13

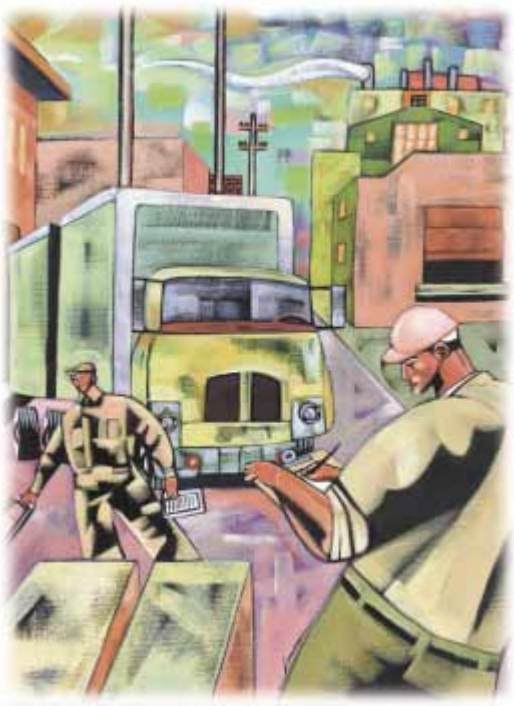

THE COSTS OF PRODUCTION

IN THIS CHAPTER YOU WILL . . .

- *Examine what items are included in a firm's costs of production*
- *Analyze the link between a firm's production process and its total costs*
- *Learn the meaning of average total cost and marginal cost and how they are related*
- *Consider the shape of a typical firm's cost curves*
- *Examine the relationship between short-run and long-run costs*

The economy is made up of thousands of firms that produce the goods and services you enjoy every day: General Motors produces automobiles, General Electric produces lightbulbs, and General Mills produces breakfast cereals. Some firms, such as these three, are large; they employ thousands of workers and have thousands of stockholders who share in the firms' profits. Other firms, such as the local barbershop or candy store, are small; they employ only a few workers and are owned by a single person or family.

In previous chapters we used the supply curve to summarize firms' production decisions. According to the law of supply, firms are willing to produce and sell a greater quantity of a good when the price of the good is higher, and this response leads to a supply curve that slopes upward. For analyzing many questions, the law of supply is all you need to know about firm behavior.

In this chapter and the ones that follow, we examine firm behavior in more detail. This topic will give you a better understanding of what decisions lie behind

the supply curve in a market. In addition, it will introduce you to a part of economics called *industrial organization*—the study of how firms' decisions regarding prices and quantities depend on the market conditions they face. The town in which you live, for instance, may have several pizzerias but only one cable television company. How does this difference in the number of firms affect the prices in these markets and the efficiency of the market outcomes? The field of industrial organization addresses exactly this question.

As a starting point for the study of industrial organization, this chapter examines the costs of production. All firms, from Delta Air Lines to your local deli, incur costs as they make the goods and services that they sell. As we will see in the coming chapters, a firm's costs are a key determinant of its production and pricing decisions. Establishing what a firm's costs are, however, is not as straightforward as it might seem.

WHAT ARE COSTS?

We begin our discussion of costs at Hungry Helen's Cookie Factory. Helen, the owner of the firm, buys flour, sugar, flavorings, and other cookie ingredients. She also buys the mixers and ovens and hires workers to run this equipment. She then sells the resulting cookies to consumers. By examining some of the issues that Helen faces in her business, we can learn some lessons that apply to all firms in the economy.

TOTAL REVENUE, TOTAL COST, AND PROFIT

We begin with the firm's objective. To understand what decisions a firm makes, we must understand what it is trying to do. It is conceivable that Helen started her firm because of an altruistic desire to provide the world with cookies or, perhaps, out of love for the cookie business. More likely, however, Helen started her business to make money. Economists normally assume that the goal of a firm is to maximize profit, and they find that this assumption works well in most cases.

total revenue
the amount a firm receives for the sale of its output

total cost
the market value of the inputs a firm uses in production

profit
total revenue minus total cost

What is a firm's profit? The amount that the firm receives for the sale of its output (cookies) is called its **total revenue.** The amount that the firm pays to buy inputs (flour, sugar, workers, ovens, etc.) is called its **total cost.** Helen gets to keep any revenue that is not needed to cover costs. We define **profit** as a firm's total revenue minus its total cost. That is,

$$\text{Profit} = \text{Total revenue} - \text{Total cost}.$$

Helen's objective is to make her firm's profit as large as possible.

To see how a firm goes about maximizing profit, we must consider fully how to measure its total revenue and its total cost. Total revenue is the easy part: It equals the quantity of output the firm produces times the price at which it sells its output. If Helen produces 10,000 cookies and sells them at $2 a cookie, her total revenue is $20,000. By contrast, the measurement of a firm's total cost is more subtle.

COSTS AS OPPORTUNITY COSTS

When measuring costs at Hungry Helen's Cookie Factory or any other firm, it is important to keep in mind one of the *Ten Principles of Economics* from Chapter 1: The cost of something is what you give up to get it. Recall that the *opportunity cost* of an item refers to all those things that must be forgone to acquire that item. When economists speak of a firm's cost of production, they include all the opportunity costs of making its output of goods and services.

A firm's opportunity costs of production are sometimes obvious and sometimes less so. When Helen pays $1,000 for flour, that $1,000 is an opportunity cost because Helen can no longer use that $1,000 to buy something else. Similarly, when Helen hires workers to make the cookies, the wages she pays are part of the firm's costs. These are **explicit costs.** By contrast, some of a firm's opportunity costs are **implicit costs.** Imagine that Helen is skilled with computers and could earn $100 per hour working as a programmer. For every hour that Helen works at her cookie factory, she gives up $100 in income, and this forgone income is also part of her costs.

explicit costs
input costs that require an outlay of money by the firm

implicit costs
input costs that do not require an outlay of money by the firm

This distinction between explicit and implicit costs highlights an important difference between how economists and accountants analyze a business. Economists are interested in studying how firms make production and pricing decisions. Because these decisions are based on both explicit and implicit costs, economists include both when measuring a firm's costs. By contrast, accountants have the job of keeping track of the money that flows into and out of firms. As a result, they measure the explicit costs but often ignore the implicit costs.

The difference between economists and accountants is easy to see in the case of Hungry Helen's Cookie Factory. When Helen gives up the opportunity to earn money as a computer programmer, her accountant will not count this as a cost of her cookie business. Because no money flows out of the business to pay for this cost, it never shows up on the accountant's financial statements. An economist, however, will count the forgone income as a cost because it will affect the decisions that Helen makes in her cookie business. For example, if Helen's wage as a computer programmer rises from $100 to $500 per hour, she might decide that running her cookie business is too costly and choose to shut down the factory in order to become a full-time computer programmer.

THE COST OF CAPITAL AS AN OPPORTUNITY COST

An important implicit cost of almost every business is the opportunity cost of the financial capital that has been invested in the business. Suppose, for instance, that Helen used $300,000 of her savings to buy her cookie factory from the previous owner. If Helen had instead left this money deposited in a savings account that pays an interest rate of 5 percent, she would have earned $15,000 per year. To own her cookie factory, therefore, Helen has given up $15,000 a year in interest income. This forgone $15,000 is one of the implicit opportunity costs of Helen's business.

As we have already noted, economists and accountants treat costs differently, and this is especially true in their treatment of the cost of capital. An economist views the $15,000 in interest income that Helen gives up every year as a cost of her business, even though it is an implicit cost. Helen's accountant, however, will not show this $15,000 as a cost because no money flows out of the business to pay for it.

To further explore the difference between economists and accountants, let's change the example slightly. Suppose now that Helen did not have the entire

$300,000 to buy the factory but, instead, used $100,000 of her own savings and borrowed $200,000 from a bank at an interest rate of 5 percent. Helen's accountant, who only measures explicit costs, will now count the $10,000 interest paid on the bank loan every year as a cost because this amount of money now flows out of the firm. By contrast, according to an economist, the opportunity cost of owning the business is still $15,000. The opportunity cost equals the interest on the bank loan (an explicit cost of $10,000) plus the forgone interest on savings (an implicit cost of $5,000).

ECONOMIC PROFIT VERSUS ACCOUNTING PROFIT

economic profit
total revenue minus total cost, including both explicit and implicit costs

accounting profit
total revenue minus total explicit cost

Now let's return to the firm's objective—profit. Because economists and accountants measure costs differently, they also measure profit differently. An economist measures a firm's **economic profit** as the firm's total revenue minus all the opportunity costs (explicit and implicit) of producing the goods and services sold. An accountant measures the firm's **accounting profit** as the firm's total revenue minus only the firm's explicit costs.

Figure 13-1 summarizes this difference. Notice that because the accountant ignores the implicit costs, accounting profit is larger than economic profit. For a business to be profitable from an economist's standpoint, total revenue must cover all the opportunity costs, both explicit and implicit.

QUICK QUIZ: Farmer McDonald gives banjo lessons for $20 an hour. One day, he spends 10 hours planting $100 worth of seeds on his farm. What opportunity cost has he incurred? What cost would his accountant measure? If these seeds will yield $200 worth of crops, does McDonald earn an accounting profit? Does he earn an economic profit?

Figure 13-1

ECONOMISTS VERSUS ACCOUNTANTS. Economists include all opportunity costs when analyzing a firm, whereas accountants measure only explicit costs. Therefore, economic profit is smaller than accounting profit.

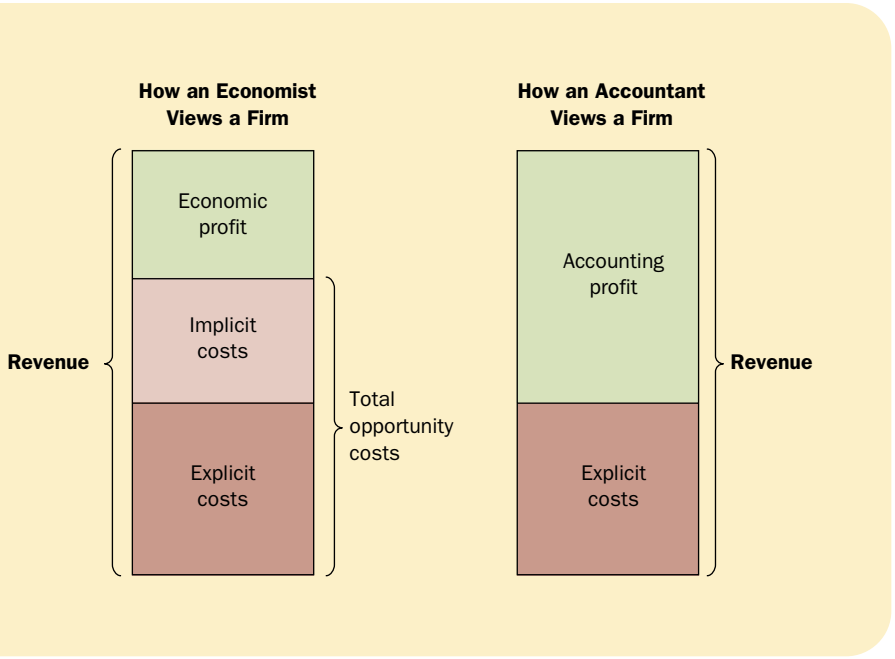

PRODUCTION AND COSTS

Firms incur costs when they buy inputs to produce the goods and services that they plan to sell. In this section we examine the link between a firm's production process and its total cost. Once again, we consider Hungry Helen's Cookie Factory.

