by choosing a quantity to produce.
- All goods sell for the same price—the market price, which is determined by the sum of the quantities produced by all the firms in the market.
- Firms choose quantities simultaneously.

When firms sell identical goods, the Bertrand competition model results in the same equilibrium that we find in a perfectly competitive market, where price equals marginal cost. Because consumers care only about the price of the good (the product is identical across firms), each firm faces a demand that is perfectly elastic. Any increase in a firm's price results in it losing all its market share. The demand for the product will go to the firm offering the lowest price.

[4] This again assumes all firms in the market have the same marginal cost. If firms have different marginal costs in an identical-product Bertrand oligopoly, then the equilibrium is for the lowest-cost firm to charge a price just under the *second-lowest* cost in the market and take the whole market.

But firms often face capacity constraints, and that limits how much demand they can fill in the short run. With capacity constraints, if a firm undercuts another's price, it can only steal as many customers as it has available capacity.

In this case, there won't be as much pressure for a firm to respond to price cuts because each firm will not lose all its customers even if it sets its price higher than that of a competitor. In fact, if the capacity of the low-price company is small enough, its competitor may not feel the need to cut prices much at all. This avoids the price-cutting spiral we saw in the Bertrand model. In this situation, the critical issue is for a firm to determine how much capacity it has and thus what quantity it can produce.

Setting Up the Cournot Model

We raise the idea of capacity constraints to motivate another major model, **Cournot competition**, an *oligopoly model in which each firm chooses its production quantity* (named after its first modeler, Augustin Cournot—yet another nineteenth-century French mathematician and economist).

Cournot competition Oligopoly model in which each firm chooses its production quantity.

In Cournot competition, firms produce identical goods and choose a quantity to produce rather than a price at which to sell the good. Individual firms do not control the price of their goods as they do in the Bertrand model. First, all firms in the industry decide how much they will produce; then based on the quantity produced by all firms, the market demand curve determines the price at which all firms' output will sell. In Chapter 9, we learned that when dealing with a monopolist, the price-quantity outcome is the same whether a firm sets the price of its product or the number of units of output it produces. In an oligopoly, however, the market outcome differs depending on whether the firm chooses to set its price or its quantity.

To be more specific, let's say there are two firms in a Cournot oligopoly, Firm 1 and Firm 2. (There can be more; we keep it at two for the sake of simplicity.) Each has a constant marginal cost of c, and both firms independently and simultaneously choose their production quantities q_1 and q_2. The good's inverse demand curve is

$$P = a - bQ$$

where Q is the *total* quantity produced in the market: $Q = q_1 + q_2$.

Firm 1's profit π_1 is the quantity q_1 it produces times the difference between the market price P and its production costs c, or

$$\pi_1 = q_1 \times (P - c)$$

Substituting the inverse demand equation for P, we find that

$$\pi_1 = q_1 \times [a - b(q_1 + q_2) - c]$$

Similarly, Firm 2's profits are given by the equation

$$\pi_2 = q_2 \times [a - b(q_1 + q_2) - c]$$

These two profit equations make clear that the firms in this oligopoly strategically interact. Firm 1's profit is not just a function of its own quantity choice q_1, but also of its competitor's quantity q_2. Likewise, Firm 2's profit is affected by Firm 1's output choice. The logic is that each firm's production choice, through its influence on the market price P, affects the other firm's profit.

An example of an industry that is like the Cournot model is the crude oil industry. Crude oil is a commodity; consumers are indifferent about oil from different sources. The price of oil is set on a worldwide market, and it depends on the total amount of oil supplied at a given time. Therefore, it's realistic to assume that oil producers, even those such as

Saudi Arabia or Iran with large oil reserves, do not choose the price of their outputs. They just choose how much to produce. Oil traders observe these production decisions for all oil producers, and they bid oil's market price up or down depending on how the total quantity produced (the market supply) compares to current demand. This price-setting process derives from the demand curve that connects total output to a market price.

Equilibrium in a Cournot Oligopoly

Finding the equilibrium for a Cournot oligopoly will be easier to follow using an example. Suppose for simplicity that only two countries pump oil, Saudi Arabia and Iran. Both have a marginal cost of production of \$20 per barrel. Also assume that the inverse demand curve for oil is $P = 200 - 3Q$, where P is in dollars per barrel and Q is in millions of barrels per day.

Finding the equilibrium for the Cournot model is similar to doing so for a monopoly, but with the one change noted above: The market quantity Q is the sum of the quantities produced in Saudi Arabia q_{SA} and Iran q_I, rather than just the monopolist's output: $Q = q_{SA} + q_I$. After recognizing this difference, we follow the same steps used to solve for a monopoly's profit-maximizing output. That is, we find each country's marginal revenue curve and then the quantity at which marginal revenue equals marginal cost.

Let's examine Saudi Arabia's profit maximization first. As we learned in Section 9.2, we can more easily find a firm's marginal revenue curve by starting with its inverse demand curve. Therefore, we start by writing the inverse demand curve equation in terms of the quantity choices of each country:

$$P = 200 - 3Q = 200 - 3(q_{SA} + q_I) = 200 - 3q_{SA} - 3q_I$$

Because the slope of the marginal revenue curve is twice the slope of the inverse demand function, Saudi Arabia's marginal revenue curve is[5]

$$MR_{SA} = 200 - 6q_{SA} - 3q_I$$

Saudi Arabia maximizes profit when it produces the quantity at which its marginal revenue equals its marginal cost:

$$200 - 6q_{SA} - 3q_I = 20$$

We can solve this equation for Saudi Arabia's profit-maximizing output:

$$q_{SA} = 30 - 0.5q_I$$

This outcome differs from the monopoly outcome: If Saudi Arabia were a monopoly, setting marginal revenue equal to marginal cost would result in a single quantity Q because its quantity supplied q_{SA} would be the market quantity supplied Q. In this example, however, Saudi Arabia's profit-maximizing output depends on the competitor's output q_I. Similarly, Iran's profit-maximizing q_I depends on q_{SA} because it faces the same market demand curve and has the same marginal cost:

$$q_I = 30 - 0.5q_{SA}$$

This result shows that one country's output choice effectively decreases the demand for the other country's output. That is, the demand curve for one country's output is shifted

[5] The inverse demand curve for Saudi Arabia, in a diagram with quantity, q_{SA}, on the horizontal axis and price on the vertical axis, has a slope of $\Delta P/\Delta q_{SA} = 3$. This means that only the coefficient on q_{SA} is used to determine the slope of the marginal revenue curve. The slope of the marginal revenue curve is $\Delta MR/\Delta q_{SA} = 6$.

by the amount of the other country's output. If the Saudis expect Iran to produce, say, 10 million barrels per day (bpd), then Saudi Arabia would effectively be facing the demand curve

$$P = 200 - 3q_{SA} - 3q_I = 200 - 3q_{SA} - 3(10) = 170 - 3q_{SA}$$

If it expected Iran to pump out 20 million bpd, Saudi Arabia would face the demand curve

$$P = 200 - 3q_{SA} - 3(20) = 140 - 3q_{SA}$$

In Cournot competition, the demand remaining for a firm's output given competitor firms' production quantities is called the **residual demand curve**. We just derived Saudi Arabia's residual demand curves for two of Iran's different production choices, 10 and 20 million bpd.

residual demand curve In Cournot competition, the demand remaining for a firm's output given competitor firms' production quantities.

In effect, a firm in a Cournot oligopoly acts like a monopolist, but one that faces its residual demand curve rather than the market demand curve. Like any regular demand curve, there's a *marginal revenue curve corresponding to the residual demand curve* (it's called . . . wait for it . . . the **residual marginal revenue curve**). The firm produces the quantity at which its residual marginal revenue equals its marginal cost. That's why Saudi Arabia's optimal quantity is the one that sets $200 - 6q_{SA} - 3q_I = 20$. The left-hand side of this equation is Saudi Arabia's residual marginal revenue (expressed in terms of any expected Iranian output level, q_I).The right-hand side is its marginal cost.

residual marginal revenue curve A marginal revenue curve corresponding to a residual demand curve.

How does the profit-maximizing output of one country change with the other country's expected production? In other words, what role do strategic interactions play in a Cournot oligopoly? This can be seen in **Figure 11.2**, which shows Saudi Arabia's residual demand, residual marginal revenue, and marginal cost curves. (The Iranian case would be the same, just with the two countries' labels switched.) The residual demand RD^1_{SA} and residual marginal revenue RMR^1_{SA} curves correspond to an Iranian output level of 10 million bpd. In other words, if Saudi Arabia expects Iran to produce 10 million bpd, Saudi Arabia's optimal output quantity is 25 million bpd. If it expects Iran to produce 30 million bpd, Saudi Arabia's residual demand and marginal revenue curves shift in to RD^2_{SA} and RMR^2_{SA}, to $P = 110 - 3q_{SA}$ and

Figure 11.2 Optimal Quantity Choices

Saudi Arabia's optimal production quantity is dependent on Iran's production quantity. If the Iranian output level is 10 million bpd, Saudi Arabia's optimal output is 25 million bpd, where its residual marginal revenue curve intersects its marginal cost. If Iranian output increases to 30 million bpd, Saudi Arabia's residual demand and residual marginal revenue curves shift to RD^2_{SA} and RMR^2_{SA}. As a result, Saudi Arabia's optimum output decreases to 15 million bpd.

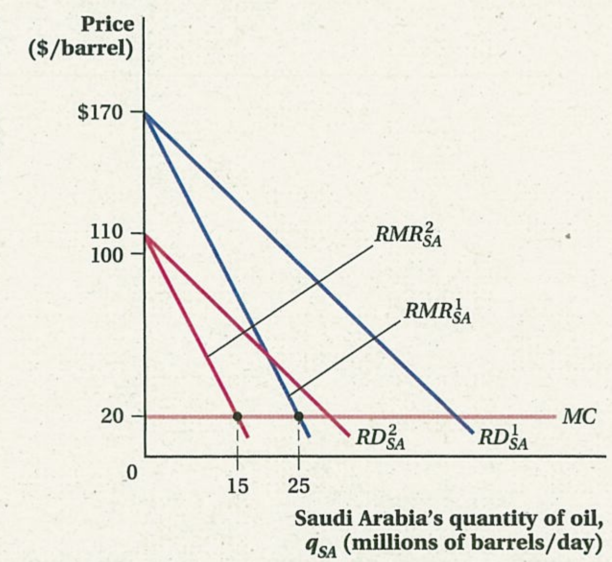

$MR = 110 - 6q_{SA}$, respectively. The Saudis' optimal quantity then falls to 15 million bpd. At Iranian output levels higher than 30 million bpd, Saudi Arabia's residual demand and marginal revenue curves would shift in further and its optimal quantity would fall.

Each competitor's profit-maximizing output depends on the other's output and in an opposite direction: If a firm expects its competitor to produce more, it will reduce its production. Although this kind of interaction seems to create a hopeless chicken-and-egg problem, we can still pin down the specific production quantities for each country if we return to the concept of Nash equilibrium: Each producer does the best it can, taking the other producer's action as given.

For a Cournot oligopoly, the equation for each country's profit-maximizing output is described given the particular output choice of the other country. The equation for q_{SA} gives Saudi Arabia's best response to any production level q_I that Iran might choose. Likewise, the q_I equation gives Iran's best response to any Saudi production decision. In other words, when both equations hold simultaneously, each country is doing the best it can given the other country's action. So, the Nash equilibrium is the combination of outputs that make both equations hold.

Cournot Equilibrium: A Graphical Approach We show this situation graphically in **Figure 11.3**. Saudi Arabia's output is on the vertical axis and Iran's output is on the horizontal axis. The curves illustrated are examples of a **reaction curve**, a *function that relates a firm's best response to its competitor's possible actions. In Cournot competition, this is the firm's best production response to its competitor's possible quantity choices.* Because both reaction curves are downward-sloping, a firm's optimal output falls as the other producer's output rises.

reaction curve A function that relates a firm's best response to its competitor's possible actions. In Cournot competition, this is the firm's best production response to its competitor's possible quantity choices.