In the analysis that follows, we make an important simplifying assumption: We assume that the size of Helen's factory is fixed and that Helen can vary the quantity of cookies produced only by changing the number of workers. This assumption is realistic in the short run, but not in the long run. That is, Helen cannot build a larger factory overnight, but she can do so within a year or so. This analysis, therefore, should be viewed as describing the production decisions that Helen faces in the short run. We examine the relationship between costs and time horizon more fully later in the chapter.

THE PRODUCTION FUNCTION

Table 13-1 shows how the quantity of cookies Helen's factory produces per hour depends on the number of workers. If there are no workers in the factory, Helen produces no cookies. When there is 1 worker, she produces 50 cookies. When there are 2 workers, she produces 90 cookies, and so on. Figure 13-2 presents a graph of these two columns of numbers. The number of workers is on the horizontal axis, and the number of cookies produced is on the vertical axis. This relationship between the quantity of inputs (workers) and quantity of output (cookies) is called the **production function.**

production function
the relationship between quantity of inputs used to make a good and the quantity of output of that good

One of the *Ten Principles of Economics* introduced in Chapter 1 is that rational people think at the margin. As we will see in future chapters, this idea is the key to understanding the decision a firm makes about how many workers to hire and how much output to produce. To take a step toward understanding these decisions, the third column in the table gives the marginal product of a worker. The **marginal product** of any input in the production process is the increase in the quantity of output obtained from an additional unit of that input. When the number of workers goes from 1 to 2, cookie production increases from 50 to 90, so the marginal product of the second worker is 40 cookies. And when the number of workers goes from 2 to 3, cookie production increases from 90 to 120, so the marginal product of the third worker is 30 cookies.

marginal product
the increase in output that arises from an additional unit of input

Notice that as the number of workers increases, the marginal product declines. The second worker has a marginal product of 40 cookies, the third worker has a marginal product of 30 cookies, and the fourth worker has a marginal product of 20 cookies. This property is called **diminishing marginal product.** At first, when only a few workers are hired, they have easy access to Helen's kitchen equipment. As the number of workers increases, additional workers have to share equipment and work in more crowded conditions. Hence, as more and more workers are hired, each additional worker contributes less to the production of cookies.

diminishing marginal product
the property whereby the marginal product of an input declines as the quantity of the input increases

Diminishing marginal product is also apparent in Figure 13-2. The production function's slope ("rise over run") tells us the change in Helen's output of

cookies ("rise") for each additional input of labor ("run"). That is, the slope of the production function measures the marginal product of a worker. As the number of workers increases, the marginal product declines, and the production function becomes flatter.

Number of Workers	Output (quantity of cookies produced per hour)	Marginal Product of Labor	Cost of Factory	Cost of Workers	Total Cost of Inputs (cost of factory + cost of workers)
0	0		$30	$ 0	$30
		50			
1	50		30	10	40
		40			
2	90		30	20	50
		30			
3	120		30	30	60
		20			
4	140		30	40	70
		10			
5	150		30	50	80

Table 13-1

A Production Function and Total Cost: Hungry Helen's Cookie Factory

Figure 13-2

Hungry Helen's Production Function. A production function shows the relationship between the number of workers hired and the quantity of output produced. Here the number of workers hired (on the horizontal axis) is from the first column in Table 13-1, and the quantity of output produced (on the vertical axis) is from the second column. The production function gets flatter as the number of workers increases, which reflects diminishing marginal product.

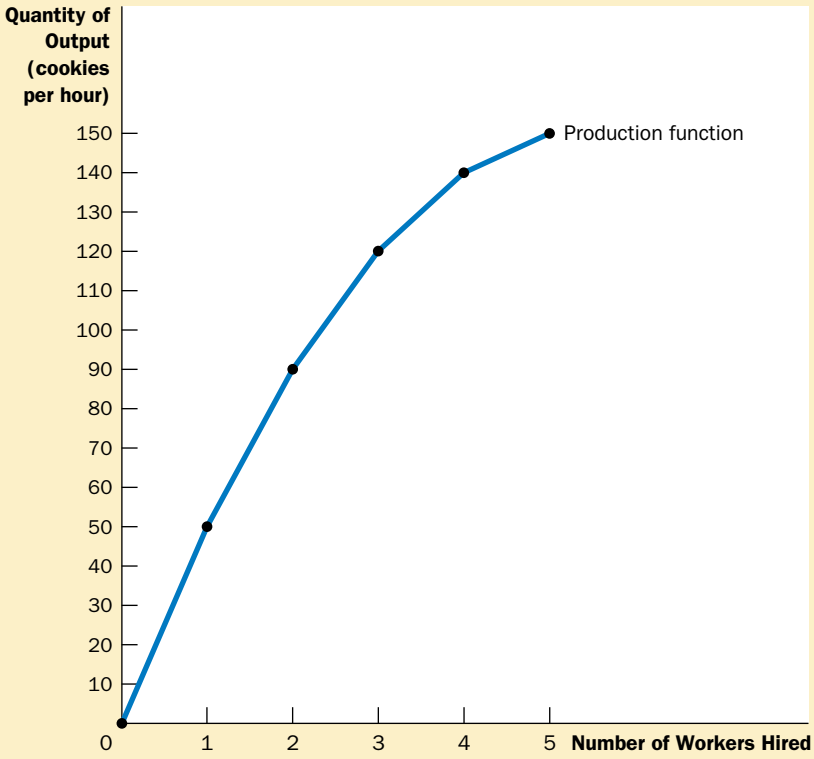

FROM THE PRODUCTION FUNCTION TO THE TOTAL-COST CURVE

The last three columns of Table 13-1 show Helen's cost of producing cookies. In this example, the cost of Helen's factory is $30 per hour, and the cost of a worker is $10 per hour. If she hires 1 worker, her total cost is $40. If she hires 2 workers, her total cost is $50, and so on. With this information, the table now shows how the number of workers Helen hires is related to the quantity of cookies she produces and to her total cost of production.

Our goal in the next several chapters is to study firms' production and pricing decisions. For this purpose, the most important relationship in Table 13-1 is between quantity produced (in the second column) and total costs (in the sixth column). Figure 13-3 graphs these two columns of data with the quantity produced on the horizontal axis and total cost on the vertical axis. This graph is called the *total-cost curve*.

Notice that the total cost gets steeper as the amount produced rises. The shape of the total-cost curve in this figure reflects the shape of the production function in Figure 13-2. Recall that when Helen's kitchen gets crowded, each additional worker adds less to the production of cookies; this property of diminishing marginal product is reflected in the flattening of the production function as the number of workers rises. But now turn this logic around: When Helen is producing a large quantity of cookies, she must have hired many workers. Because her kitchen is already crowded, producing an additional cookie is quite costly. Thus, as the quantity produced rises, the total-cost curve becomes steeper.

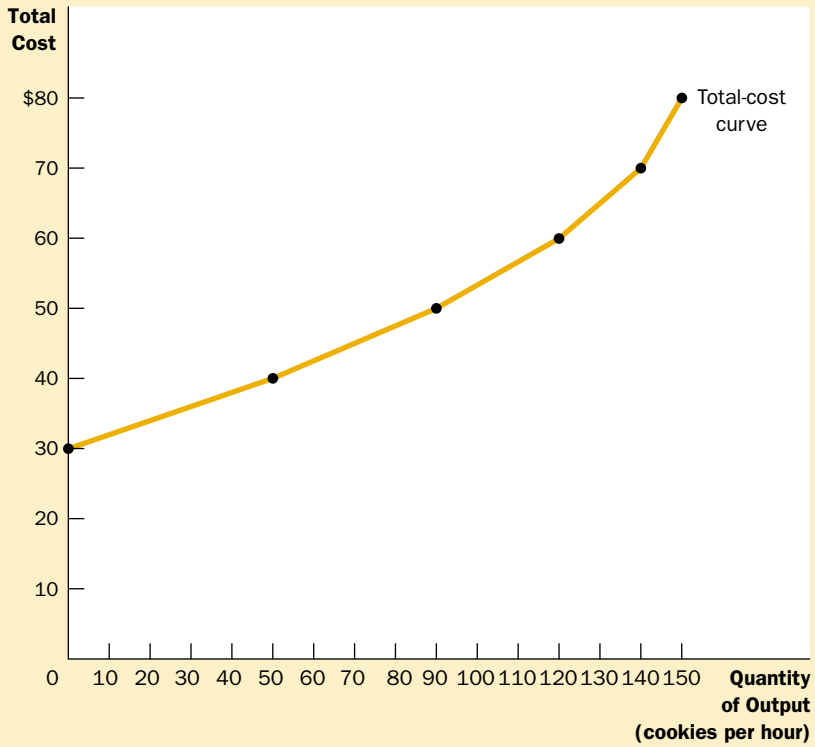

Figure 13-3

HUNGRY HELEN'S TOTAL-COST CURVE. A total-cost curve shows the relationship between the quantity of output produced and total cost of production. Here the quantity of output produced (on the horizontal axis) is from the second column in Table 13-1, and the total cost (on the vertical axis) is from the sixth column. The total-cost curve gets steeper as the quantity of output increases because of diminishing marginal product.

QUICK QUIZ: If Farmer Jones plants no seeds on his farm, he gets no harvest. If he plants 1 bag of seeds, he gets 3 bushels of wheat. If he plants 2 bags, he gets 5 bushels. If he plants 3 bags, he gets 6 bushels. A bag of seeds costs $100, and seeds are his only cost. Use these data to graph the farmer's production function and total-cost curve. Explain their shapes.

THE VARIOUS MEASURES OF COST

Our analysis of Hungry Helen's Cookie Factory demonstrated how a firm's total cost reflects its production function. From data on a firm's total cost, we can derive several related measures of cost, which will turn out to be useful when we analyze production and pricing decisions in future chapters. To see how these related measures are derived, we consider the example in Table 13-2. This table presents cost data on Helen's neighbor: Thirsty Thelma's Lemonade Stand.

The first column of the table shows the number of glasses of lemonade that Thelma might produce, ranging from 0 to 10 glasses per hour. The second column shows Thelma's total cost of producing lemonade. Figure 13-4 plots Thelma's total-cost curve. The quantity of lemonade (from the first column) is on the horizontal axis, and total cost (from the second column) is on the vertical axis. Thirsty

Figure 13-4

THIRSTY THELMA'S TOTAL-COST CURVE. Here the quantity of output produced (on the horizontal axis) is from the first column in Table 13-2, and the total cost (on the vertical axis) is from the second column. As in Figure 13-3, the total-cost curve gets steeper as the quantity of output increases because of diminishing marginal product.

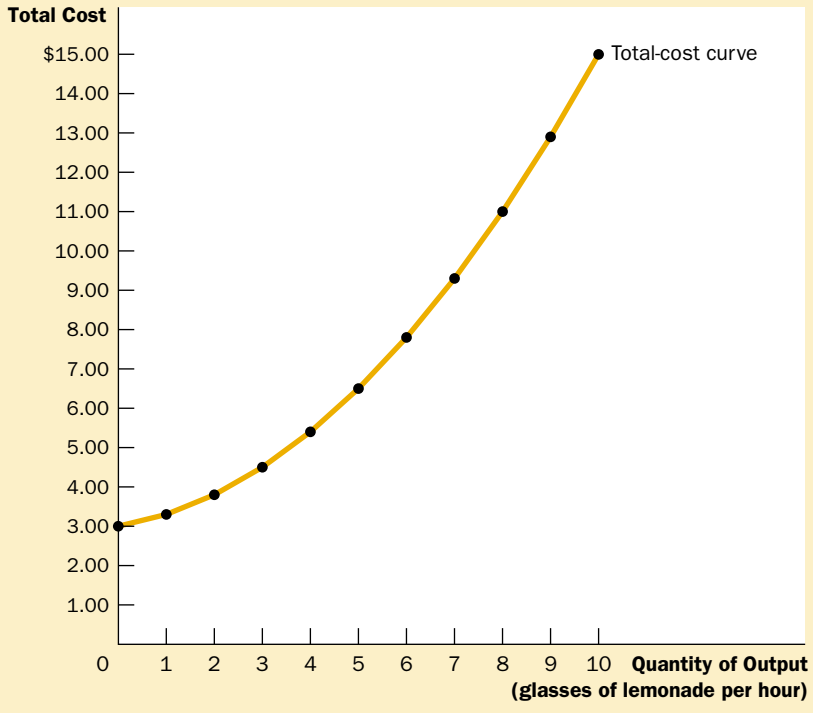

Quantity of Lemonade (glasses per hour)	Total Cost	Fixed Cost	Variable Cost	Average Fixed Cost	Average Variable Cost	Average Total Cost	Marginal Cost
0	$ 3.00	$3.00	$ 0.00	—	—	—	
1	3.30	3.00	0.30	$3.00	$0.30	$3.30	$0.30
2	3.80	3.00	0.80	1.50	0.40	1.90	0.50
3	4.50	3.00	1.50	1.00	0.50	1.50	0.70
4	5.40	3.00	2.40	0.75	0.60	1.35	0.90
5	6.50	3.00	3.50	0.60	0.70	1.30	1.10
6	7.80	3.00	4.80	0.50	0.80	1.30	1.30
7	9.30	3.00	6.30	0.43	0.90	1.33	1.50
8	11.00	3.00	8.00	0.38	1.00	1.38	1.70
9	12.90	3.00	9.90	0.33	1.10	1.43	1.90
10	15.00	3.00	12.00	0.30	1.20	1.50	2.10

THE VARIOUS MEASURES OF COST: THIRSTY THELMA'S LEMONADE STAND

Table 13-2

Thelma's total-cost curve has a shape similar to Hungry Helen's. In particular, it becomes steeper as the quantity produced rises, which (as we have discussed) reflects diminishing marginal product.

FIXED AND VARIABLE COSTS

Thelma's total cost can be divided into two types. Some costs, called **fixed costs,** do not vary with the quantity of output produced. They are incurred even if the firm produces nothing at all. Thelma's fixed costs include the rent she pays because this cost is the same regardless of how much lemonade Thelma produces. Similarly, if Thelma needs to hire a full-time bookkeeper to pay bills, regardless of the quantity of lemonade produced, the bookkeeper's salary is a fixed cost. The third column in Table 13-2 shows Thelma's fixed cost, which in this example is $3.00 per hour.

fixed costs
costs that do not vary with the quantity of output produced

Some of the firm's costs, called **variable costs,** change as the firm alters the quantity of output produced. Thelma's variable costs include the cost of lemons and sugar: The more lemonade Thelma makes, the more lemons and sugar she needs to buy. Similarly, if Thelma has to hire more workers to make more lemonade, the salaries of these workers are variable costs. The fourth column of the table shows Thelma's variable cost. The variable cost is 0 if she produces nothing, $0.30 if she produces 1 glass of lemonade, $0.80 if she produces 2 glasses, and so on.

variable costs
costs that do vary with the quantity of output produced

A firm's total cost is the sum of fixed and variable costs. In Table 13-2, total cost in the second column equals fixed cost in the third column plus variable cost in the fourth column.

AVERAGE AND MARGINAL COST

As the owner of her firm, Thelma has to decide how much to produce. A key part of this decision is how her costs will vary as she changes the level of production. In making this decision, Thelma might ask her production supervisor the following two questions about the cost of producing lemonade:

- How much does it cost to make the typical glass of lemonade?
- How much does it cost to increase production of lemonade by 1 glass?

Although at first these two questions might seem to have the same answer, they do not. Both answers will turn out to be important for understanding how firms make production decisions.

To find the cost of the typical unit produced, we would divide the firm's costs by the quantity of output it produces. For example, if the firm produces 2 glasses per hour, its total cost is \$3.80, and the cost of the typical glass is \$3.80/2, or \$1.90. Total cost divided by the quantity of output is called **average total cost.** Because total cost is just the sum of fixed and variable costs, average total cost can be expressed as the sum of average fixed cost and average variable cost. **Average fixed cost** is the fixed cost divided by the quantity of output, and **average variable cost** is the variable cost divided by the quantity of output.

average total cost
total cost divided by the quantity of output

average fixed cost
fixed costs divided by the quantity of output

average variable cost
variable costs divided by the quantity of output

marginal cost
the increase in total cost that arises from an extra unit of production

Although average total cost tells us the cost of the typical unit, it does not tell us how much total cost will change as the firm alters its level of production. The last column in Table 13-2 shows the amount that total cost rises when the firm increases production by 1 unit of output. This number is called **marginal cost.** For example, if Thelma increases production from 2 to 3 glasses, total cost rises from \$3.80 to \$4.50, so the marginal cost of the third glass of lemonade is \$4.50 minus \$3.80, or \$0.70.

It may be helpful to express these definitions mathematically. If Q stands for quantity, TC for total cost, ATC for average total cost, and MC for marginal cost, then we can then write:

$$ATC = \text{Total cost}/\text{Quantity} = TC/Q$$

and

$$MC = (\text{Change in total cost})/(\text{Change in quantity}) = \Delta TC/\Delta Q.$$

Here Δ, the Greek letter *delta*, represents the change in a variable. These equations show how average total cost and marginal cost are derived from total cost.

As we will see more fully in the next chapter, Thelma, our lemonade entrepreneur, will find the concepts of average total cost and marginal cost extremely useful when deciding how much lemonade to produce. Keep in mind, however, that these concepts do not actually give Thelma new information about her costs of production. Instead, average total cost and marginal cost express in a new way information that is already contained in her firm's total cost. *Average total cost tells us the cost of a typical unit of output if total cost is divided evenly over all the units produced. Marginal cost tells us the increase in total cost that arises from producing an additional unit of output.*

COST CURVES AND THEIR SHAPES

Just as in previous chapters we found graphs of supply and demand useful when analyzing the behavior of markets, we will find graphs of average and marginal cost useful when analyzing the behavior of firms. Figure 13-5 graphs Thelma's costs using the data from Table 13-2. The horizontal axis measures the quantity the firm produces, and the vertical axis measures marginal and average costs. The graph shows four curves: average total cost (*ATC*), average fixed cost (*AFC*), average variable cost (*AVC*), and marginal cost (*MC*).

The cost curves shown here for Thirsty Thelma's Lemonade Stand have some features that are common to the cost curves of many firms in the economy. Let's examine three features in particular: the shape of marginal cost, the shape of average total cost, and the relationship between marginal and average total cost.

Rising Marginal Cost Thirsty Thelma's marginal cost rises with the quantity of output produced. This reflects the property of diminishing marginal product. When Thelma is producing a small quantity of lemonade, she has few workers, and much of her equipment is not being used. Because she can easily put these idle resources to use, the marginal product of an extra worker is large, and the marginal cost of an extra glass of lemonade is small. By contrast, when Thelma is producing a large quantity of lemonade, her stand is crowded with workers, and most of her equipment is fully utilized. Thelma can produce more lemonade by adding workers, but these new workers have to work in crowded conditions and may have to

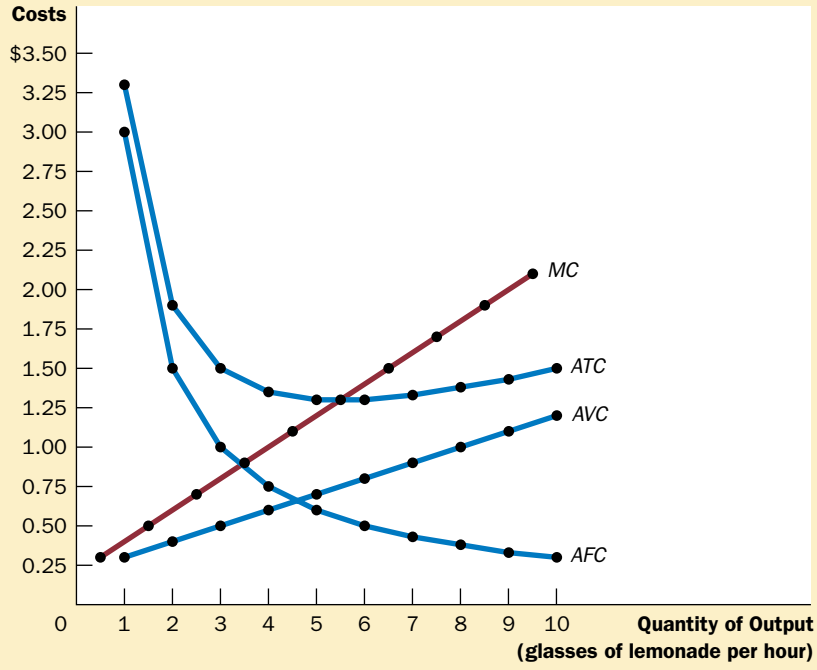

Figure 13-5

THIRSTY THELMA'S AVERAGE-COST AND MARGINAL-COST CURVES. This figure shows the average total cost (*ATC*), average fixed cost (*AFC*), average variable cost (*AVC*), and marginal cost (*MC*) for Thirsty Thelma's Lemonade Stand. All of these curves are obtained by graphing the data in Table 13-2. These cost curves show three features that are considered common: (1) Marginal cost rises with the quantity of output. (2) The average-total-cost curve is U-shaped. (3) The marginal-cost curve crosses the average-total-cost curve at the minimum of average total cost.

wait to use the equipment. Therefore, when the quantity of lemonade being produced is already high, the marginal product of an extra worker is low, and the marginal cost of an extra glass of lemonade is large.

U-Shaped Average Total Cost Thirsty Thelma's average-total-cost curve is U-shaped. To understand why this is so, remember that average total cost is the sum of average fixed cost and average variable cost. Average fixed cost always declines as output rises because the fixed cost is getting spread over a larger number of units. Average variable cost typically rises as output increases because of diminishing marginal product. Average total cost reflects the shapes of both average fixed cost and average variable cost. At very low levels of output, such as 1 or 2 glasses per hour, average total cost is high because the fixed cost is spread over only a few units. Average total cost then declines as output increases until the firm's output reaches 5 glasses of lemonade per hour, when average total cost falls to $1.30 per glass. When the firm produces more than 6 glasses, average total cost starts rising again because average variable cost rises substantially.

efficient scale
the quantity of output that minimizes average total cost

The bottom of the U-shape occurs at the quantity that minimizes average total cost. This quantity is sometimes called the **efficient scale** of the firm. For Thirsty Thelma, the efficient scale is 5 or 6 glasses of lemonade. If she produces more or less than this amount, her average total cost rises above the minimum of $1.30.

The Relationship between Marginal Cost and Average Total Cost If you look at Figure 13-5 (or back at Table 13-2), you will see something that may be surprising at first. *Whenever marginal cost is less than average total cost, average total cost is falling. Whenever marginal cost is greater than average total cost, average total cost is rising.* This feature of Thirsty Thelma's cost curves is not a coincidence from the particular numbers used in the example: It is true for all firms.