Reaction curve *SA* shows Saudi Arabia's best response to any production choice of Iran—it shows the points at which $q_{SA} = 30 - 0.5q_I$. If it expects Iran to produce no oil ($q_I = 0$), for example, the profit-maximizing Saudi response is to produce $q_{SA} = 30$ million bpd. This combination occurs at point *A*. The optimal q_{SA} falls as Iranian production rises. If Iran produces $q_I = 10$, then Saudi Arabia maximizes its profit by producing

Figure 11.3 Reaction Curves and Cournot Equilibrium

A reaction curve represents a firm's optimal production response given its competitor's production quantity. *SA* and *I* are the reaction curves for Saudi Arabia and Iran, respectively. At point *E*, where Iran and Saudi Arabia each produce 20 million bpd ($q_I = q_{SA} = 20$), the market has reached a Nash equilibrium. Here, the two countries are simultaneously producing optimally given the other's actions.

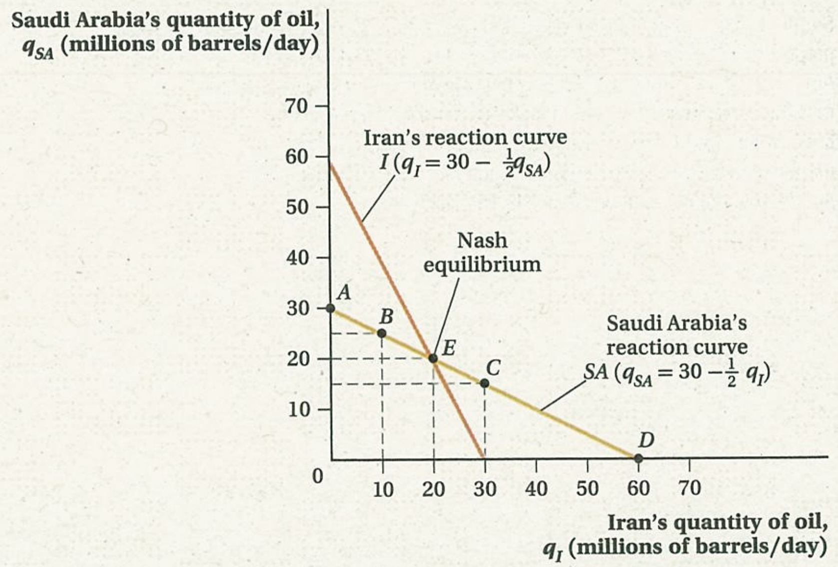

$q_{SA} = 30 - 0.5(10) = 25$ million bpd (point *B*). If $q_I = 30$, then the optimal q_{SA} is 15 million bpd (point *C*). Saudi Arabia's optimal production continues to fall as Iran's production rises until it hits zero at $q_I = 60$ million bpd (point *D*). At any q_I greater than 60 million bpd, the market price is below $20 per barrel (see the demand curve). Because this price is below the marginal cost of production, it wouldn't be profitable for Saudi Arabia to pump any oil if Iran produced 60 million bpd.

Line *I* is the corresponding reaction curve for Iran's profit-maximizing quantity $q_I = 30 - 0.5q_{SA}$. It's essentially the same as *SA*, except with the axes flipped. Just as Saudi Arabia's profit-maximizing output falls with Iran's production choice, Iran's optimal production decreases with expected Saudi production q_{SA}. The optimal q_I is 30 million bpd if $q_{SA} = 0$, and it falls toward 0 as q_{SA} rises toward 60 million bpd.

Each country realizes that its actions affect the desired actions of its competitor, which, in turn, affect its own optimal action, and so on. This back-and-forth strategic interaction is captured in firms' reaction curves and is why the equilibrium is found where reaction curves intersect. The intersection of the two reaction curves at point *E* shows the quantities at which both competitors are simultaneously producing optimally given the other's actions. That is, point *E* is the Nash equilibrium of the Cournot oligopoly—the mutual best response. If one country is producing at point *E*, the other country would only reduce its profits by unilaterally producing at some other point. At this equilibrium, each country produces 20 million bpd, and total output is 40 million bpd.

Cournot Equilibrium: A Mathematical Approach Instead of finding the Cournot equilibrium graphically, we algebraically solve for the output levels that equate the two reaction curves. One way to do this is to substitute one equation into the other to get rid of one quantity variable and solve for the remaining one. For example, if we substitute Iran's reaction curve into Saudi Arabia's reaction curve for q_I, we find

The online appendix finds the equilibrium for Cournot competition using calculus.

$$q_{SA} = 30 - 0.5q_I = 30 - 0.5(30 - 0.5q_{SA})$$
$$= 30 - 15 + 0.25q_{SA}$$
$$0.75q_{SA} = 15$$
$$q_{SA} = 20$$

Thus, the equilibrium output for Saudi Arabia is 20 million bpd. If we substitute this value back into Iran's reaction curve, we find that $q_I = 30 - 0.5q_S = 30 - 0.5(20) = 20$. Iran's optimal production is also 20 million bpd. Equilibrium point *E* in Figure 11.3 has the coordinates (20, 20), and total industry output is 40 million bpd.

The equilibrium price of oil at this point can be found by plugging these production decisions into the inverse market demand curve. Doing so gives $P = 200 - 3(q_{SA} + q_I) = 200 - 3(20 + 20) = \$80$ per barrel. Each country's profit is 20 million bpd × ($80 − $20) = $1,200 million = $1.2 billion per day, so the industry's total profit is $2.4 billion per day.

figure it out 11.2

Go online for interactive, step-by-step help in solving the following problem.

OilPro and GreaseTech are the only two firms that provide oil changes in a local market in a Cournot duopoly (a two-firm oligopoly). The oil changes performed by the two firms are identical, and consumers are indifferent about which firm they will purchase an oil change from. The market inverse demand for the oil changes is $P = 100 - 2Q$, where Q is the total number of oil changes (in thousands per year) produced by the two firms, $q_O + q_G$. OilPro has a marginal cost of $12 per oil change, while GreaseTech has a marginal cost of \$20. Assume that neither firm has any fixed cost.

a. Determine each firm's reaction curve and graph it.

b. How many oil changes will each firm produce in Cournot equilibrium?

c. What will the market price for an oil change be?

d. How much profit does each firm earn?

Solution:

a. Start by substituting $Q = q_O + q_G$ into the market inverse demand curve:

$$P = 100 - 2Q = 100 - 2(q_O + q_G) = 100 - 2q_O - 2q_G$$

From this inverse demand cure, we can derive each firm's marginal revenue curve:

$$MR_O = 100 - 4q_O - 2q_G$$
$$MR_G = 100 - 2q_O - 4q_G$$

Each firm will set its marginal revenue equal to its marginal cost to maximize profit. From this, we can obtain each firm's reaction curve:

$$MR_O = 100 - 4q_O - 2q_G = 12$$
$$4q_O = 88 - 2q_G$$
$$q_O = 22 - 0.5q_G$$

$$MR_G = 100 - 2q_O - 4q_G = 20$$
$$4q_G = 80 - 2q_O$$
$$q_G = 20 - 0.5q_O$$

These reaction curves are shown in the figure in answer b.

b. To solve for equilibrium, we need to substitute one firm's reaction curve into the reaction curve for the other firm:

$$q_O = 22 - 0.5q_G$$
$$q_O = 22 - 0.5(20 - 0.5q_O) = 22 - 10 + 0.25q_O = 12 + 0.25q_O$$
$$0.75q_O = 12$$
$$q_O = 16$$
$$q_G = 20 - 0.5q_O = 20 - 0.5(16) = 20 - 8 = 12$$

Therefore, OilPro produces 16,000 oil changes per year, while GreaseTech produces 12,000.

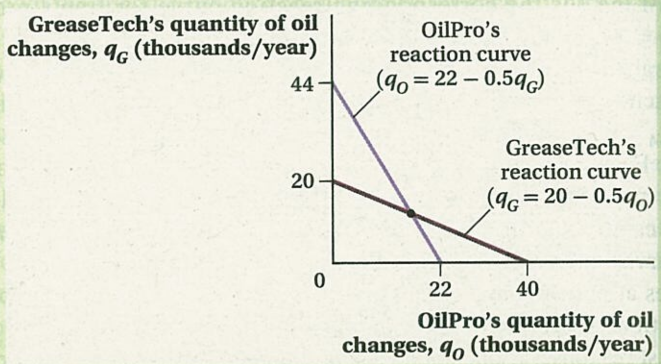

c. We can use the market inverse demand curve to determine the market price:

$$P = 100 - 2Q = 100 - 2(q_O + q_G) = 100 - 2(16 + 12)$$
$$= 100 - 56 = 44$$

The price will be \$44 per oil change.

d. OilPro sells 16,000 oil changes at a price of \$44 for a total revenue $TR = 16,000 \times \$44 = \$704,000$. Total cost $TC = 16,000 \times \$12 = \$192,000$. Therefore, profit for OilPro is $\pi = \$704,000 - \$192,000 = \$512,000$.

GreaseTech sells 12,000 oil changes at a price of \$44 for a total revenue of $TR = 12,000 \times \$44 = \$528,000$. Total cost $TC = 12,000 \times \$20 = \$240,000$. Thus, GreaseTech's profit is $\pi = \$528,000 - \$240,000 = \$288,000$.

Note that the firm with the lower marginal cost produces more output and earns a greater profit.

Comparing Cournot to Collusion and to Bertrand Oligopoly

Let's compare this equilibrium in a Cournot oligopoly ($Q = 40$ million bpd at $P = \$80$) and profit (\$2.4 billion per day) to the outcomes in other oligopoly models we've analyzed. These results are described in **Table 11.2**.

Table 11.2 Comparing Equilibria across Oligopolies

Oligopoly Structure	Total Output (million bpd)	Price (\$ per barrel)	Industry Profit (per day)
Collusion	30	\$110	\$2.7 billion
Bertrand (identical products)	60	20	0
Cournot	40	80	2.4 billion

Collusion Let's first suppose Saudi Arabia and Iran can actually get their acts together and collude to behave like a monopolist. In that case, they would treat their separate production decisions q_I and q_{SA} as a single total output, $Q = q_{SA} + q_I$. Following the normal marginal-revenue-equals-marginal-cost procedure, we would find that $Q = 30$ million bpd. Presumably, the two countries would split this output evenly at 15 million bpd because they have the same marginal costs. This is less than the total Cournot oligopoly production of 40 million bpd that we just derived. Furthermore, because monopoly production is lower, the price is higher, too: Plugging this monopoly quantity into the demand curve, price becomes $P = 200 - 3(30) = \$110$ per barrel. We also know that total industry profit must be higher in the collusive monopoly outcome. In this case, it's 30 million bpd $\times (\$110 - \$20) = \$2.7$ billion per day (or \$1.35 billion for each country). This total is \$300 million per day higher than the Cournot competition outcome. At the collusive monopoly equilibrium, output is lower than at the Cournot equilibrium, and price and profit are higher.

Bertrand Oligopoly with Identical Products Next, let's consider the Nash equilibrium if the two countries competed as in the Bertrand model with identical products. In this case, we know that price will equal marginal cost, so $P = \$20$. Total demand at this price is determined by plugging \$20 into the demand curve: $P = 20 = 200 - 3Q$, or $Q = 60$ million bpd. The two countries would split this demand equally, with each selling 30 million bpd. Because both countries sell at a price equal to their marginal cost, each earns zero profit. At the Bertrand equilibrium, output quantity is higher than at the Cournot equilibrium, price is lower, and there is no profit.

Summary To summarize then in terms of total industry output, the lowest is the collusive monopoly outcome, followed by Cournot, then Bertrand:

$$Q_m < Q_c < Q_b$$

The order is the opposite for prices, with Bertrand prices the lowest and the collusive price the highest:

$$P_b < P_c < P_m$$

Similarly, profit is lowest in the Bertrand case (at zero), highest under collusion, with Cournot in the middle:

$$\pi_b < \pi_c < \pi_m$$

Therefore, the Cournot oligopoly outcome is something between those for a monopoly and Bertrand oligopoly (for which the outcome is equivalent to perfect competition). And, unlike the collusive and Bertrand outcomes, the price and output in the Cournot equilibrium depend on the number of firms in the industry.