To see why, consider an analogy. Average total cost is like your cumulative grade point average. Marginal cost is like the grade in the next course you will take. If your grade in your next course is less than your grade point average, your grade point average will fall. If your grade in your next course is higher than your grade point average, your grade point average will rise. The mathematics of average and marginal costs is exactly the same as the mathematics of average and marginal grades.

This relationship between average total cost and marginal cost has an important corollary: *The marginal-cost curve crosses the average-total-cost curve at the efficient scale.* Why? At low levels of output, marginal cost is below average total cost, so average total cost is falling. But after the two curves cross, marginal cost rises above average total cost. For the reason we have just discussed, average total cost must start to rise at this level of output. Hence, this point of intersection is the minimum of average total cost. As you will see in the next chapter, this point of minimum average total cost plays a key role in the analysis of competitive firms.

TYPICAL COST CURVES

In the examples we have studied so far, the firms exhibit diminishing marginal product and, therefore, rising marginal cost at all levels of output. Yet actual firms are often a bit more complicated than this. In many firms, diminishing marginal product does not start to occur immediately after the first worker is hired. Depending on the

production process, the second or third worker might have higher marginal product than the first because a team of workers can divide tasks and work more productively than a single worker. Such firms would first experience increasing marginal product for a while before diminishing marginal product sets in.

Table 13-3 shows the cost data for such a firm, called Big Bob's Bagel Bin. These data are graphed in Figure 13-6. Panel (a) shows how total cost (*TC*) depends on the quantity produced, and panel (b) shows average total cost (*ATC*), average fixed cost (*AFC*), average variable cost (*AVC*), and marginal cost (*MC*). In the range of output from 0 to 4 bagels per hour, the firm experiences increasing marginal product, and the marginal-cost curve falls. After 5 bagels per hour, the firm starts to experience diminishing marginal product, and the marginal-cost curve starts to rise. This combination of increasing then diminishing marginal product also makes the average-variable-cost curve U-shaped.

Despite these differences from our previous example, Big Bob's cost curves share the three properties that are most important to remember:

- Marginal cost eventually rises with the quantity of output.
- The average-total-cost curve is U-shaped.
- The marginal-cost curve crosses the average-total-cost curve at the minimum of average total cost.

Quantity of Bagels (per hour)	Total Cost	Fixed Cost	Variable Cost	Average Fixed Cost	Average Variable Cost	Average Total Cost	Marginal Cost
0	$ 2.00	$2.00	$ 0.00	—	—	—	
1	3.00	2.00	1.00	$2.00	$1.00	$3.00	$1.00
2	3.80	2.00	1.80	1.00	0.90	1.90	0.80
3	4.40	2.00	2.40	0.67	0.80	1.47	0.60
4	4.80	2.00	2.80	0.50	0.70	1.20	0.40
5	5.20	2.00	3.20	0.40	0.64	1.04	0.40
6	5.80	2.00	3.80	0.33	0.63	0.96	0.60
7	6.60	2.00	4.60	0.29	0.66	0.95	0.80
8	7.60	2.00	5.60	0.25	0.70	0.95	1.00
9	8.80	2.00	6.80	0.22	0.76	0.98	1.20
10	10.20	2.00	8.20	0.20	0.82	1.02	1.40
11	11.80	2.00	9.80	0.18	0.89	1.07	1.60
12	13.60	2.00	11.60	0.17	0.97	1.14	1.80
13	15.60	2.00	13.60	0.15	1.05	1.20	2.00
14	17.80	2.00	15.80	0.14	1.13	1.27	2.20

THE VARIOUS MEASURES OF COST: BIG BOB'S BAGEL BIN

Table 13-3

Figure 13-6

BIG BOB'S COST CURVES. Many firms, like Big Bob's Bagel Bin, experience increasing marginal product before diminishing marginal product and, therefore, have cost curves like those in this figure. Panel (a) shows how total cost (*TC*) depends on the quantity produced. Panel (b) shows how average total cost (*ATC*), average fixed cost (*AFC*), average variable cost (*AVC*), and marginal cost (*MC*) depend on the quantity produced. These curves are derived by graphing the data from Table 13-3. Notice that marginal cost and average variable cost fall for a while before starting to rise.

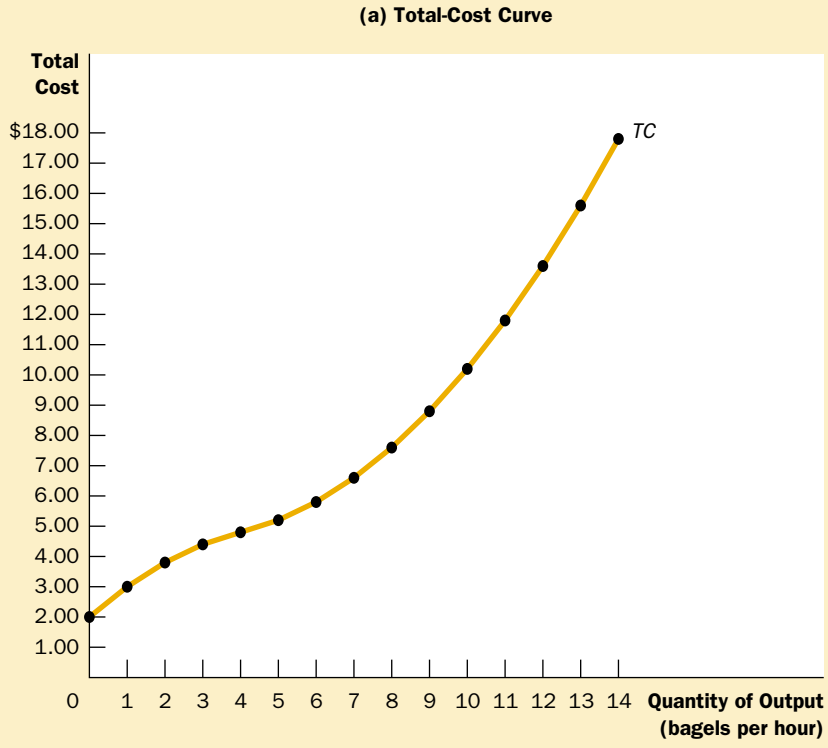

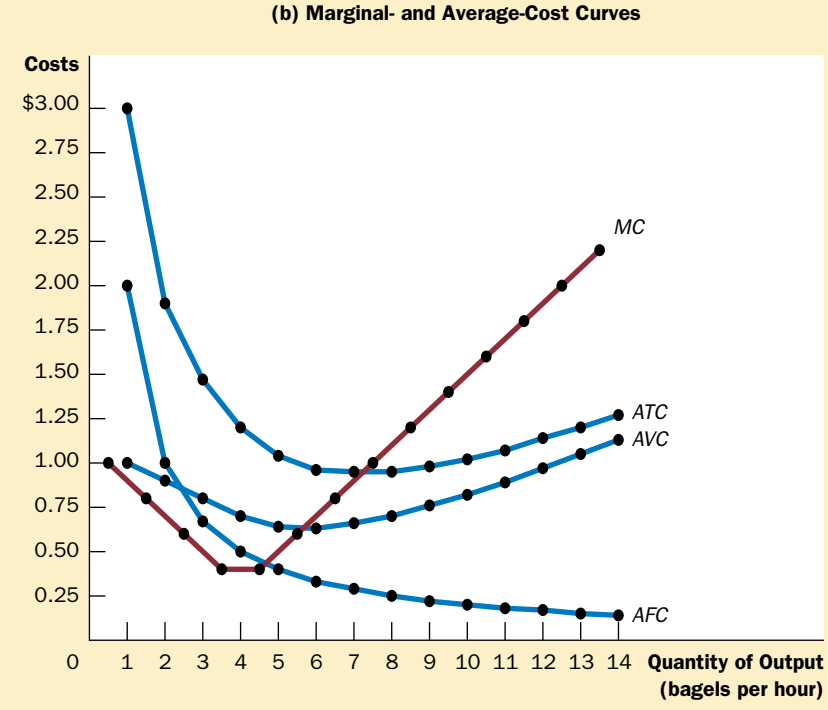

QUICK QUIZ: Suppose Honda's total cost of producing 4 cars is $225,000 and its total cost of producing 5 cars is $250,000. What is the average total cost of producing 5 cars? What is the marginal cost of the fifth car? ◆ Draw the marginal-cost curve and the average-total-cost curve for a typical firm, and explain why these curves cross where they do.

COSTS IN THE SHORT RUN AND IN THE LONG RUN

We noted at the beginning of this chapter that a firm's costs might depend on the time horizon being examined. Let's examine more precisely why this might be the case.

THE RELATIONSHIP BETWEEN SHORT-RUN AND LONG-RUN AVERAGE TOTAL COST

For many firms, the division of total costs between fixed and variable costs depends on the time horizon. Consider, for instance, a car manufacturer, such as Ford Motor Company. Over a period of only a few months, Ford cannot adjust the number or sizes of its car factories. The only way it can produce additional cars is to hire more workers at the factories it already has. The cost of these factories is, therefore, a fixed cost in the short run. By contrast, over a period of several years, Ford can expand the size of its factories, build new factories, or close old ones. Thus, the cost of its factories is a variable cost in the long run.

Because many decisions are fixed in the short run but variable in the long run, a firm's long-run cost curves differ from its short-run cost curves. Figure 13-7 shows an example. The figure presents three short-run average-total-cost curves—for a small, medium, and large factory. It also presents the long-run average-total-cost curve. As the firm moves along the long-run curve, it is adjusting the size of the factory to the quantity of production.

This graph shows how short-run and long-run costs are related. The long-run average-total-cost curve is a much flatter U-shape than the short-run average-total-cost curve. In addition, all the short-run curves lie on or above the long-run curve. These properties arise because of the greater flexibility firms have in the long run. In essence, in the long run, the firm gets to choose which short-run curve it wants to use. But in the short run, it has to use whatever short-run curve it chose in the past.

The figure shows an example of how a change in production alters costs over different time horizons. When Ford wants to increase production from 1,000 to 1,200 cars per day, it has no choice in the short run but to hire more workers at its existing medium-sized factory. Because of diminishing marginal product, average total cost rises from $10,000 to $12,000 per car. In the long run, however, Ford can expand both the size of the factory and its workforce, and average total cost remains at $10,000.

How long does it take for a firm to get to the long run? The answer depends on the firm. It can take a year or longer for a major manufacturing firm, such as a

Figure 13-7

AVERAGE TOTAL COST IN THE SHORT AND LONG RUNS. Because fixed costs are variable in the long run, the average-total-cost curve in the short run differs from the average-total-cost curve in the long run.

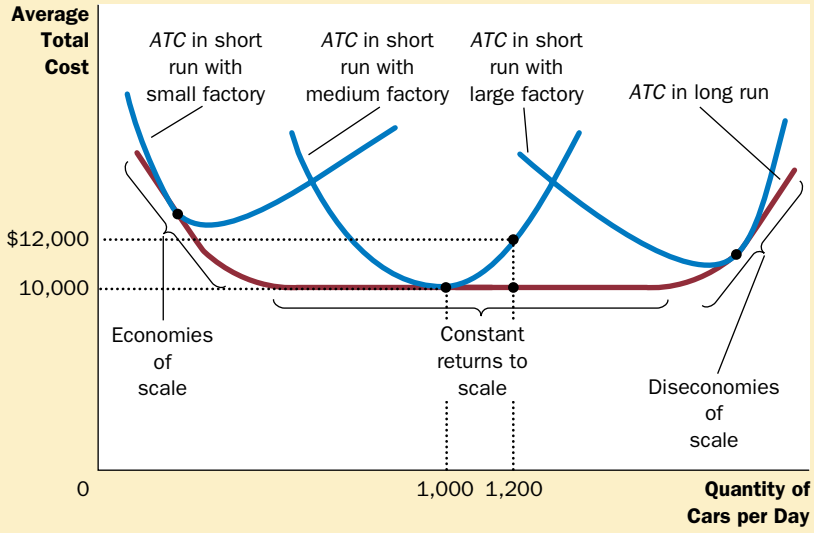

car company, to build a larger factory. By contrast, a person running a lemonade stand can go and buy a larger pitcher within an hour or less. There is, therefore, no single answer about how long it takes a firm to adjust its production facilities.

ECONOMIES AND DISECONOMIES OF SCALE

economies of scale
the property whereby long-run average total cost falls as the quantity of output increases

diseconomies of scale
the property whereby long-run average total cost rises as the quantity of output increases

constant returns to scale
the property whereby long-run average total cost stays the same as the quantity of output changes

The shape of the long-run average-total-cost curve conveys important information about the technology for producing a good. When long-run average total cost declines as output increases, there are said to be **economies of scale.** When long-run average total cost rises as output increases, there are said to be **diseconomies of scale.** When long-run average total cost does not vary with the level of output, there are said to be **constant returns to scale.** In this example, Ford has economies of scale at low levels of output, constant returns to scale at intermediate levels of output, and diseconomies of scale at high levels of output.

What might cause economies or diseconomies of scale? Economies of scale often arise because higher production levels allow *specialization* among workers, which permits each worker to become better at his or her assigned tasks. For instance, modern assembly-line production requires a large number of workers. If Ford were producing only a small quantity of cars, it could not take advantage of this approach and would have higher average total cost. Diseconomies of scale can arise because of *coordination problems* that are inherent in any large organization. The more cars Ford produces, the more stretched the management team becomes, and the less effective the managers become at keeping costs down.

This analysis shows why long-run average-total-cost curves are often U-shaped. At low levels of production, the firm benefits from increased size because it can take advantage of greater specialization. Coordination problems,

FYI

Lessons from a Pin Factory

"Jack of all trades, master of none." This well-known adage helps explain why firms sometimes experience economies of scale. A person who tries to do everything usually ends up doing nothing very well. If a firm wants its workers to be as productive as they can be, it is often best to give them a limited task that they can master. But this is possible only if a firm employs a large number of workers and produces a large quantity of output.

In his celebrated book, *An Inquiry into the Nature and Causes of the Wealth of Nations*, Adam Smith described an example of this based on a visit he made to a pin factory. Smith was impressed by the specialization among the workers that he observed and the resulting economies of scale. He wrote,

> "One man draws out the wire, another straightens it, a third cuts it, a fourth points it, a fifth grinds it at the top for receiving the head; to make the head requires two or three distinct operations; to put it on is a peculiar business; to whiten it is another; it is even a trade by itself to put them into paper."

Smith reported that because of this specialization, the pin factory produced thousands of pins per worker every day. He conjectured that if the workers had chosen to work separately, rather than as a team of specialists, "they certainly could not each of them make twenty, perhaps not one pin a day." In other words, because of specialization, a large pin factory could achieve higher output per worker and lower average cost per pin than a small pin factory.

The specialization that Smith observed in the pin factory is prevalent in the modern economy. If you want to build a house, for instance, you could try to do all the work yourself. But most people turn to a builder, who in turn hires carpenters, plumbers, electricians, painters, and many other types of worker. These workers specialize in particular jobs, and this allows them to become better at their jobs than if they were generalists. Indeed, the use of specialization to achieve economies of scale is one reason modern societies are as prosperous as they are.

meanwhile, are not yet acute. By contrast, at high levels of production, the benefits of specialization have already been realized, and coordination problems become more severe as the firm grows larger. Thus, long-run average total cost is falling at low levels of production because of increasing specialization and rising at high levels of production because of increasing coordination problems.

QUICK QUIZ: If Boeing produces 9 jets per month, its long-run total cost is $9.0 million per month. If it produces 10 jets per month, its long-run total cost is $9.5 million per month. Does Boeing exhibit economies or diseconomies of scale?

CONCLUSION

The purpose of this chapter has been to develop some tools that we can use to study how firms make production and pricing decisions. You should now understand what economists mean by the term *costs* and how costs vary with the quantity of output a firm produces. To refresh your memory, Table 13-4 summarizes some of the definitions we have encountered.

Table 13-4

THE MANY TYPES OF COST: A SUMMARY

TERM	DEFINITION	MATHEMATICAL DESCRIPTION
Explicit costs	Costs that require an outlay of money by the firm	—
Implicit costs	Costs that do not require an outlay of money by the firm	—
Fixed costs	Costs that do not vary with the quantity of output produced	FC
Variable costs	Costs that do vary with the quantity of output produced	VC
Total cost	The market value of all the inputs that a firm uses in production	$TC = FC + VC$
Average fixed cost	Fixed costs divided by the quantity of output	$AFC = FC/Q$
Average variable cost	Variable costs divided by the quantity of output	$AVC = VC/Q$
Average total cost	Total cost divided by the quantity of output	$ATC = TC/Q$
Marginal cost	The increase in total cost that arises from an extra unit of production	$MC = \Delta TC/\Delta Q$

By themselves, of course, a firm's cost curves do not tell us what decisions the firm will make. But they are an important component of that decision, as we will begin to see in the next chapter.

Summary

- The goal of firms is to maximize profit, which equals total revenue minus total cost.
- When analyzing a firm's behavior, it is important to include all the opportunity costs of production. Some of the opportunity costs, such as the wages a firm pays its workers, are explicit. Other opportunity costs, such as the wages the firm owner gives up by working in the firm rather than taking another job, are implicit.
- A firm's costs reflect its production process. A typical firm's production function gets flatter as the quantity of an input increases, displaying the property of diminishing marginal product. As a result, a firm's total-cost curve gets steeper as the quantity produced rises.
- A firm's total costs can be divided between fixed costs and variable costs. Fixed costs are costs that do not change when the firm alters the quantity of output produced. Variable costs are costs that do change when the firm alters the quantity of output produced.
- From a firm's total cost, two related measures of cost are derived. Average total cost is total cost divided by the quantity of output. Marginal cost is the amount by which total cost would rise if output were increased by 1 unit.
- When analyzing firm behavior, it is often useful to graph average total cost and marginal cost. For a typical firm, marginal cost rises with the quantity of output. Average total cost first falls as output increases and then

rises as output increases further. The marginal-cost curve always crosses the average-total-cost curve at the minimum of average total cost.

- A firm's costs often depend on the time horizon being considered. In particular, many costs are fixed in the short run but variable in the long run. As a result, when the firm changes its level of production, average total cost may rise more in the short run than in the long run.

Key Concepts

total revenue, p. 270
total cost, p. 270
profit, p. 270
explicit costs, p. 271
implicit costs, p. 271
economic profit, p. 272
accounting profit, p. 272
production function, p. 273
marginal product, p. 273
diminishing marginal product, p. 273
fixed costs, p. 277
variable costs, p. 277
average total cost, p. 278
average fixed cost, p. 278
average variable cost, p. 278
marginal cost, p. 278
efficient scale, p. 280
economies of scale, p. 284
diseconomies of scale, p. 284
constant returns to scale, p. 284

Questions for Review

1. What is the relationship between a firm's total revenue, profit, and total cost?
2. Give an example of an opportunity cost that an accountant might not count as a cost. Why would the accountant ignore this cost?
3. What is marginal product, and what does it mean if it is diminishing?
4. Draw a production function that exhibits diminishing marginal product of labor. Draw the associated total-cost curve. (In both cases, be sure to label the axes.) Explain the shapes of the two curves you have drawn.
5. Define total cost, average total cost, and marginal cost. How are they related?
6. Draw the marginal-cost and average-total-cost curves for a typical firm. Explain why the curves have the shapes that they do and why they cross where they do.
7. How and why does a firm's average-total-cost curve differ in the short run and in the long run?
8. Define *economies of scale* and explain why they might arise. Define *diseconomies of scale* and explain why they might arise.

Problems and Applications

1. This chapter discusses many types of costs: opportunity cost, total cost, fixed cost, variable cost, average total cost, and marginal cost. Fill in the type of cost that best completes the phrases below:
   a. The true cost of taking some action is its _____.
   b. _____ is falling when marginal cost is below it, and rising when marginal cost is above it.
   c. A cost that does not depend on the quantity produced is a _____.
   d. In the ice-cream industry in the short run, _____ includes the cost of cream and sugar, but not the cost of the factory.
   e. Profits equal total revenue less _____.
   f. The cost of producing an extra unit of output is _____.
2. Your aunt is thinking about opening a hardware store. She estimates that it would cost $500,000 per year to rent the location and buy the stock. In addition, she would have to quit her $50,000 per year job as an accountant.
   a. Define opportunity cost.
   b. What is your aunt's opportunity cost of running a hardware store for a year? If your aunt thought she could sell $510,000 worth of merchandise in a year, should she open the store? Explain.

3. Suppose that your college charges you separately for tuition and for room and board.
   a. What is a cost of attending college that is not an opportunity cost?
   b. What is an explicit opportunity cost of attending college?
   c. What is an implicit opportunity cost of attending college?
4. A commercial fisherman notices the following relationship between hours spent fishing and the quantity of fish caught:

Hours	Quantity of Fish (in pounds)
0	0
1	10
2	18
3	24
4	28
5	30

   a. What is the marginal product of each hour spent fishing?
   b. Use these data to graph the fisherman's production function. Explain its shape.
   c. The fisherman has a fixed cost of $10 (his pole). The opportunity cost of his time is $5 per hour. Graph the fisherman's total-cost curve. Explain its shape.
5. Nimbus, Inc., makes brooms and then sells them door-to-door. Here is the relationship between the number of workers and Nimbus's output in a given day:

Workers	Output	Marginal Product	Total Cost	Average Total Cost	Marginal Cost
0	0		_____	_____	
		_____			_____
1	20		_____	_____	
		_____			_____
2	50		_____	_____	
		_____			_____
3	90		_____	_____	
		_____			_____
4	120		_____	_____	
		_____			_____
5	140		_____	_____	
		_____			_____
6	150		_____	_____	
		_____			_____
7	155		_____	_____	

   a. Fill in the column of marginal products. What pattern do you see? How might you explain it?
   b. A worker costs $100 a day, and the firm has fixed costs of $200. Use this information to fill in the column for total cost.
   c. Fill in the column for average total cost. (Recall that $ATC = TC/Q$.) What pattern do you see?
   d. Now fill in the column for marginal cost. (Recall that $MC = \Delta TC/\Delta Q$.) What pattern do you see?
   e. Compare the column for marginal product and the column for marginal cost. Explain the relationship.
   f. Compare the column for average total cost and the column for marginal cost. Explain the relationship.
6. Suppose that you and your roommate have started a bagel delivery service on campus. List some of your fixed costs and describe why they are fixed. List some of your variable costs and describe why they are variable.
7. Consider the following cost information for a pizzeria:

Q (dozens)	Total Cost	Variable Cost
0	$300	$ 0
1	350	50
2	390	90
3	420	120
4	450	150
5	490	190
6	540	240

   a. What is the pizzeria's fixed cost?
   b. Construct a table in which you calculate the marginal cost per dozen pizzas using the information on total cost. Also calculate the marginal cost per dozen pizzas using the information on variable cost. What is the relationship between these sets of numbers? Comment.
8. You are thinking about setting up a lemonade stand. The stand itself costs $200. The ingredients for each cup of lemonade cost $0.50.
   a. What is your fixed cost of doing business? What is your variable cost per cup?
   b. Construct a table showing your total cost, average total cost, and marginal cost for output levels varying from zero to 10 gallons. (Hint: There are 16 cups in a gallon.) Draw the three cost curves.
9. Your cousin Vinnie owns a painting company with fixed costs of $200 and the following schedule for variable costs:

	Quantity of Houses Painted per Month						
	1	2	3	4	5	6	7
Variable costs	$10	$20	$40	$80	$160	$320	$640

Calculate average fixed cost, average variable cost, and average total cost for each quantity. What is the efficient scale of the painting company?