What Happens if There Are More Than Two Firms in a Cournot Oligopoly?

These intermediate outcomes are for a market with two firms. If more than two firms are in a Cournot oligopoly, the total quantity, profits, and price remain between the monopoly and perfectly competitive extremes. However, the more firms there are, the closer these outcomes get to the perfectly competitive case, with price equaling marginal cost and economic profits being zero. Having more competitors means that any single firm's supply decision becomes a smaller and smaller part of the total market. Its output choice therefore affects the market price less and less. With a very large number of firms in the market, a producer essentially becomes a price taker. It therefore behaves like a firm in a perfectly competitive industry, producing where the market price equals its marginal cost.

Cournot versus Bertrand: Extensions

The fact that the intensity of competition changes with the number of firms in the market is a nice feature of the Cournot model. This prediction is more in line with many people's view of oligopoly than the Bertrand model's prediction that anything more than a single firm leads to a perfectly competitive outcome. The downside of the Cournot framework is that it's a bit more of a stretch than usual to assume that companies can only compete in their quantity choices and have no ability to charge different prices. Oil seems a very special case.

Economists David Kreps and José Scheinkman examined this scenario in more detail and showed (using math that's a bit beyond our level here) that even if firms actually set their prices instead of quantities, the industry equilibrium will still look like a Cournot model as long as the firms first choose their production capacity before they set prices. The firms are then constrained to produce at or below that capacity level once they make their price decisions.[6]

As an example of a market described by the Cournot model, imagine that a few real estate developers in a college town build student apartments that are identical in quality and size. Once these developers have constructed their apartment buildings, they can charge whatever price the market will bear for the apartments, but their choice of prices will be constrained by the total number of apartments they have built. If, for some reason, the developers want to charge a ridiculously low rent of, say, $50 per month, they would probably not be able to satisfy all the quantity demanded at that low price because they only have a fixed number of apartments to rent. If the developers first choose the number of apartments in their buildings and then sell their fixed capacity at whatever prices they select, Kreps and Scheinkman show that the equilibrium price and quantity (which, as it turns out, will equal the developers' capacity choice in this case) will be much like those in a Cournot oligopoly.

This result means that in industries where there are large costs of investing in capacity so that firms don't change their capacity very often, the Cournot model will probably be a good predictor of market outcomes even if firms choose their prices in the short run. (In the long run, the firms could both change their capacity by building more apartment buildings and alter the prices they choose.)

[6] David M. Kreps and José A. Scheinkman, "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics* 14, no. 2 (1983): 326–337.

11.5 Oligopoly with Identical Goods but with a First-Mover: Stackelberg Competition

Model Assumptions Stackelberg Competition with Identical Goods

1. Firms sell identical products.
2. Firms compete by choosing a quantity to produce.
3. All goods sell for the same price (which is determined by the sum total of quantities produced by all the firms combined).
4. Firms do *not* choose quantities simultaneously. One firm chooses its quantity first. The next firm observes this and then chooses its quantity.

The Cournot model gave us a way to analyze oligopolistic markets that are somewhere between collusion/monopoly and Bertrand/perfect competition. As in most oligopoly models, equilibrium in the Cournot model came from firms rationally thinking through how other firms in the market are likely to behave in response to their production decisions.

Importantly, the Cournot model also relies on another assumption whose implications we didn't consider much, namely, that the firms choose simultaneously. That is, each firm chooses its optimal quantity based on what the firm believes its competitor(s) *might* do. If it expects its competitor(s) to produce some other quantity, its own optimal action changes—that was the logic of the reaction curve.

If you more fully contemplate this situation, though, each company has an incentive to try to choose its output level first and force its competitors to be the ones who have to react. The first firm to make its decision could increase its output and say, "Oh well, I have already made more than Cournot says I am supposed to produce. What are you going to do about it?"

An *oligopoly model in which firms make product decisions sequentially*—first one, then another, then (if there are more than two firms) another, and so on—is referred to as **Stackelberg competition**. (Heinrich Freiherr von Stackelberg was an early-twentieth-century German economist who first analyzed this type of oligopoly.) Because the competitor's reaction curve slopes downward in this model, the competitor, seeing the high quantity the original firm is producing, would want to reduce its output. Therefore, this creates a **first-mover advantage**, an *advantage gained in a Stackelberg competition by the initial firm in setting its production quantity.* The firm that moves first is sometimes called the Stackelberg leader. To see how sequential competition changes things, let's revisit our oil producers, Saudi Arabia and Iran.

Stackelberg competition Oligopoly model in which firms make production decisions sequentially.

first-mover advantage The advantage gained in a Stackelberg competition by the initial firm in setting its production quantity.

The market inverse demand for oil was $P = 200 - 3Q$, and both countries had a constant marginal cost of $20 per barrel. Each firm produced where marginal revenue equaled marginal cost:

$$MR_{SA} = 200 - 6q_{SA} - 3q_I = 20$$
$$MR_I = 200 - 6q_I - 3q_{SA} = 20$$

In Cournot competition, we rearranged these equations to solve for each country's reaction curve:

$$q_{SA} = 30 - 0.5q_I$$
$$q_I = 30 - 0.5q_{SA}$$

We know that these formulas give the best output each country can choose, taking as given the other country's output level. Plugging one reaction curve into the other gave us the Nash equilibrium, in which each country produced 20 million bpd at a market price of $80 per barrel.

Stackelberg Competition and the First-Mover Advantage

Now suppose Saudi Arabia is a Stackelberg leader: It chooses its quantity first. Iran's incentives remain unchanged. It still has the same residual demand and reaction curve, and the reaction curve continues to show Iran's best response to any choice by Saudi Arabia. In Stackelberg competition, however, Iran will know with certainty what Saudi Arabia's production decision is before it makes its own. Iran reacts optimally to any production choice that Saudi Arabia makes by plugging this value for q_{SA} into its reaction function. Importantly, *Saudi Arabia realizes Iran will do this before it makes its first move.*

Because Saudi Arabia knows that Iran's output is going to be a function of whatever Saudi Arabia chooses first, the Saudis want to take that impact into account when they make their *initial* production decision. In this way, Saudi Arabia can take advantage of being the first mover. To do so, it plugs Iran's best response function into its own demand and marginal revenue curve equations. The fact that the Saudi marginal revenue curve changes means that Saudi Arabia will no longer have the same reaction curve it had in the Cournot model. In that model, Saudi Arabia faced the demand curve

$$P = 200 - 3(q_{SA} + q_I)$$

Now that it is a first-mover in a Stackelberg oligopoly, Saudi Arabia's demand is

$$P = 200 - 3q_{SA} - 3q_I = 200 - 3q_{SA} - 3(30 - 0.5q_{SA}) = 200 - 3q_{SA} - 90 + 1.5q_{SA}$$

Do you see what happened? We substituted Iran's reaction function $(q_I = 30 - 0.5q_{SA})$ directly into the Saudi demand curve. We did this because Saudi Arabia recognizes that, by going first, its output choice affects its demand (and therefore its marginal revenue) both directly and indirectly through its effect on Iran's production decision. The direct effect is captured by the term $-3q_{SA}$ in the equation; this effect is the same as in the Cournot model. The indirect effect comes from the impact of Saudi Arabia's output choice on Iran's production response, embodied in the equation's second q_{SA} term $(1.5q_{SA})$.

We can further simplify this demand curve:

$$P = 110 - 1.5q_{SA}$$

We know from Chapter 9 that Saudi Arabia's marginal revenue curve is then $MR_{SA} = 110 - 3q_{SA}$. Setting this equal to marginal cost ($20 per barrel) and solving for q_{SA} give Saudi Arabia's profit-maximizing output in this Stackelberg oligopoly:

$$MR_{SA} = 110 - 3q_{SA} = 20$$
$$3q_{SA} = 90$$
$$q_{SA} = 30$$

As the first-mover, Saudi Arabia finds it optimal to produce 30 million bpd, 10 million more than the Cournot oligopoly output (20 million bpd).

Next, we need to see how Saudi Arabia's decision affects Iran's optimal production level. To do this, we plug Saudi Arabia's output level into Iran's reaction curve:

$$q_I = 30 - 0.5q_{SA} = 30 - 0.5(30) = 15$$

Iran now produces 15 million bpd, rather than 20 as in the Cournot case. By moving first, Saudi Arabia gets the jump on Iran, leaving Iran no choice but to drop its output level from 20 to 15 million bpd.

Therefore, total production is 45 million bpd in the Stackelberg case. This is more than the output produced in the Cournot oligopoly (40 million). And, because production is higher, the market price must be lower under sequential production decisions than under Cournot's simultaneous-decision framework. Specifically, the price is $200 - 3(30 + 15) = \$65$ per barrel (instead of the Cournot equilibrium price of $80).

Profit changes, too. For Saudi Arabia, profit is $30 \times (65 - 20) = \$1{,}350$ million/day, $150 million more than its $1,200 million/day profit in the (simultaneous-move) Cournot oligopoly. Such an outcome shows us the advantage of being the first mover. Iran, on the other hand, makes a profit of only $15 \times (65 - 20) = \$675$ million/day, well below its Cournot profit level of $1,200 million/day. In the next chapter on game theory, we will discuss the role of first-mover advantage in strategic decision making in more detail, but you can already see why firms might want to enter a market early.

Although it's somewhat abstract, the idea of Stackelberg competition in which one firm moves first and obtains an advantage that leads later firms to adjust their strategy and reduce their output is very true to life, as we will see in the next chapter.

figure it out 11.3

Go online for interactive, step-by-step help in solving the following problem.

Consider again the case of the two oil change producers OilPro and GreaseTech from Figure It Out 11.2. Recall that the market inverse demand for the oil changes is $P = 100 - 2Q$, where Q is the total number of oil changes (in thousands per year) produced by the two firms, $q_O + q_G$. OilPro has a marginal cost of $12 per oil change, while GreaseTech has a marginal cost of $20.

a. Suppose this market is a Stackelberg oligopoly and OilPro is the first-mover. How much does each firm produce? What will the market price of an oil change be? How much profit does each firm earn?

b. Now suppose that GreaseTech is the first-mover in this Stackelberg oligopoly. How much will each firm produce, and what will the market price be? How much profit does each firm earn?

Solution:

a. We need to start by reconsidering the demand for OilPro's product. It is going to move first and we assume that it knows from previous experience that GreaseTech's output is a function of OilPro's output. Thus, we need to substitute GreaseTech's reaction curve, from the illustration in prior Figure It Out 11.2, into the market inverse demand curve to solve for the inverse demand for OilPro.

GreaseTech's reaction curve is $q_G = 20 - 0.5q_O$. Substituting this into the inverse market demand curve, we get

$$P = 100 - 2Q = 100 - 2(q_O + q_G) = 100 - 2q_O - 2q_G$$
$$= 100 - 2q_O - 2(20 - 0.5q_O) = 100 - 2q_O - 40 + q_O = 60 - q_O$$

So, the inverse demand curve for OilPro oil changes is $P = 60 - q_O$. This means that the marginal revenue curve for OilPro is

$$MR_O = 60 - 2q_O$$

Setting $MR = MC$ will provide us with OilPro's profit-maximizing output:

$$MR_O = 60 - 2q_O = 12$$
$$2q_O = 48$$
$$q_O = 24$$

Now that we know q_O, we can substitute it into GreaseTech's reaction curve to find q_G:

$$q_G = 20 - 0.5q_O = 20 - 0.5(24) = 20 - 12 = 8$$

OilPro will produce 24,000 oil changes, while GreaseTech will only produce 8,000. Using the inverse market demand, we can determine the market price:

$$P = 100 - 2(q_O + q_G) = 100 - 2(32) = 100 - 64 = \$36$$

OilPro's profit will be $\pi_O = (\$36 - \$12) \times 24{,}000 = \$576{,}000$. GreaseTech's profit will be $\pi_G = (\$36 - \$20) \times 8{,}000 = \$128{,}000$.

b. If GreaseTech is the first-mover, we can use OilPro's reaction curve (from the figure in prior Figure It Out 11.2) to find the inverse market demand for GreaseTech. OilPro's reaction curve is $q_O = 22 - 0.5q_G$. Substituting into the market inverse demand, we get

$$P = 100 - 2q_O - 2q_G = 100 - 2(22 - 0.5q_G) - 2q_G$$
$$= 100 - 44 + q_G - 2q_G$$
$$= 56 - q_G$$

This is the inverse demand for GreaseTech's oil changes. Its marginal revenue is therefore

$$MR_G = 56 - 2q_G$$

Setting $MR = MC$, we can see that

$$MR_G = 56 - 2q_G = 20$$
$$2q_G = 36$$
$$q_G = 18$$

To find OilPro's output, we substitute q_G into OilPro's reaction curve:

$$q_O = 22 - 0.5q_G = 22 - 0.5(18) = 22 - 9 = 13$$

So, when GreaseTech is the first-mover, OilPro only produces 13,000 oil changes, while GreaseTech produces 18,000. We can determine the price using the inverse market demand:

$$P = 100 - 2(q_O + q_G) = 100 - 2(31) = \$38$$

GreaseTech's profit will be $\pi_G = (\$38 - \$20) \times 18{,}000 = \$324{,}000$. OilPro's profit will be $\pi_O = (\$38 - \$12) \times 13{,}000 = \$338{,}000$.

11.6 Oligopoly with Differentiated Goods: Bertrand Competition

Model Assumptions Bertrand Competition with Differentiated Goods

1. Firms do *not* sell identical products. They sell differentiated products, meaning consumers do not view them as perfect substitutes.
2. Each firm chooses the price at which it sells its product.
3. Firms set prices simultaneously.

Every model of imperfect competition that we've looked at so far—collusion, Bertrand, Cournot, and Stackelberg—has assumed that the industry's producers all sell the same product. Often, however, a more realistic description of an industry is a set of firms that make similar but not identical products. When consumers buy cars, breakfast cereals, or even light bulbs, they must choose between competing versions and brands, each with its own unique features, produced by a small number of companies. A *market in which multiple varieties of a common product type are available* is called a **differentiated product market**.

differentiated product market A market in which multiple varieties of a common product type are available.

It is often possible to treat the differentiated products as interacting in a single market, even when it seems as if each one could be considered its own separate market. The key is to explicitly account for the way consumers are willing to substitute among the products.

To see how a Bertrand oligopoly works with differentiated products, think back to the Bertrand model we studied in Section 11.3. There, two companies (Walmart and Target in our example) competed by setting prices for an identical product (the Nintendo Switch). Now, however, instead of thinking of the firms' products as identical as we did in Section 11.3, we assume that consumers view the products as being somewhat distinct. Even though a game console is the same regardless of where customers buy it, the stores offering such a product may have, for example, different locations and customers might care about travel costs. The specific source of the product distinction isn't important. Regardless of its source, any differentiation counts that helps the stores exert more market power and earn more profit. When products are identical, the incentive to undercut price is so intense that firms compete the market price right down to marginal cost and earn zero economic profit as a result. That is not how it works in the differentiated-product Bertrand model, as we see in the following example.

Equilibrium in a Differentiated-Products Bertrand Market

Suppose there are two main manufacturers of snowboards, Burton and K2. Because many snowboarders view the two companies' products as similar but not identical, if either firm cuts its prices, it will gain market share from the other. But because the firms' products aren't *perfect* substitutes, the price-cutting company won't take all the business away from the other company just because it sets its price a bit lower. Some people are still going to prefer the competitor's product, even at a higher price.

This product differentiation means that each firm faces its own demand curve, and each product's price has a different effect on the firm's demand curve. So, Burton's demand curve might be

$$q_B = 900 - 2p_B + p_K$$

As you can see, the quantity of boards Burton sells goes down when it raises the price it charges for its own boards, p_B. On the other hand, Burton's quantity demanded goes up when K2 raises its price, p_K. In this example, we've assumed that Burton's demand is more sensitive to changes in its own price than to changes in K2's price. (For every \$1 change in p_B, there is a 2-unit decrease in quantity demanded; this ratio is 1-to-1—and positive—for changes in p_K.) Our assumption is a realistic one in many markets.

K2 has a demand curve that looks similar, but with the roles of the two firms' prices reversed:

$$q_K = 900 - 2p_K + p_B$$

The responses of each company's quantity demanded to price changes reflect consumers' willingness to substitute across varieties of the industry's product. But this substitution is limited; a firm can't take over the entire market with a 1 cent price cut, as it can in the identical-products Bertrand model.

To determine the equilibrium in a Bertrand oligopoly model with differentiated products, we follow the same steps we used for all the other models: Assume each company sets its price to maximize its profit, taking the prices of its competitors as given. That is, we look for a Nash equilibrium. To keep things simple, we assume that both firms have a marginal cost of zero.[7]

[7] We assume zero marginal cost in this example because the concept of marginal cost is a little different when firms choose prices rather than quantities. Remember that marginal cost is the change in total cost driven by changing output by 1 unit: $MC = \Delta TC/\Delta Q$. As in all other market structures, a firm in a differentiated-product Bertrand oligopoly maximizes profit by setting its marginal revenue equal to its marginal cost. But the expression for marginal revenue in a Bertrand setup is the change in revenue resulting from small *price* changes, or $MR = \Delta TR/\Delta P$, rather than from small *quantity* changes, or $MR = \Delta TR/\Delta Q$. Therefore, the profit-maximizing price in a differentiated-products Bertrand oligopoly sets this price-based marginal revenue equal to a price-based marginal cost: $\Delta TR/\Delta P = \Delta TC/\Delta P$. We could go through some extra algebra to tie the two together, but it's easier for our purposes here to just assume that marginal costs are zero.

Burton's total revenue is

$$TR_B = p_B \times q_B = p_B \times (900 - 2p_B + p_K)$$

Notice that we've written total revenue in terms of Burton's price, rather than its quantity. This is because in a Bertrand oligopoly, Burton chooses the price it will charge rather than how much it will produce. Writing total revenue in price terms lets us derive the marginal revenue curve in price terms as well. Namely, marginal revenue is

$$MR_B = 900 - 4p_B + p_K$$

(Recall that the marginal revenue curve of a linear inverse demand curve is just the inverse demand curve with the quantity coefficient doubled. The same logic holds when marginal revenue is expressed in terms of price.) We can solve for Burton's profit-maximizing price through the usual step of setting this marginal revenue equal to the marginal cost (zero in this case). Doing so and rearranging gives

$$MR_B = 900 - 4p_B + p_K = 0$$
$$4p_B = 900 + p_K$$
$$p_B = 225 + 0.25p_K$$

Notice how this again gives a firm's (Burton's) optimal action as a function of the other firm's action (K2's). In other words, this equation describes Burton's reaction curve. But here, the actions are price choices rather than quantity choices as in the Cournot model.

K2 has a reaction curve, too. It looks similar to, but is a little different than, Burton's because K2's demand curve is slightly different. Repeating the same steps as above, we have

$$MR_K = 900 - 4p_K + p_B = 0$$
$$4p_K = 900 + p_B$$
$$p_K = 225 + 0.25p_B$$

An interesting detail to note about these reaction curves in the Bertrand differentiated-product model is that a firm's optimal price *increases* when its competitor's price increases. If Burton believes K2 will charge a higher price, for instance, Burton wants to raise its price. That is, the reaction curves are upward-sloping. This is the opposite of the quantity reaction curves in the Cournot model (review Figure 11.3). There, a firm's optimal response to a competitor's output change is to do the opposite: If a firm expects its competitor to produce more, then it should produce less.

Differentiated Bertrand Equilibrium: A Graphical Approach Figure 11.4 plots Burton's and K2's reaction curves. The vertical axis shows Burton's optimal profit-maximizing price; the horizontal axis represents K2's optimal profit-maximizing price. The positive slope of Burton's reaction curve indicates that Burton's profit-maximizing price rises when K2 charges more. The positive slope of K2's reaction curve indicates that K2's profit-maximizing price rises when Burton charges more. If Burton expects K2 to charge \$100, then Burton should price its boards at \$250 (point *A*). If instead Burton believes K2 will price at \$200, then it should price at \$275 (point *B*). A K2 price of \$400 will make Burton's optimal response \$325 (point *C*), and so on. K2's reaction curve works the same way.

The point where the two reaction curves cross, *E*, is the Nash equilibrium. There, both firms are doing as well as they can given the other's actions. If either were to decide on its own to change its price, that firm's profit would decline.

The online appendix finds the equilibrium for differentiated Bertrand competition using calculus.

Differentiated Bertrand Equilibrium: A Mathematical Approach We can algebraically solve for this Nash equilibrium as we did in the Cournot model—by finding the point at which the reaction curve equations equal one another. Mechanically, that means we substitute one reaction curve into the other, solve for one firm's optimal price, and then use it to solve for the other firm's optimal price.

Figure 11.4 Nash Equilibrium in a Bertrand Market

This shows Burton and K2's reaction curves. At point *E*, when each company sells 600 snowboards at a market price of \$300 per snowboard, the market is at a Nash equilibrium, and the two companies are producing optimally.

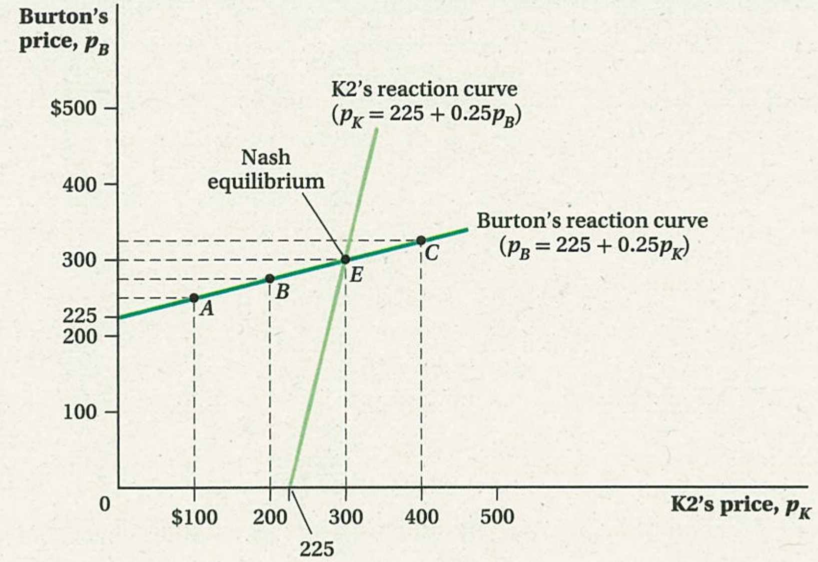

First, we plug K2's reaction curve into Burton's and solve for Burton's equilibrium price:

$$p_B = 225 + 0.25p_K$$
$$p_B = 225 + 0.25 \times (225 + 0.25p_B)$$
$$p_B = 225 + 56.25 + 0.0625p_B$$
$$0.9375p_B = 281.25$$
$$p_B = 300$$

Substituting this price into K2's reaction curve gives its equilibrium price:

$$p_K = 225 + 0.25p_B = 225 + (0.25 \times 300) = 225 + 75 = 300$$

In equilibrium, both firms charge the same price, \$300. This isn't too surprising. After all, the two firms face similar-looking demand curves and have the same (zero) marginal costs. Interestingly, the particular implication of the identical-products Bertrand oligopoly that we looked at in Section 11.3 (that both firms charge the same price in equilibrium) holds here. The difference is that the price no longer equals marginal cost. Instead, equilibrium prices are above marginal cost (\$300 is a lot more than zero).