10. Healthy Harry's Juice Bar has the following cost schedules:

Q (vats)	Variable Cost	Total Cost
0	$ 0	$ 30
1	10	40
2	25	55
3	45	75
4	70	100
5	100	130
6	135	165

a. Calculate average variable cost, average total cost, and marginal cost for each quantity.
b. Graph all three curves. What is the relationship between the marginal-cost curve and the average-total-cost curve? Between the marginal-cost curve and the average-variable-cost curve? Explain.

11. Consider the following table of long-run total cost for three different firms:

	Quantity						
	1	2	3	4	5	6	7
Firm A	$60	$70	$80	$90	$100	$110	$120
Firm B	11	24	39	56	75	96	119
Firm C	21	34	49	66	85	106	129

Does each of these firms experience economies of scale or diseconomies of scale?

14

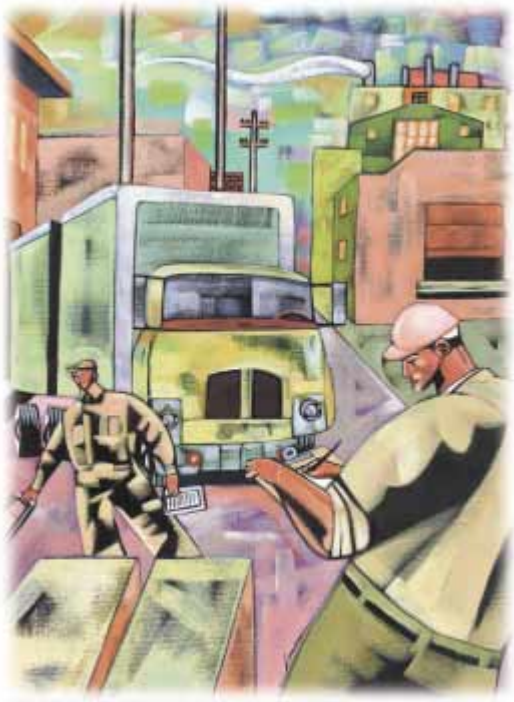

FIRMS IN COMPETITIVE MARKETS

IN THIS CHAPTER YOU WILL . . .

- *Learn what characteristics make a market competitive*
- *Examine how competitive firms decide how much output to produce*
- *Examine how competitive firms decide when to shut down production temporarily*
- *Examine how competitive firms decide whether to exit or enter a market*
- *See how firm behavior determines a market's short-run and long-run supply curves*

If your local gas station raised the price it charges for gasoline by 20 percent, it would see a large drop in the amount of gasoline it sold. Its customers would quickly switch to buying their gasoline at other gas stations. By contrast, if your local water company raised the price of water by 20 percent, it would see only a small decrease in the amount of water it sold. People might water their lawns less often and buy more water-efficient shower heads, but they would be hard-pressed to reduce water consumption greatly and would be unlikely to find another supplier. The difference between the gasoline market and the water market is obvious: There are many firms pumping gasoline, but there is only one firm pumping water. As you might expect, this difference in market structure shapes the pricing and production decisions of the firms that operate in these markets.

In this chapter we examine the behavior of competitive firms, such as your local gas station. You may recall that a market is competitive if each buyer and seller

is small compared to the size of the market and, therefore, has little ability to influence market prices. By contrast, if a firm can influence the market price of the good it sells, it is said to have *market power.* In the three chapters that follow this one, we examine the behavior of firms with market power, such as your local water company.

Our analysis of competitive firms in this chapter will shed light on the decisions that lie behind the supply curve in a competitive market. Not surprisingly, we will find that a market supply curve is tightly linked to firms' costs of production. (Indeed, this general insight should be familiar to you from our analysis in Chapter 7.) But among a firm's various costs—fixed, variable, average, and marginal—which ones are most relevant for its decision about the quantity to supply? We will see that all these measures of cost play important and interrelated roles.

WHAT IS A COMPETITIVE MARKET?

Our goal in this chapter is to examine how firms make production decisions in competitive markets. As a background for this analysis, we begin by considering what a competitive market is.

THE MEANING OF COMPETITION

competitive market
a market with many buyers and sellers trading identical products so that each buyer and seller is a price taker

Although we have already discussed the meaning of competition in Chapter 4, let's review the lesson briefly. A **competitive market,** sometimes called a *perfectly competitive market,* has two characteristics:

- There are many buyers and many sellers in the market.
- The goods offered by the various sellers are largely the same.

As a result of these conditions, the actions of any single buyer or seller in the market have a negligible impact on the market price. Each buyer and seller takes the market price as given.

An example is the market for milk. No single buyer of milk can influence the price of milk because each buyer purchases a small amount relative to the size of the market. Similarly, each seller of milk has limited control over the price because many other sellers are offering milk that is essentially identical. Because each seller can sell all he wants at the going price, he has little reason to charge less, and if he charges more, buyers will go elsewhere. Buyers and sellers in competitive markets must accept the price the market determines and, therefore, are said to be *price takers*.

In addition to the foregoing two conditions for competition, there is a third condition sometimes thought to characterize perfectly competitive markets:

- Firms can freely enter or exit the market.

If, for instance, anyone can decide to start a dairy farm, and if any existing dairy farmer can decide to leave the dairy business, then the dairy industry would satisfy this condition. It should be noted that much of the analysis of competitive firms does not rely on the assumption of free entry and exit because this condition is not necessary for firms to be price takers. But as we will see later in this chapter, entry and exit are often powerful forces shaping the long-run outcome in competitive markets.

THE REVENUE OF A COMPETITIVE FIRM

A firm in a competitive market, like most other firms in the economy, tries to maximize profit, which equals total revenue minus total cost. To see how it does this, we first consider the revenue of a competitive firm. To keep matters concrete, let's consider a specific firm: the Smith Family Dairy Farm.

The Smith Farm produces a quantity of milk Q and sells each unit at the market price P. The farm's total revenue is $P \times Q$. For example, if a gallon of milk sells for \$6 and the farm sells 1,000 gallons, its total revenue is \$6,000.

Because the Smith Farm is small compared to the world market for milk, it takes the price as given by market conditions. This means, in particular, that the price of milk does not depend on the quantity of output that the Smith Farm produces and sells. If the Smiths double the amount of milk they produce, the price of milk remains the same, and their total revenue doubles. As a result, total revenue is proportional to the amount of output.

Table 14-1 shows the revenue for the Smith Family Dairy Farm. The first two columns show the amount of output the farm produces and the price at which it sells its output. The third column is the farm's total revenue. The table assumes that the price of milk is \$6 a gallon, so total revenue is simply \$6 times the number of gallons.

Just as the concepts of *average* and *marginal* were useful in the preceding chapter when analyzing costs, they are also useful when analyzing revenue. To see what these concepts tell us, consider these two questions:

Table 14-1

TOTAL, AVERAGE, AND MARGINAL REVENUE FOR A COMPETITIVE FIRM

QUANTITY (IN GALLONS)	PRICE	TOTAL REVENUE	AVERAGE REVENUE	MARGINAL REVENUE
(Q)	(P)	($TR = P \times Q$)	($AR = TR/Q$)	($MR = \Delta TR/\Delta Q$)
1	\$6	\$ 6	\$6	
2	6	12	6	\$6
3	6	18	6	6
4	6	24	6	6
5	6	30	6	6
6	6	36	6	6
7	6	42	6	6
8	6	48	6	6

- How much revenue does the farm receive for the typical gallon of milk?
- How much additional revenue does the farm receive if it increases production of milk by 1 gallon?

The last two columns in Table 14-1 answer these questions.

average revenue
total revenue divided by the quantity sold

The fourth column in the table shows **average revenue,** which is total revenue (from the third column) divided by the amount of output (from the first column). Average revenue tells us how much revenue a firm receives for the typical unit sold. In Table 14-1, you can see that average revenue equals $6, the price of a gallon of milk. This illustrates a general lesson that applies not only to competitive firms but to other firms as well. Total revenue is the price times the quantity ($P \times Q$), and average revenue is total revenue ($P \times Q$) divided by the quantity (Q). Therefore, *for all firms, average revenue equals the price of the good.*

marginal revenue
the change in total revenue from an additional unit sold

The fifth column shows **marginal revenue,** which is the change in total revenue from the sale of each additional unit of output. In Table 14-1, marginal revenue equals $6, the price of a gallon of milk. This result illustrates a lesson that applies only to competitive firms. Total revenue is $P \times Q$, and P is fixed for a competitive firm. Therefore, when Q rises by 1 unit, total revenue rises by P dollars. *For competitive firms, marginal revenue equals the price of the good.*

QUICK QUIZ: When a competitive firm doubles the amount it sells, what happens to the price of its output and its total revenue?

PROFIT MAXIMIZATION AND THE COMPETITIVE FIRM'S SUPPLY CURVE

The goal of a competitive firm is to maximize profit, which equals total revenue minus total cost. We have just discussed the firm's revenue, and in the last chapter we discussed the firm's costs. We are now ready to examine how the firm maximizes profit and how that decision leads to its supply curve.

A SIMPLE EXAMPLE OF PROFIT MAXIMIZATION

Let's begin our analysis of the firm's supply decision with the example in Table 14-2. In the first column of the table is the number of gallons of milk the Smith Family Dairy Farm produces. The second column shows the farm's total revenue, which is $6 times the number of gallons. The third column shows the farm's total cost. Total cost includes fixed costs, which are $3 in this example, and variable costs, which depend on the quantity produced.

The fourth column shows the farm's profit, which is computed by subtracting total cost from total revenue. If the farm produces nothing, it has a loss of $3. If it produces 1 gallon, it has a profit of $1. If it produces 2 gallons, it has a profit of $4, and so on. To maximize profit, the Smith Farm chooses the quantity that makes profit as large as possible. In this example, profit is maximized when the farm produces 4 or 5 gallons of milk, when the profit is $7.

Quantity (in gallons)	Total Revenue	Total Cost	Profit	Marginal Revenue	Marginal Cost
(*Q*)	(*TR*)	(*TC*)	(*TR* − *TC*)	(*MR* = Δ*TR*/Δ*Q*)	(*MC* = Δ*TC*/Δ*Q*)
0	$ 0	$ 3	−$3		
1	6	5	1	$6	$2
2	12	8	4	6	3
3	18	12	6	6	4
4	24	17	7	6	5
5	30	23	7	6	6
6	36	30	6	6	7
7	42	38	4	6	8
8	48	47	1	6	9

Profit Maximization: A Numerical Example

Table 14-2

There is another way to look at the Smith Farm's decision: The Smiths can find the profit-maximizing quantity by comparing the marginal revenue and marginal cost from each unit produced. The last two columns in Table 14-2 compute marginal revenue and marginal cost from the changes in total revenue and total cost. The first gallon of milk the farm produces has a marginal revenue of $6 and a marginal cost of $2; hence, producing that gallon increases profit by $4 (from −$3 to $1). The second gallon produced has a marginal revenue of $6 and a marginal cost of $3, so that gallon increases profit by $3 (from $1 to $4). As long as marginal revenue exceeds marginal cost, increasing the quantity produced raises profit. Once the Smith Farm has reached 5 gallons of milk, however, the situation is very different. The sixth gallon would have marginal revenue of $6 and marginal cost of $7, so producing it would reduce profit by $1 (from $7 to $6). As a result, the Smiths would not produce beyond 5 gallons.