To figure out the quantity each firm sells, we plug each firm's price into its demand curve equation. Burton's quantity demanded is $q_B = 900 - 2(300) + 300 = 600$ boards. K2 also sells $q_K = 900 - 2(300) + 300 = 600$ boards. Again, the fact that both firms sell the same quantity is not surprising because they have similar demand curves and charge the same price. Total industry production is therefore 1,200 boards, which is two-thirds of what it would be if both firms charged their marginal costs (each firm in that case would make 900 boards, meaning total production of 1,800 boards). In the Bertrand model where the firms produce differentiated products, each firm earns a profit of $600 \times (300 - 0) = \$180,000$.

In this example, both firms had demand curves that were mirror images of each other. If instead the firms had different demand curves, we would go about solving for equilibrium prices, quantities, and profits the same way, but these probably wouldn't be the same for each firm.

figure it out 11.4

Consider our example of the two snowboard manufacturers, Burton and K2. We just determined that at the Nash equilibrium for these two firms, each firm produced 600 snowboards at a price of $300 per board. Now let's suppose that Burton launches a successful advertising campaign to convince snowboarders that its product is superior to K2's, so the demand for Burton snowboards rises to $q_B = 1{,}000 - 1.5p_B + 1.5p_K$, while the demand for K2 boards falls to $q_K = 800 - 2p_K + 0.5p_B$. (For simplicity, assume that the marginal cost is still zero for both firms.)

a. Derive each firm's reaction curve.

b. What happens to each firm's optimal price?

c. What happens to each firm's optimal output?

d. Draw the reaction curves in a diagram and indicate the equilibrium.

Solution:

a. To determine the firms' reaction curves, we first need to solve for each firm's marginal revenue curve:

$$MR_B = 1{,}000 - 3p_B + 1.5p_K$$
$$MR_K = 800 - 4p_K + 0.5p_B$$

By setting each firm's marginal cost equal to marginal revenue, we can find the firm's reaction curve:

$$MR_B = 1{,}000 - 3p_B + 1.5p_K = 0$$
$$3p_B = 1{,}000 + 1.5p_K$$
$$p_B = 333.33 + 0.5p_K$$

$$MR_K = 800 - 4p_K + 0.5p_B = 0$$
$$4p_K = 800 + 0.5p_B$$
$$p_K = 200 + 0.125p_B$$

b. We can solve for the equilibrium by substituting one firm's reaction curve into the other's:

$$p_B = 333.33 + 0.5p_K$$
$$p_B = 333.33 + 0.5(200 + 0.125p_B)$$
$$= 333.33 + 100 + 0.0625p_B$$
$$p_B = 433.33 + 0.0625p_B$$
$$0.9375p_B = 433.33$$
$$p_B = \$462.22$$

We can then substitute p_B back into the reaction function for K2 to get the K2 price:

$$p_K = 200 + 0.125p_B$$
$$= 200 + 0.125(462.22) = 200 + 57.78 = \$257.78$$

Thus, the successful advertising campaign means that Burton can increase its price from the original equilibrium price of $300 (which we determined in our initial analysis of this market) to $462.22, while K2 will have to lower its own price from $300 to $257.78.

c. To find each firm's optimal output, we need to substitute the firms' prices into the inverse demand curves for each firm's product. For Burton,

$$q_B = 1{,}000 - 1.5p_B + 1.5p_K = 1{,}000 - 1.5(462.22) + 1.5(257.78)$$
$$= 1{,}000 - 693.33 + 386.67 = 693.34$$

For K2,

$$q_K = 800 - 2p_K + 0.5p_B = 800 - 2(257.78) + 0.5(462.22)$$
$$= 800 - 515.56 + 231.11 = 515.55$$

Burton now produces more snowboards (693.34 instead of 600), while K2 produces fewer (515.55 instead of 600).

d. The reaction curves are shown in the diagram below:

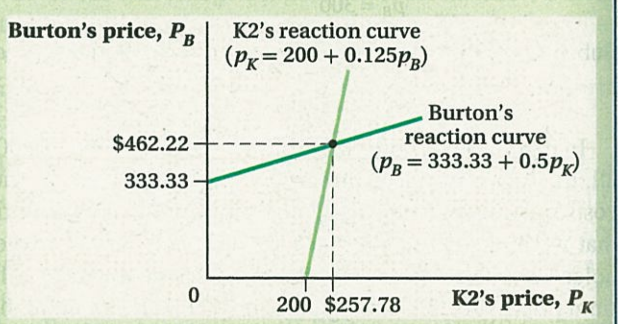

Application: Computer Parts—Differentiation Out of Desperation

Bertrand competition with identical products is extremely intense. In equilibrium, firms set price equal to marginal cost and earn no profit. This is a situation that most firms want to avoid. And, as we just saw, firms can earn profits if their products are differentiated. This gives firms a huge incentive to try to differentiate their products from their competitors' products, even if an outsider to the market might not believe any important differences really exist among them.

This sort of behavior was documented by economists Glenn Ellison and Sara Ellison in an online market for computer chips.[8] In this market, high-tech customers who like to build their own computers shop for CPUs and memory chips using an online price search engine that tracks down and lists the products of various electronic parts retailers.

Ellison and Ellison documented how some computer parts retailers in the market used a little economic know-how to get away with setting their prices above marginal cost. Those firms realized that the key to getting more producer surplus was to differentiate their products, thus shifting the structure of competition from a Bertrand oligopoly with identical products to one with differentiated products.

Just how could these firms differentiate what were otherwise identical computer chips? They couldn't do this the way K2 and Burton can with the snowboards they sell, by varying designs, materials, and so on. So, they turned to slightly more, well, creative methods—methods that Ellison and Ellison categorized as "obfuscation."

Ellison and Ellison found that online firms rely on two primary means of obfuscation. In the first, the firm lists a cheap but inferior product that the price search engine displays at the beginning of its listings. Customers click on this product and are redirected to the firm's website, where the company then offers a more expensive product upgrade. Once one firm undercuts its competitors with this "loss leader" strategy, all firms will list similarly cheap products or risk having their product listing buried deep in the last pages of the listings. As a result, it becomes more time consuming for the customer to compare the prices of the product "upgrades," and the firm can charge a price higher than marginal cost without the risk of being priced out.

Another common strategy is the use of product add-ons. As with the first method, firms list artificially cheap products that bait consumers into visiting their website. This time, instead of upgraded products, customers are offered product add-ons, such as additional screws to fasten the chip to the circuit board or a snazzy mouse pad. Often, these products are added on automatically; that is, to purchase only the original product, the consumer has to unselect a number of additional products. Although the product the consumer initially selected may be selling at or even below marginal cost, the add-ons often sell at inflated prices—the mouse pad one online firm offered cost nearly $12 according to Ellison and Ellison. This practice allows the firm to sell the entire bundle of products at a price above marginal cost.

Obfuscation methods such as these are part of the reason the Bertrand model with identical products that we first studied is so unusual in the real world. Even products that aren't obviously differentiable can be made to stand out through some clever strategies devised by the firms selling them. Given that firms selling such products would otherwise expect to earn something close to nothing, they have a massive incentive to figure out differentiation strategies, and thus try to reduce competition. ■

[8] Glenn Ellison and Sara Fisher Ellison, "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica* 77, no. 2 (2009): 427–452.

11.7 Monopolistic Competition

Model Assumptions Monopolistic Competition

- Industry firms sell differentiated products that consumers do not view as perfect substitutes.
- Other firms' choices affect a firm's residual demand curve, but the firm ignores any strategic interactions between its own quantity or price choice and that of its competitors.
- There is free entry into the market.

monopolistic competition A market structure characterized by many firms selling a differentiated product with no barriers to entry.

In the models we've studied thus far, we haven't considered the possibility that other firms might want to enter markets in which firms are earning positive economic profits. Presumably, other firms exist that would like a piece of such action. If there are no barriers to entering a market like the snowboard market, an additional firm will cause Burton and K2's profits to decline. We saw in the Cournot model that adding more firms to the industry drove the equilibrium closer to perfect competition. In this section, we look at our last model of imperfect competition, and see what happens when there is entry into a market with differentiated products. **Monopolistic competition** is a *market structure characterized by many firms selling a differentiated product with no barriers to entry.* This term might sound like an oxymoron—competitive monopoly?—and in a way it is, but the term reflects the basic tension between market power and competitive forces that exists in these types of markets.

Every firm in a monopolistically competitive industry faces a downward-sloping demand curve, so it has some market power and every firm follows the monopoly pricing rule. That's where the "monopolistic" comes from. What is *competitive* about such markets is that there are no restrictions on entry as exist in monopoly markets—any number of firms can come into the industry at any time. This means that the firms in a monopolistically competitive industry, despite having market power, earn zero economic profit. (If they were making a profit, more firms would enter to acquire some of it. Entry only stops when profit is driven to zero for every firm in the market.)

Many markets are monopolistically competitive. For example, there are hundreds of fast-food restaurants in Chicago. Some differences between them exist, but basically people view such restaurants as largely interchangeable. Because travel is costly, though, each restaurant has a bit of market power in its local neighborhood. So, a restaurant does have some ability to set its own prices. At the same time, however, there's little to stop a new restaurant from opening. If people in a neighborhood become more enthralled with eating out, an existing restaurant might be able to raise its prices and earn economic profit for a brief period, but if the trend lasts, it is likely that a new restaurant will open up to grab some of that profit.

Keep in mind that, although monopolistic competition is categorized as "imperfect competition" along with oligopoly, there are differences between these two market structures. One is that oligopoly markets have barriers to entry, while monopolistically competitive markets do not. However, the key distinction between oligopoly and monopolistic competition is the assumption about strategic interaction. In an oligopoly, firms know that their production decisions affect their competitors' optimal choices, and all oligopolistic firms take this feedback effect into account when making their decisions. On the other hand, in monopolistic competition, firms do not worry about the production decisions of their competitors because the impact of any competitor on another is assumed to be too small for these firms to be concerned about.

Equilibrium in Monopolistically Competitive Markets

To analyze monopolistically competitive markets, let's look at a single company with market power—say for a moment that, for some reason, a city has only one fast-food restaurant. In this city, that restaurant has a monopoly on fast food. The firm faces a downward-sloping demand curve for meals served per day, as in **Figure 11.5**. We'll label this demand D_{ONE} (for one firm). The figure also shows the marginal revenue curve that corresponds to this demand, as well as the firm's average total and marginal cost curves.

Because the restaurant in Figure 11.5 is a monopolist, it produces where its marginal revenue equals marginal cost, Q^*_{ONE}. The price it charges is P^*_{ONE}. In addition to the marginal cost of production, however, the restaurant has to pay a fixed cost equal to F (this fixed cost is the reason why the firm's average total cost curve is U-shaped). The monopolist restaurant's profit is shown by the shaded rectangle: the difference between the price and the average total cost at the quantity produced, multiplied by that quantity. Because average total cost includes both variable and fixed costs, the average total cost at Q^*_{ONE}—that is, ATC^*—fully reflects all the firm's production costs.

Now suppose another restaurateur notices this firm's profit and decides to compete and open a second, slightly different fast-food restaurant. The new restaurant may differ in location, type of food served, anything that differentiates it from the existing restaurant.

The key to understanding monopolistically competitive markets is to recognize what happens to the demand curve(s) of the market's existing firm(s) when another firm enters. We know that when more substitutes for a good are available, the demand curve for the initial good becomes more elastic (less steep). Having another restaurant open up means that more substitution possibilities now exist for consumers. Instead of there being one firm with a demand curve, as in Figure 11.5, the entry of a second firm means each restaurant now has a demand curve that is a bit flatter than the monopolist firm's demand curve. And, because the demand is being split across two firms, not only is the monopolist firm's demand curve flatter, but it has shifted in as well. **Figure 11.6** shows this change from one

Figure 11.5 Demand and Cost Curves for a Monopoly

A monopolist restaurant has demand D_{ONE}, marginal revenue MR_{ONE}, average total cost ATC, and marginal cost MC. The restaurant produces where marginal revenue equals marginal cost, at quantity Q^*_{ONE}. The restaurant's profit, represented by the shaded rectangle, is the difference between the firm's price P^*_{ONE} and average total cost ATC^*, multiplied by Q^*_{ONE}.

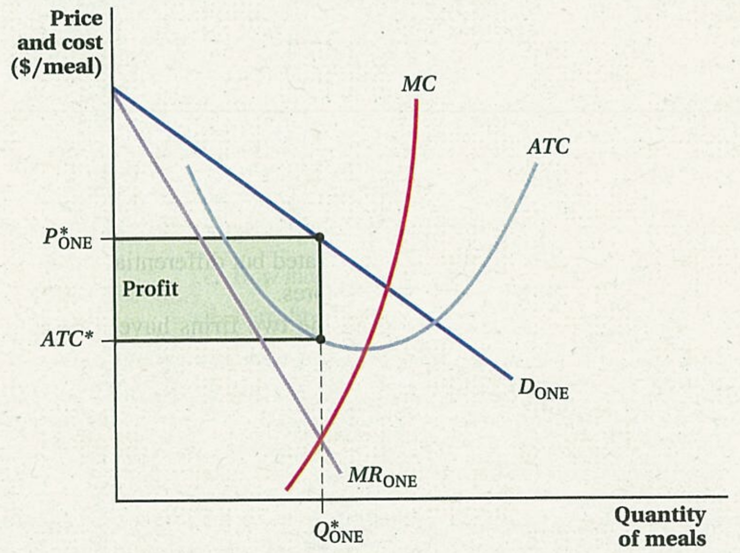

Figure 11.6 Effect of Firm Entry on Demand for a Monopolistically Competitive Firm

When a second restaurant enters the market, the original restaurant's demand curve shifts left from D_{ONE} to the more elastic residual demand curve D_{TWO}, and the marginal revenue curve MR_{ONE} shifts to MR_{TWO}. The restaurant now sells quantity Q^*_{TWO} at price P^*_{TWO} and earns profit represented by the shaded rectangle.

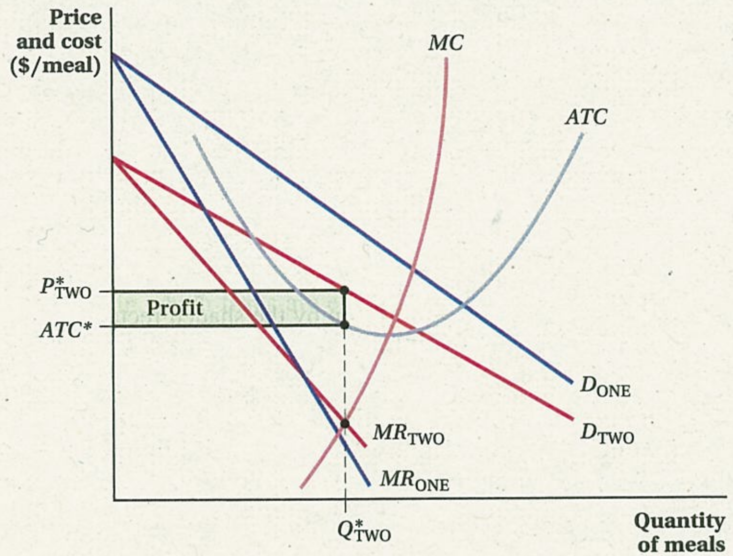

to two firms, as the initial (monopolist) firm's demand curve (now it is a residual demand curve) shifts from D_{ONE} to D_{TWO}. Notice how D_{TWO} is both flatter than D_{ONE} and to the left of it. The marginal revenue curves also shift accordingly. (The figure illustrates only what's going on for one of the two firms in the market; the picture is exactly the same for the other firm.)

Even after entry, however, both firms are essentially monopolists over their own residual demand curves. Each individual firm's demand curve reflects the fact that (1) it is splitting the market with another firm and (2) the presence of a substitute product makes the firm's demand more elastic. The competitor's presence *is* accounted for, but it is incorporated in the firm's residual demand curve. In monopolistic competition, the firm takes this residual demand as given. This is different from the oligopoly models we covered, in which firms realize that their actions affect the desired actions of their competitors, which, in turn, affect their own optimal action, and so on. This strategic interaction is captured in firms' reaction curves. A monopolistically competitive firm, on the other hand, acts as if it is in its own little monopoly world, even though its competitors' actions affect the residual demand it faces. This assumption about monopolistically competitive firms' ignorance of strategic interactions is more likely to hold in industries where there are a large number of firms selling related but differentiated products, such as car washes, self-storage facilities, and mattress stores.

Assuming the two firms have identical residual demand curves, both produce the quantity Q^*_{TWO} at which marginal revenue equals marginal cost and charge the profit-maximizing price P^*_{TWO} at that quantity. Each firm earns the profit given by the shaded rectangle in the figure.

Because two firms in the market make positive economic profit, still more firms will want to enter. Each new firm that enters will further shift the other individual companies' demand curves to the left and make them more elastic (flatter).

Entry will cease only when industry firms are no longer making economic profit. At that point, the market will look like **Figure 11.7**. When there are N firms in the market, each firm's residual demand curve eventually shifts back to D_N. Faced with this demand

Figure 11.7 Long-Run Equilibrium for a Monopolistically Competitive Market

In a monopolistically competitive market with N firms, firms face long-run demand D_N, marginal revenue MR_N, marginal costs M_C, and average total cost ATC. At the long-run equilibrium, the firm's quantity is Q^*_N, price P^*_N is equal to average cost ATC^*, and each firm earns zero economic profit.

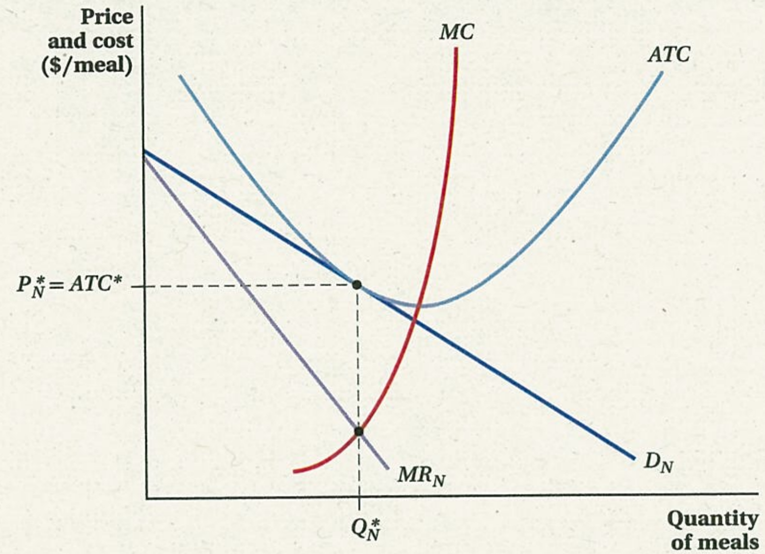

curve, the firm produces the quantity Q^*_N at which marginal revenue equals marginal cost, charges a price of P^*_N, and earns zero economic profit.

Economic profit equals zero at this point because the firm's average total cost curve and its demand curve are tangent at Q^*_N and P^*_N. When price equals average total cost, profit is zero. The firm is just covering its costs of operation (variable and fixed) at this point.

Here's an important point about monopolistically competitive markets: Even though entry occurs until profits are zero, the entry process does not ultimately lead to a perfectly competitive outcome in which price equals marginal cost (because there are fixed costs). Firms in a monopolistically competitive market face a downward-sloping demand curve, so marginal revenue is always less than price. At the profit-maximizing output, marginal cost will equal marginal revenue, which means that marginal cost will also be less than price. Free entry ensures that this markup over marginal cost is just enough to cover the firm's fixed cost, and no more.

figure it out 11.5

Sticky Stuff produces cases of taffy in a monopolistically competitive market. The inverse demand curve for its product is $P = 50 - Q$, where Q is in thousands of cases per year and P is dollars per case.

Sticky Stuff can produce each case of taffy at a constant marginal cost of $10 per case and has no fixed cost. Its total cost curve is therefore $TC = 10Q$.

a. To maximize profit, how many cases of taffy should Sticky Stuff produce each month?

b. What price will Sticky Stuff charge for a case of taffy?

c. How much profit will Sticky Stuff earn each year?

d. In reality, firms in monopolistic competition generally face fixed costs in the short run. Given the information above, what would Sticky Stuff's fixed costs have to be in order for this industry to be in long-run equilibrium? Explain.

Solution:

a. Sticky Stuff maximizes its profit by producing where $MR = MC$. Since the demand curve is linear, we know from Chapter 9 that the *MR* curve will be linear with twice the slope. Therefore, $MR = 50 - 2Q$. Setting $MR = MC$, we get

$$50 - 2Q = 10$$
$$Q = 20$$

Sticky Stuff should produce 20,000 cases of taffy each year.

b. We can find the price Sticky Stuff will charge by substituting the quantity into the demand curve:

$$P = 50 - Q = 50 - 20 = \$30 \text{ per case}$$

c. Total revenue for Sticky Stuff will be $TR = P \times Q = \$30 \times 20{,}000 = \$600{,}000$. Total cost will be $TC = 10Q = (10 \times 20{,}000) = \$200{,}000$. Therefore, Sticky Stuff will earn an annual profit of $\pi = TR - TC = \$600{,}000 - \$200{,}000 = \$400{,}000$.

d. Long-run equilibrium occurs when firms have no incentive to enter or exit. Therefore, firms must be earning zero economic profit. From (c), we know that Sticky Stuff is earning a profit of \$400,000. In order for its profit to be zero, Sticky Stuff must face annual fixed costs equal to \$400,000.

11.8 Conclusion

In this chapter, we've looked at multiple models of imperfect competition—that middle ground between perfect competition (which we studied in Chapter 8) and monopoly (which we studied in Chapter 9). We started with the reminder that the number of firms in a market is only one of many factors that can determine market prices, quantities, and producer profits. So, it's no surprise that there are different models of imperfect competition, each of which offers different predictions about market outcomes. Which model is the most applicable to any market situation requires some judgment on the part of the economist. Are the products essentially identical, or slightly or completely differentiated? Are the firms setting prices or quantities? Are firms making their choices simultaneously or in sequence? Are there barriers to entry or is entry into the market free? These and other questions need to be considered when choosing the imperfect competition model most applicable to the industry being analyzed. In the next chapter, we will examine how individuals and firms may act strategically to achieve a greater outcome (such as increased utility or higher profits).