One of the *Ten Principles of Economics* in Chapter 1 is that rational people think at the margin. We now see how the Smith Family Dairy Farm can apply this principle. If marginal revenue is greater than marginal cost—as it is at 1, 2, or 3 gallons—the Smiths should increase the production of milk. If marginal revenue is less than marginal cost—as it is at 6, 7, or 8 gallons—the Smiths should decrease production. If the Smiths think at the margin and make incremental adjustments to the level of production, they are naturally led to produce the profit-maximizing quantity.

THE MARGINAL-COST CURVE AND THE FIRM'S SUPPLY DECISION

To extend this analysis of profit maximization, consider the cost curves in Figure 14-1. These cost curves have the three features that, as we discussed in Chapter 13, are thought to describe most firms: The marginal-cost curve (*MC*) is upward sloping. The average-total-cost curve (*ATC*) is U-shaped. And the marginal-cost curve

Figure 14-1

PROFIT MAXIMIZATION FOR A COMPETITIVE FIRM. This figure shows the marginal-cost curve (*MC*), the average-total-cost curve (*ATC*), and the average-variable-cost curve (*AVC*). It also shows the market price (*P*), which equals marginal revenue (*MR*) and average revenue (*AR*). At the quantity Q_1, marginal revenue MR_1 exceeds marginal cost MC_1, so raising production increases profit. At the quantity Q_2, marginal cost MC_2 is above marginal revenue MR_2, so reducing production increases profit. The profit-maximizing quantity Q_{MAX} is found where the horizontal price line intersects the marginal-cost curve.

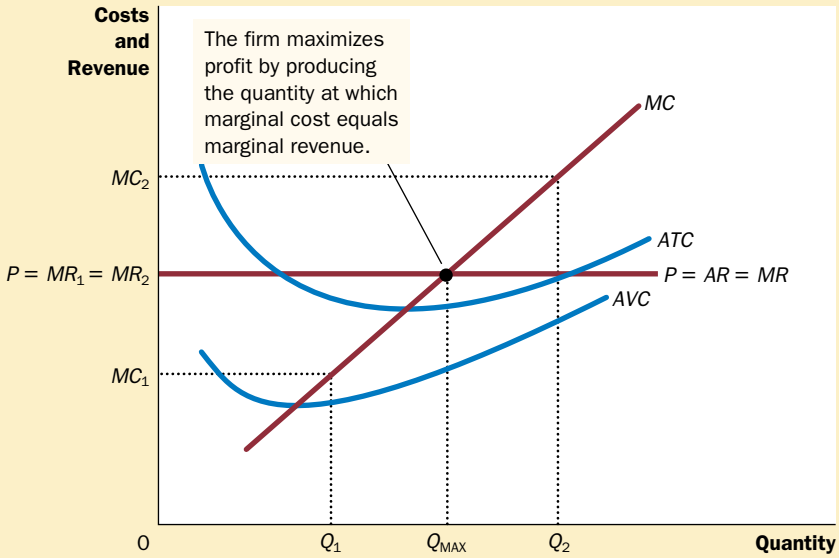

crosses the average-total-cost curve at the minimum of average total cost. The figure also shows a horizontal line at the market price (*P*). The price line is horizontal because the firm is a price taker: The price of the firm's output is the same regardless of the quantity that the firm decides to produce. Keep in mind that, for a competitive firm, the firm's price equals both its average revenue (*AR*) and its marginal revenue (*MR*).

We can use Figure 14-1 to find the quantity of output that maximizes profit. Imagine that the firm is producing at Q_1. At this level of output, marginal revenue is greater than marginal cost. That is, if the firm raised its level of production and sales by 1 unit, the additional revenue (MR_1) would exceed the additional costs (MC_1). Profit, which equals total revenue minus total cost, would increase. Hence, if marginal revenue is greater than marginal cost, as it is at Q_1, the firm can increase profit by increasing production.

A similar argument applies when output is at Q_2. In this case, marginal cost is greater than marginal revenue. If the firm reduced production by 1 unit, the costs saved (MC_2) would exceed the revenue lost (MR_2). Therefore, if marginal revenue is less than marginal cost, as it is at Q_2, the firm can increase profit by reducing production.

Where do these marginal adjustments to level of production end? Regardless of whether the firm begins with production at a low level (such as Q_1) or at a high level (such as Q_2), the firm will eventually adjust production until the quantity produced reaches Q_{MAX}. This analysis shows a general rule for profit maximization: *At the profit-maximizing level of output, marginal revenue and marginal cost are exactly equal.*

We can now see how the competitive firm decides the quantity of its good to supply to the market. Because a competitive firm is a price taker, its marginal

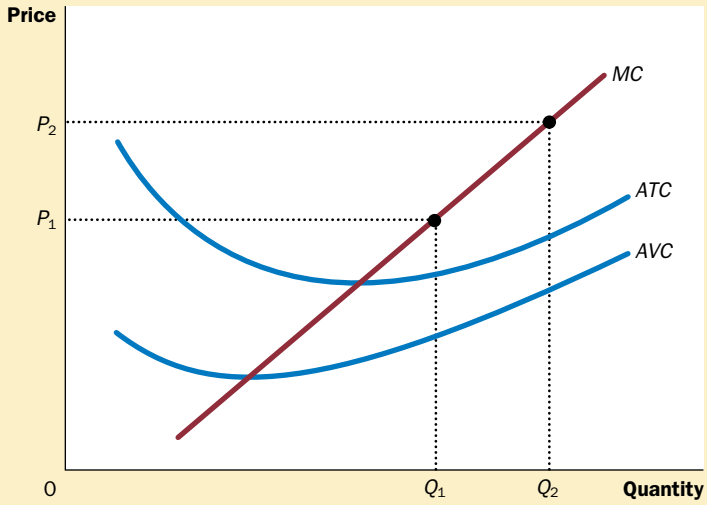

Figure 14-2

MARGINAL COST AS THE COMPETITIVE FIRM'S SUPPLY CURVE. An increase in the price from P_1 to P_2 leads to an increase in the firm's profit-maximizing quantity from Q_1 to Q_2. Because the marginal-cost curve shows the quantity supplied by the firm at any given price, it is the firm's supply curve.

revenue equals the market price. For any given price, the competitive firm's profit-maximizing quantity of output is found by looking at the intersection of the price with the marginal-cost curve. In Figure 14-1, that quantity of output is Q_{MAX}.

Figure 14-2 shows how a competitive firm responds to an increase in the price. When the price is P_1, the firm produces quantity Q_1, which is the quantity that equates marginal cost to the price. When the price rises to P_2, the firm finds that marginal revenue is now higher than marginal cost at the previous level of output, so the firm increases production. The new profit-maximizing quantity is Q_2, at which marginal cost equals the new higher price. *In essence, because the firm's marginal-cost curve determines the quantity of the good the firm is willing to supply at any price, it is the competitive firm's supply curve.*

THE FIRM'S SHORT-RUN DECISION TO SHUT DOWN

So far we have been analyzing the question of how much a competitive firm will produce. In some circumstances, however, the firm will decide to shut down and not produce anything at all.

Here we should distinguish between a temporary shutdown of a firm and the permanent exit of a firm from the market. A *shutdown* refers to a short-run decision not to produce anything during a specific period of time because of current market conditions. *Exit* refers to a long-run decision to leave the market. The short-run and long-run decisions differ because most firms cannot avoid their fixed costs in the short run but can do so in the long run. That is, a firm that shuts down temporarily still has to pay its fixed costs, whereas a firm that exits the market saves both its fixed and its variable costs.

For example, consider the production decision that a farmer faces. The cost of the land is one of the farmer's fixed costs. If the farmer decides not to produce any

crops one season, the land lies fallow, and he cannot recover this cost. When making the short-run decision whether to shut down for a season, the fixed cost of land is said to be a *sunk cost*. By contrast, if the farmer decides to leave farming altogether, he can sell the land. When making the long-run decision whether to exit the market, the cost of land is not sunk. (We return to the issue of sunk costs shortly.)

Now let's consider what determines a firm's shutdown decision. If the firm shuts down, it loses all revenue from the sale of its product. At the same time, it saves the variable costs of making its product (but must still pay the fixed costs). Thus, *the firm shuts down if the revenue that it would get from producing is less than its variable costs of production.*

A small bit of mathematics can make this shutdown criterion more useful. If TR stands for total revenue, and VC stands for variable costs, then the firm's decision can be written as

$$\text{Shut down if } TR < VC.$$

The firm shuts down if total revenue is less than variable cost. By dividing both sides of this inequality by the quantity Q, we can write it as

$$\text{Shut down if } TR/Q < VC/Q.$$

Notice that this can be further simplified. TR/Q is total revenue divided by quantity, which is average revenue. As we discussed previously, average revenue for any firm is simply the good's price P. Similarly, VC/Q is average variable cost AVC. Therefore, the firm's shutdown criterion is

$$\text{Shut down if } P < AVC.$$

That is, a firm chooses to shut down if the price of the good is less than the average variable cost of production. This criterion is intuitive: When choosing to produce, the firm compares the price it receives for the typical unit to the average variable cost that it must incur to produce the typical unit. If the price doesn't cover the average variable cost, the firm is better off stopping production altogether. The firm can reopen in the future if conditions change so that price exceeds average variable cost.

We now have a full description of a competitive firm's profit-maximizing strategy. If the firm produces anything, it produces the quantity at which marginal cost equals the price of the good. Yet if the price is less than average variable cost at that quantity, the firm is better off shutting down and not producing anything. These results are illustrated in Figure 14-3. *The competitive firm's short-run supply curve is the portion of its marginal-cost curve that lies above average variable cost.*

SPILT MILK AND OTHER SUNK COSTS

sunk cost
a cost that has already been committed and cannot be recovered

Sometime in your life, you have probably been told, "Don't cry over spilt milk," or "Let bygones be bygones." These adages hold a deep truth about rational decision-making. Economists say that a cost is a **sunk cost** when it has already been committed and cannot be recovered. In a sense, a sunk cost is the opposite of an opportunity cost: An opportunity cost is what you have to give up if you choose to

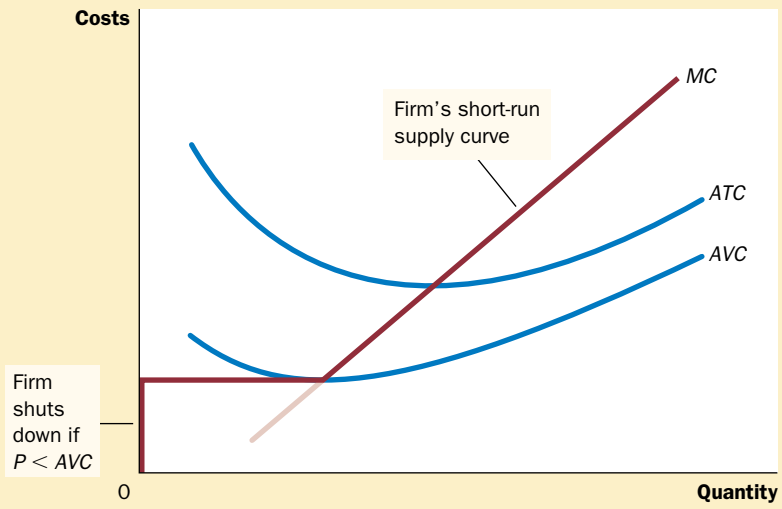

Figure 14-3

THE COMPETITIVE FIRM'S SHORT-RUN SUPPLY CURVE. In the short run, the competitive firm's supply curve is its marginal-cost curve (*MC*) above average variable cost (*AVC*). If the price falls below average variable cost, the firm is better off shutting down.

do one thing instead of another, whereas a sunk cost cannot be avoided, regardless of the choices you make. Because nothing can be done about sunk costs, you can ignore them when making decisions about various aspects of life, including business strategy.