Summary

1. In oligopolistic markets, each firm makes production decisions conditional on its competitors' actions. The resulting market equilibrium is known as a **Nash equilibrium**, one of the cornerstones of economic game theory. A Nash equilibrium occurs when each firm is doing its best given the actions of other firms. **[Section 11.1]**
2. Oligopolistic firms may be able to form cartels, in which all participating firms coordinate their production decisions and act collectively as a monopoly. The resulting market quantity and price are equal to those from a monopoly, and industry profit is maximized. While collusive behavior allows firms to capture monopoly profits, **collusion** and **cartels** are rarely stable because every firm has the incentive to increase its own profit by producing more (pricing lower). **[Section 11.2]**
3. In **Bertrand competition,** products are identical and firms compete on price. Each firm simultaneously sets the price of its good, and consumers then choose to purchase all the quantity demanded from whichever firm has the lowest price, even if the price is only one penny lower. The Bertrand model shows that only two firms need to be in a market to achieve the perfectly competitive market outcome where price equals marginal cost. This result arises because firms in these situations have such a strong incentive to try to undercut the prices of their rivals. Market output is equal to the competitive level of output and firm profits are zero. **[Section 11.3]**
4. In contrast to firms in Bertrand competition, firms in **Cournot competition** simultaneously choose the quantity of a good to produce, and not the price at which the good sells. The Cournot equilibrium price is generally above the price in Bertrand competition, but below the monopoly price. The Cournot output is less than the Bertrand level of output, but greater than the output generated by a cartel. Firms in a Cournot oligopoly earn greater profits than those in the Bertrand model, but less than the monopoly profit. **[Section 11.4]**
5. In **Stackelberg competition,** firms make production decisions sequentially. Because the first firm in an industry can make production decisions independently of other firms and may be able to capture larger profits, a **first-mover advantage** exists for these firms. **[Section 11.5]**
6. In the Bertrand model with differentiated products, consumers in these markets are willing to substitute across goods, but do not consider them identical, or perfect substitutes. As a result, small differences in prices do not lead to all demand being satisfied by the producer with the lowest price (as in the Bertrand oligopoly with identical products). **[Section 11.6]**
7. **Monopolistic competition** is a market structure in which firms sell differentiated products, and firms have some characteristics of both monopolies and perfectly competitive firms. Because there are no barriers to entry in a monopolistically competitive market, economic profit is driven to zero through the entry of firms. **[Section 11.7]**

Review Questions

1. Name some different forms of imperfect competition.
2. Define Nash equilibrium. Why do firms in oligopoly situations reach Nash equilibria?
3. Why are collusions and cartels often unstable?
4. What is the market equilibrium in Bertrand competition with identical goods?
5. Contrast Bertrand and Cournot competition. Why do they reach different market equilibria?
6. What does the residual demand curve tell us about a firm's output in Cournot competition?
7. How can reaction curves be used to find a firm's equilibrium in Cournot competition?
8. What causes the first-mover advantage in Stackelberg competition?
9. Contrast the market equilibria in Betrand competition with identical products and with differentiated products.
10. What are the characteristics of a monopolistically competitive firm?
11. When will firms enter a monopolistically competitive industry? At what point will firms stop entering a monopolistically competitive industry?
12. Why do firms in monopolistic competition not reach the perfectly competitive equilibrium?

Problems

(Solutions to problems marked with an asterisk appear at the back of this book. Problems adapted to use calculus are available online.)

1. In 1969, tobacco companies were the largest single product advertisers on television. That same year, the Surgeon General of the United States released a report linking smoking to adverse health consequences. Shortly thereafter, the federal government banned cigarette companies from advertising on television. Suppose you were an executive at a tobacco company at that time. Read Section 11.1 carefully, then explain how you might react to the federal government's advertising ban, and why.
2. Suppose you and a rival are the only producers of oysters in an isolated town. Every morning you both dive for oysters that you will sell in the market that afternoon. Each morning, you both have a choice of bringing up 10 or 20 dozen oysters; each dozen you bring up has a marginal cost of \$10. If 20 dozen oysters are brought to market in total, they will sell for \$35 each. If 30 dozen oysters are brought to market, they will sell for \$25 each. If 40 dozen oysters are brought to market, they will sell for \$20 each. The following table shows the profit you and a rival can expect to earn based on your choice of bringing up 10 or 20 dozen oysters each:

		Your Rival	
		10 dozen	**20 dozen**
You	**10 dozen**	250, 250	150, 300
	20 dozen	300, 150	200, 200

   a. Verify that the profits represented in the table above are accurate.
   b. Where is the Nash equilibrium in the game you and your rival play?
   c. Is the Nash equilibrium one that you and your rival would agree to if you were to discuss production before diving each day? If not, explain why that agreement is unlikely to be honored.
   d. Draw a parallel between the game described in this problem and the advertising game Disney and Warner Brothers played in Table 11.1.

*3. Because cooking soufflés is incredibly difficult, the supply of soufflés in a small French town is controlled by two bakers, Gaston and Pierre. The demand for soufflés is given by $P = 30 - 2Q$, and the marginal and average total cost of producing soufflés is \$6. Because baking a soufflé requires a great deal of work and preparation, each morning Gaston and Pierre make a binding decision about how many soufflés to bake.
   a. Suppose that Pierre and Gaston agree to collude, evenly splitting the output a monopolist would make and charging the monopoly price.
      i. Derive the equation for the monopolist's marginal revenue curve.

ii. Determine the profit-maximizing collective output for the cartel.

iii. Determine the price Pierre and Gaston will be able to charge.

iv. Determine profits for Pierre and Gaston individually, as well as for the cartel as a whole.

b. Suppose that Pierre cheats on the cartel agreement by baking one extra soufflé each morning.

i. What does the extra production do to the price of soufflés in the marketplace?

ii. Calculate Pierre's profit. How much did he gain by cheating?

iii. Calculate Gaston's profit. How much did Pierre's cheating cost him?

iv. How much potential profit does the group lose as a result of Pierre's cheating?

c. Suppose that Gaston, fed up with Pierre's behavior, also begins baking one extra soufflé each morning.

i. How does the extra production affect the price of soufflés in the marketplace?

ii. Calculate Gaston's profit. How much did he gain by cheating?

iii. Calculate Pierre's profit. How much did Gaston's cheating cost him?

iv. How much potential profit does the group lose as a result of Pierre's and Gaston's cheating?

v. Demonstrate that it is in neither Pierre's nor Gaston's best interest to cheat further on their agreement.

4. Consider the soufflé bakers Gaston and Pierre, described in Problem 3. Suppose Gaston and Pierre are each faced with the choice of baking 3 or 4 souffles each morning.

a. Calculate the profits each would earn based on the decision each makes. Insert those numbers into the table below, with Gaston's number on the left and Pierre's on the right.

		Pierre	
		Bake 4 Souffles	Bake 3 Souffles
Gaston	Bake 4 Souffles	___, ___	___, ___
	Bake 3 Souffles	___, ___	___, ___

b. Explain how the profits Gaston earns are conditional on the decision Pierre makes.

c. Draw parallels between this problem and the advertising game Disney and Warner Brothers play in Table 11.1.

5. Suppose in Problem 3 that Gaston can produce soufflés at a constant marginal cost of $5, but Pierre produces soufflés for $7. Together, they collude to produce 3 units each.

a. How much profit will each producer earn? What will be the total profit of the cartel?

b. Gaston observes that he is a more efficient producer than Pierre, and suggests that if they are going to produce 6 units, the cartel's interests are better served if Gaston produces all of the soufflés.

i. If Gaston produces and sells all the soufflés and Pierre produces nothing, what happens to the profit of the cartel?

ii. Is Pierre likely to agree not to produce any soufflés?

iii. Suppose Gaston offers to pay Pierre not to produce any soufflés. How much would Gaston potentially be willing to offer? What is the minimum offer that Pierre should accept?

iv. Suppose that the deal in part (iii) is reached for Pierre's minimum price. What happens to Pierre's profit if he cheats on his agreement with Gaston and increases his output from zero soufflés to 1? What happens to Gaston's profit?

v. Compare Pierre's incentive to cheat under this arrangement with the incentive that exists when they split production equally. Also compare Gaston's vulnerability to Pierre's cheating under both arrangements. Why might this cartel choose to use the less profitable method of each member producing 3 units to the potentially more profitable method of having Gaston produce everything?

6. The Organization of Petroleum Exporting Countries (OPEC) is a cartel that attempts to keep oil prices high by restricting output. As part of that process, each member nation is assigned a production quota; most members have nationalized their oil industry so that the government controls overall production. However, member nations routinely exceed their production targets. Read "What Makes Collusion Easier" in Section 11.2; then explain why OPEC often has difficulty keeping output low and prices high. Do you think that violators are more likely to emerge from politically stable countries or unstable countries? From monarchies or democracies?

7. Suppose that the inverse market demand for pumpkins is given by $P = \$10 - 0.05Q$. Pumpkins can be grown by anyone at a constant marginal cost of $1.

a. If there are lots of pumpkin growers in town so that the pumpkin industry is competitive, how many pumpkins will be sold, and what price will they sell for?

b. Suppose that a freak weather event wipes out the pumpkins of all but two producers, Linus and Lucy. Both Linus and Lucy produced bumper crops and have more than enough pumpkins available to satisfy the demand at even a zero price. If Linus and Lucy collude to generate monopoly profits, how many pumpkins will they sell, and what price will they sell for?

c. Suppose that the predominant form of competition in the pumpkin industry is price competition. In other words, suppose that Linus and Lucy are Bertrand competitors. What will be the final price of pumpkins in this market—in other words, what is the Bertrand equilibrium price?

d. At the Bertrand equilibrium price, what will be the final quantity of pumpkins sold by both Linus and Lucy individually, and for the industry as a whole? How profitable will Linus and Lucy be?

e. Would the results you found in parts (c) and (d) be likely to hold if Linus let it be known that his pumpkins were the most orange in town, and Lucy let it be known that hers were the tastiest? Explain.

f. Would the results you found in parts (c) and (d) hold if Linus could grow pumpkins at a marginal cost of $0.95?

8. Andres and Julian have the only liquor licenses in a small resort town. The inverse demand for mimosas (a favorite adult breakfast drink) is given by $P = 8 - 0.5Q$. Mimosas can be produced at a constant average and marginal cost of $2.

a. Adam Smith once wrote, "People of the same trade seldom meet together . . . but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices." Suppose Andres and Julian were to follow Smith's advice and conspire to raise prices to the level a monopolist would charge. What price should they agree to set in order to maximize profits?

b. Suppose both Andres and Julian each have two choices: charge the price you found in (a), or charge $1 less. Fill in the prices each may choose to charge. Then determine the profits each will earn given the choice each makes, and put them in the table below, with Andres's profits before the comma and Julian's after it.

		Julian	
		Price from (a):	$1 less:
Andres	Price from (a):	___, ___	___, ___
	$1 less:	___, ___	___, ___

c. Draw parallels between this problem and the advertising game Disney and Warner Brothers play in Table 11.1.

9. Suppose that three grocery stores sell Bubba's Gourmet Red Beans and Rice. Bullseye Market is able to acquire, stock, and market them for $2.00 per package. OKMart can acquire, stock, and market them for $1.98 per package. SamsMart can acquire, stock, and market them for $1.96 per package.

a. If the three competitors are located in close proximity to one another, so the cost of going to a different store to purchase red beans and rice is negligible, and if the market for prepackaged gourmet red beans and rice is characterized by Bertrand competition, what will the prevailing market price be?

b. Where will customers buy their red beans and rice? Bullseye Market, OKMart, or SamsMart? What does your answer suggest about the potential rewards to small improvements in efficiency via cost-cutting?

c. Suppose that each day, equal numbers of customers begin their shopping at each of the three stores. If the cost of going to a different store to purchase red beans and rice is 2 cents, is the Bertrand result likely to hold in this case? Where will customers purchase red beans and rice? Where will they not purchase them?

*10. The platypus is a shy and secretive animal that does not breed well in captivity. But two breeders, Sydney and Adelaide, have discovered the secret to platypus fertility and have effectively cornered the market. Zoos across the globe come to them to purchase their output; the world inverse demand for baby platypuses is given by $P = 1{,}000 - 2Q$, where Q is the combined output of Sydney (q_S) and Adelaide (q_A).

a. Sydney wishes to produce the profit-maximizing quantity of baby platypus. Given Adelaide's choice of output, q_A, write an equation for the residual demand faced by Sydney.

b. Derive Sydney's residual marginal revenue curve.

c. Assume that the marginal and average total cost of raising a baby platypus to an age at which it can be sold is $200. Derive Sydney's reaction function.

d. Repeat steps (a), (b), and (c) to find Adelaide's reaction function to Sydney's output choice.

e. Substitute Sydney's reaction function into Adelaide's to solve for Adelaide's profit-maximizing level of output. Then use your answer to find Sydney's profit-maximizing level of output.

f. Determine industry output, the price of platypus, and the profits of both Sydney and Adelaide.

g. If Adelaide were hit by a bus on her way home from work, and Sydney were to become a monopolist, what would happen to industry quantity, price, and profit?

11. Suppose that two firms are Cournot competitors. Industry demand is given by $P = 200 - q_1 - q_2$, where q_1 is the output of Firm 1 and q_2 is the output of Firm 2. Both Firm 1 and Firm 2 face constant marginal and average total costs of $20.
    a. Solve for the Cournot price, quantity, and firm profits.
    b. Firm 1 is considering investing in costly technology that will enable it to reduce its costs to $15 per unit. How much should Firm 1 be willing to pay if such an investment can guarantee that Firm 2 will not be able to acquire it?
    c. How does your answer to (b) change if Firm 1 knows the technology is available to Firm 2?
12. Consider the demand for boccie balls shown in the diagram below. Demand is given by $P = 80 - Q$. Boccie balls can be produced at a constant marginal and average total cost of $20.

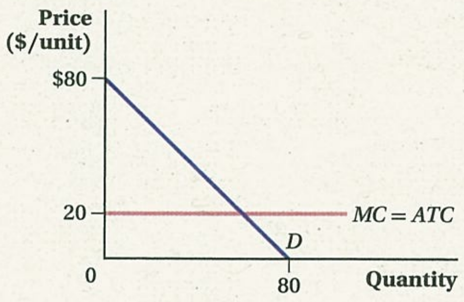

    a. If the boccie ball industry were perfectly competitive, what quantity would be sold, and what price would prevail in the market?
    b. Suppose that the boccie ball industry were a monopoly. Draw in a marginal revenue curve and determine the profit-maximizing quantity.
        i. Divide the monopoly (one-firm) quantity by the competitive quantity to determine the proportion of competitive output that a monopolist provides. Present your answer in reduced fractional form.
        ii. Determine the price and draw a dot on the demand curve indicating the monopolist's price and quantity.
    c. Suppose the boccie ball industry were a Cournot duopoly, with two firms. Use the procedures developed in this chapter to determine the industry output.
        i. Divide the duopoly quantity by the competitive quantity to determine the proportion of competitive output that a duopoly provides. Present your answer in reduced fractional form.
        ii. Determine the price and draw a dot on the demand curve indicating the duopoly's price and quantity.
    d. Hypothesize as to the fraction of competitive output that would be sold if the boccie ball industry had three identical Cournot competitors. Then check your answer by deriving reaction functions for a three-firm oligopoly and solving for each firm's output.
    e. In general, what fraction of the competitive output level will be brought to market if there are N identical firms in the industry?
    f. What happens to the quantity sold as more competitors are added to the industry? The price? What happens to consumer surplus and deadweight loss? Does this provide support for the government's desire to ensure competitive industries rather than monopolies or small oligopolies?
13. Two organic emu ranchers, Bill and Ted, serve a small metropolitan market. Bill and Ted are Cournot competitors, making a conscious decision each year regarding how many emus to breed. The price they can charge depends on how many emus they collectively raise, and demand in this market is given by $Q = 150 - P$. Bill raises emus at a constant marginal and average total cost of $10; Ted raises emus at a constant marginal and average total cost of $20.
    a. Find the Cournot equilibrium price, quantity, profits, and consumer surplus.
    b. Suppose that Bill and Ted merge and become a monopoly provider of emus. Furthermore, suppose that Ted adopts Bill's production techniques. Find the monopoly price, quantity, profits, and consumer surplus.
    c. Suppose that instead of merging, Bill considers buying Ted's operation for cash. How much should Bill be willing to offer Ted to purchase his emu ranch? (Assume that the combined firms are only going to operate for one period.)
    d. Has the combination of the two ranches discussed above been good for society or bad for society? Discuss how the forces of monopoly power and increased efficiency tend to push social well-being in opposite directions.

*14. The market for nutmeg is controlled by two small island economies, Penang and Grenada. The market demand for bottled nutmeg is given by $P = 100 - q_P - q_G$, where q_P is the quantity Penang produces and q_G is the quantity Grenada produces. Both Grenada and Penang produce nutmeg at a constant marginal and average cost of $20 per bottle.
    a. Verify that the reaction function for Grenada is given by $q_G = 40 - 0.5q_P$. Then verify that the reaction function for Penang is given by $q_P = 40 - 0.5q_G$.
    b. Find the Cournot equilibrium quantity for each island. Then solve for the market price of nutmeg and for each firm's profit.
    c. Suppose that Grenada transforms the nature of competition to Stackelberg competition by announcing its production targets publicly in an attempt to seize a first-mover advantage.
        i. Grenada must first decide how much to produce, and to do this, it needs to know the demand conditions it faces. Substitute Penang's reaction function into the market demand curve to find the demand faced by Grenada.
        ii. Based on your answer to the problem above, find the marginal revenue curve faced by Grenada.
        iii. Equate marginal revenue with marginal cost to find Grenada's output.
        iv. Plug Grenada's output into Penang's reaction function to determine Penang's output.
        v. Plug the combined output of Grenada and Penang into the market demand curve to determine the price. How do the industry quantity and price compare to those under Cournot competition?
        vi. Determine profits in Grenada and Penang. How do the profits of each compare to profits under Cournot competition? Is there an advantage to being the first-mover?
15. Two farmers, Tito and Helen, supply a chain of islands with kale. The inverse demand for kale in the islands is given by $P = 60 - 0.5Q$, where Q is the combined output of Tito (q_T) and Helen (q_H), measured in 10-pound bunches. Tito grows kale at a constant marginal and average cost of $12 per bunch; Helen grows kale at a constant marginal and average cost of $10 per bunch.
    a. Suppose this market is a Stackelberg oligopoly and Tito is the first-mover. How much will he and Helen produce? What will the market price of kale be? How much profit will each farmer earn?
    b. Now suppose that Helen is the first-mover in this Stackelberg oligopoly. How much will each farmer produce? What will the market price of kale be? How much profit will each farmer earn?
    c. Quantify the value to Helen of being the first-mover in this Stackelberg game.
16. In each case below, identify the type of competition and determine if there is likely to be a first-mover advantage.
    a. Saudi Arabia, a major oil producer, announces its annual oil production target to the world.
    b. L.L.Bean and Land's End sell nearly identical outerwear via mail order. Each is anxious to publish its fall catalog; once that catalog is published, the firm cannot change its prices without undertaking another costly mailing.
17. August and François are the only sellers of sparkling water at a market in a small, rural French town. They obtain their sparkling water for free from wells in their backyards and transport it to the market in wheelbarrows; neither has access to motorized transportation. Identify the type of oligopoly (Cournot, Bertrand, Stackelberg) that is the best fit for each situation below and explain your reasoning:
    a. August and François both live 4 hours' walk from the market.
    b. August and François both live half a block from the market.
    c. August lives a long walk away, but is an early riser who always arrives at 8:00 A.M.; François lives quite close and never shows up until 8:30.
18. Internet users in a small Colorado town can access the Web in two ways: via their television cable or via a digital subscriber line (DSL) from their telephone company. The cable and telephone companies are Bertrand competitors, but because changing providers is slightly costly (waiting for the cable repairman can eat up at least small amounts of time!), customers have some slight resistance to switching from one to another. The demand for cable Internet services is given by $q_C = 100 - 3p_C + 2p_T$, where q_C is the number of cable Internet subscribers in town, p_C is the monthly price of cable Internet service, and p_T is the price of a DSL line from the telephone company. The demand for DSL Internet service is similarly given by $q_T = 100 - 3p_T + 2p_C$. Assume that both sellers can produce broadband service at zero marginal cost.
    a. Derive the cable company's reaction curve. Your answer should express p_C as a function of p_T.
    b. Derive the telephone company's reaction curve. Your answer should express p_T as a function of p_C.
    c. Combine reaction functions to determine the price each competitor should charge. Then determine each competitor's quantity and profits, assuming that the average total costs are zero.
    d. Suppose that the cable company begins to offer slightly faster service than the telephone company, which alters demands for the two products. Now $q_C = 100 - 2p_C + 3p_T$ and $q_T = 100 - 4p_T + p_C$. Show what effect this increase in service has on the prices and profit of each competitor.
19. Consider two Bertrand competitors in the market for brie, François and Babette. The cheeses of François and Babette are differentiated, with the demand for François' cheese given by $q_F = 30 - p_F + p_B$, where q_F is the quantity François sells, p_F is the price François charges, and p_B is the price charged by Babette. The demand for Babette's cheese is similarly given as $q_B = 30 - p_B + p_F$. Assume that the marginal cost of producing cheese is zero.
    a. Find the Bertrand equilibrium prices and quantities for these two competitors.
    b. Now consider a situation in which François sets his price first and Babette responds. Follow procedures similar to those you used for Stackelberg quantity competition to solve for François's profit-maximizing price, quantity, and profit.

c. Solve for Babette's profit-maximizing price, quantity, and profit.

d. Was François's attempt to seize the first-mover advantage worthwhile?

20. There are only three big tobacco companies, but they produce dozens of brands of cigarettes. Compare and contrast Bertrand competition with undifferentiated and differentiated products to explain why the big three tobacco companies devote many resources to support so many different brands instead of each producing just a single type of generic cigarette. Do you think supporting all these different brands is good for society, or bad?

21. Consider a monopolistically competitive industry. A graph of demand and cost conditions for a typical firm is depicted in the diagram below:

a. Is this firm generating producer surplus? Is this firm earning a profit? How can you reconcile your answers?

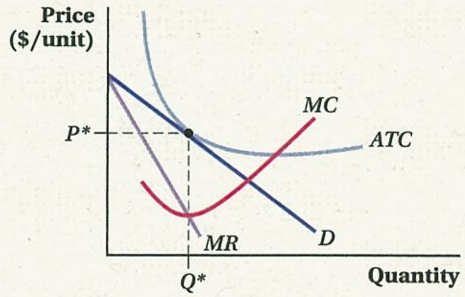

b. Do you expect any entry into or exit from this industry to occur? Explain.

c. Suppose that the government reduces annual licensing fees, causing the fixed cost of the typical firm to fall. Make appropriate shifts of all curves that might be affected. What happens to producer surplus? What happens to profit? Do you expect the fall in fixed costs to cause entry into or exit from this industry? Explain.

d. Shift the demand and marginal revenue curves to reflect the entry/exit you indicated in (c). Find the new equilibrium.

e. Continue to reduce fixed cost. What happens to the demand curve as fixed cost continues to fall? What happens to producer surplus and profit?

f. Find the equilibrium as fixed cost falls to zero.

22. Sally sells brilliant economics lectures to knowledge-seeking students. (This industry is monopolistically competitive: There are at least two other brilliant lecturers Sally competes with.) The inverse demand for Sally's lectures is given by $P = 60 - 0.5Q$, where Q measures the number of lectures Sally gives each week. The total cost of her delivering lectures is given by $TC = 4Q + F$, where F represents her fixed costs. The marginal cost of each lecture is therefore \$4.

a. To maximize profit, how many lectures should Sally deliver each week?

b. What price will Sally charge for her lectures?

c. How much producer surplus will Sally earn?

d. What must Sally's fixed costs be for the industry to be in long-run equilibrium? If Sally's fixed costs were lower than this, what would you expect to happen to the demand for Sally's lectures in the long run?

23. One big question economics ponders is how to produce the greatest material well-being using the fewest resources. Compare and contrast perfect competition and monopolistic competition in achieving that end. (*Hint:* You may want to consider a particular monopolistically competitive industry such as clothing or restaurant meals, and imagine what that industry would look like if it were perfectly competitive instead.) How does your answer depend on your definition of material well-being?

24. When competition between firms is based on quantities (Cournot competition), the reaction functions we derive tell us that when Firm A increases its output, Firm B's best response is to cut its own. However, when competition between firms is based on price (Bertrand competition), reaction functions tell us that Firm B's response to a cut in Firm A's price (which will lead to an increase in the quantity A sells) should be a corresponding cut in B's price (and a corresponding increase in its own output). Reconcile these two results.

*25. Suppose that the market demand for rose hips is given by $P = 100 - Q$. There are two firms, A and B, producing rose hips, each at a constant marginal and average total cost of \$5. Fill in the table below for each market structure.

	Collusive Monopoly	Cournot Oligopoly	Bertrand Oligopoly	Stackelberg Oligopoly (A is first-mover)
A's Quantity				
B's Quantity				
Industry Quantity				
Price				
A's Profit				
B's Profit				
Industry Profit				