Our analysis of the firm's shutdown decision is one example of the irrelevance of sunk costs. We assume that the firm cannot recover its fixed costs by temporarily stopping production. As a result, the firm's fixed costs are sunk in the short run, and the firm can safely ignore these costs when deciding how much to produce. The firm's short-run supply curve is the part of the marginal-cost curve that lies above average variable cost, and the size of the fixed cost does not matter for this supply decision.

The irrelevance of sunk costs explains how real businesses make decisions. In the early 1990s, for instance, most of the major airlines reported large losses. In one year, American Airlines, Delta, and USAir each lost more than $400 million. Yet despite the losses, these airlines continued to sell tickets and fly passengers. At first, this decision might seem surprising: If the airlines were losing money flying planes, why didn't the owners of the airlines just shut down their businesses?

To understand this behavior, we must acknowledge that many of the airlines' costs are sunk in the short run. If an airline has bought a plane and cannot resell it, then the cost of the plane is sunk. The opportunity cost of a flight includes only the variable costs of fuel and the wages of pilots and flight attendants. As long as the total revenue from flying exceeds these variable costs, the airlines should continue operating. And, in fact, they did.

The irrelevance of sunk costs is also important for personal decisions. Imagine, for instance, that you place a $10 value on seeing a newly released movie. You buy a ticket for $7, but before entering the theater, you lose the ticket. Should you buy another ticket? Or should you now go home and refuse to pay a total of $14 to see the movie? The answer is that you should buy another ticket. The benefit of seeing

the movie ($10) still exceeds the opportunity cost (the $7 for the second ticket). The $7 you paid for the lost ticket is a sunk cost. As with spilt milk, there is no point in crying about it.

STAYING OPEN CAN BE PROFITABLE, EVEN WITH MANY TABLES EMPTY.

CASE STUDY NEAR-EMPTY RESTAURANTS AND OFF-SEASON MINIATURE GOLF

Have you ever walked into a restaurant for lunch and found it almost empty? Why, you might have asked, does the restaurant even bother to stay open? It might seem that the revenue from the few customers could not possibly cover the cost of running the restaurant.

In making the decision whether to open for lunch, a restaurant owner must keep in mind the distinction between fixed and variable costs. Many of a restaurant's costs—the rent, kitchen equipment, tables, plates, silverware, and so on—are fixed. Shutting down during lunch would not reduce these costs. In other words, these costs are sunk in the short run. When the owner is deciding whether to serve lunch, only the variable costs—the price of the additional food and the wages of the extra staff—are relevant. The owner shuts down the restaurant at lunchtime only if the revenue from the few lunchtime customers fails to cover the restaurant's variable costs.

An operator of a miniature-golf course in a summer resort community faces a similar decision. Because revenue varies substantially from season to season, the firm must decide when to open and when to close. Once again, the fixed costs—the costs of buying the land and building the course—are irrelevant. The miniature-golf course should be open for business only during those times of year when its revenue exceeds its variable costs.

THE FIRM'S LONG-RUN DECISION TO EXIT OR ENTER A MARKET

The firm's long-run decision to exit the market is similar to its shutdown decision. If the firm exits, it again will lose all revenue from the sale of its product, but now it saves on both fixed and variable costs of production. Thus, *the firm exits the market if the revenue it would get from producing is less than its total costs.*

We can again make this criterion more useful by writing it mathematically. If TR stands for total revenue, and TC stands for total cost, then the firm's criterion can be written as

$$\text{Exit if } TR < TC.$$

The firm exits if total revenue is less than total cost. By dividing both sides of this inequality by quantity Q, we can write it as

$$\text{Exit if } TR/Q < TC/Q.$$

We can simplify this further by noting that TR/Q is average revenue, which equals the price P, and that TC/Q is average total cost ATC. Therefore, the firm's exit criterion is

$$\text{Exit if } P < ATC.$$

That is, a firm chooses to exit if the price of the good is less than the average total cost of production.

A parallel analysis applies to an entrepreneur who is considering starting a firm. The firm will enter the market if such an action would be profitable, which occurs if the price of the good exceeds the average total cost of production. The entry criterion is

$$\text{Enter if } P > ATC.$$

The criterion for entry is exactly the opposite of the criterion for exit.

We can now describe a competitive firm's long-run profit-maximizing strategy. If the firm is in the market, it produces the quantity at which marginal cost equals the price of the good. Yet if the price is less than average total cost at that quantity, the firm chooses to exit (or not enter) the market. These results are illustrated in Figure 14-4. *The competitive firm's long-run supply curve is the portion of its marginal cost curve that lies above average total cost.*

MEASURING PROFIT IN OUR GRAPH FOR THE COMPETITIVE FIRM

As we analyze exit and entry, it is useful to be able to analyze the firm's profit in more detail. Recall that profit equals total revenue (*TR*) minus total cost (*TC*):

$$\text{Profit} = TR - TC.$$

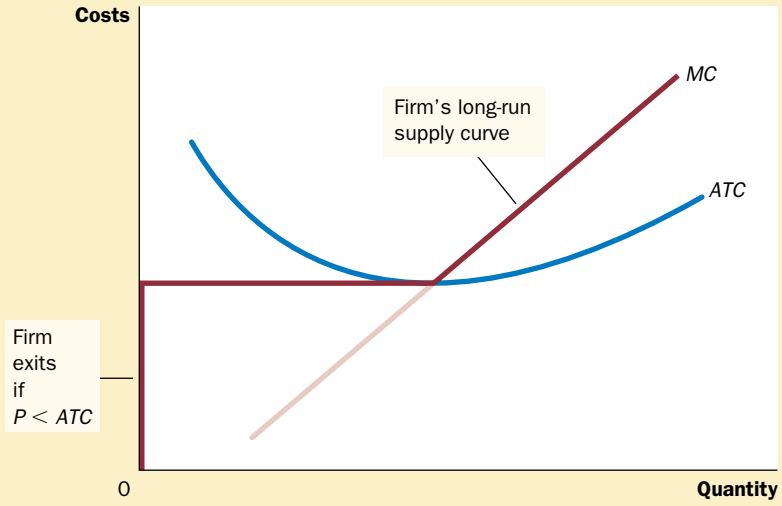

Figure 14-4

THE COMPETITIVE FIRM'S LONG-RUN SUPPLY CURVE. In the long run, the competitive firm's supply curve is its marginal-cost curve (*MC*) above average total cost (*ATC*). If the price falls below average total cost, the firm is better off exiting the market.

IN THE NEWS

Entry and Exit in Transition Economies

IN THE 1990s, MANY COUNTRIES THAT HAD previously relied on communist theories of central planning tried to make the transition to free-market capitalism. According to this article, Poland succeeded because it encouraged free entry and exit, and Russia failed because it didn't.

Russia Is Not Poland, and That's Too Bad

BY MICHAEL M. WEINSTEIN

Put aside for a moment the frightening crash of the ruble and the collapse of Russia's stock and bond markets last week. They are symptoms of something larger—a deformed economy in which the Government sets business taxes that few firms ever pay, enterprises promise wages that employees never see, loans go unpaid, people barter with pots, pans and socks, and shady dealing runs rampant.

It didn't have to be this way. The Russians need only look to Poland to behold the better road untraveled. Poland too began the decade saddled with paltry living standards bequeathed by a sclerotic, centrally controlled economy run by discredited Communists. It reached out to the West for help creating monetary, budget, trade and legal regimes, and unlike Russia it followed through with sustained political will. It now ranks among Europe's fastest-growing economies.

Key to Poland's steady success have been two policy decisions, and discussing them helps to illuminate by contrast what is going wrong with Russia.

First, Poland adopted what might be called the Balcerowicz rule, named after Leszek Balcerowicz, the Finance Minister who masterminded Poland's market reforms. Mr. Balcerowicz invited thousands of would-be entrepreneurs to sell, within loose limits, anything they wanted anywhere they wanted at whatever price they wanted. Economists called this liberalization. The Poles called it competition.

The Balcerowicz rule helped break the chokehold of Communist-dominated, state-owned enterprises and Government bureaucracies over economic activity. Also, encouraging small start-ups denies organized crime opportunities for large prey.

When Poland broke away from communism, Western economists had wrung their hands trying to figure out what to do with its sprawling state-owned factories, which operated more like social welfare agencies than production units. The solution, it turned out, was benign neglect. Rather than convert factories, the Poles allowed them to

We can rewrite this definition by multiplying and dividing the right-hand side by Q:

$$\text{Profit} = (TR/Q - TC/Q) \times Q.$$

But note that TR/Q is average revenue, which is the price P, and TC/Q is average total cost ATC. Therefore,

$$\text{Profit} = (P - ATC) \times Q.$$

This way of expressing the firm's profit allows us to measure profit in our graphs.

Panel (a) of Figure 14-5 shows a firm earning positive profit. As we have already discussed, the firm maximizes profit by producing the quantity at which price equals marginal cost. Now look at the shaded rectangle. The height of the rectangle is $P - ATC$, the difference between price and average total cost. The

shrivel. Workers peeled away to set up retail shops and other small enterprises largely free of Government interference.

The second major decision was scarier. Poland forced insolvent firms into bankruptcy, preventing them from draining resources from productive parts of the economy. That also ended a drain on the Federal budget by firms that had to be propped up by one disguised subsidy or another.

There were moments when the post-Communist Government in Russia appeared headed in the same direction. In early 1992, the Yeltsin Government embraced the Balcerowicz rule. Russians were invited to take to the streets and set up kiosks and curbside tables, selling whatever they wanted at whatever price consumers would pay. But then Communist antibodies, in the form of the oligarchs who controlled the state-owned factories and natural resources, were activated. They detected foreign tissue and attacked. Local government buried the Balcerowicz rule, imposing licensing and other requirements and eventually strangling start-ups. Professor Marshall Goldman of Harvard points to revealing comments by Viktor S. Chernomyrdin, the off-again, on-again Prime Minister whom President Boris N. Yeltsin restored to his post last week. Mr. Chernomyrdin observed that street vendors were an unattractive, chaotic blight on a proud country. The Russian authorities cracked down.

The impact was severe. Anders Aslund, a former adviser to the Russian Government now at the Carnegie Endowment for International Peace, estimates that since the middle of 1994, the number of enterprises in Russia has stagnated. In a typical Western economy, he estimates, there is 1 business for every 10 residents. In Russia, the ratio is 1 for every 55.

By snuffing out start-ups, Russia lost the remarkable device by which Poland drained workers out of worthless factories into units that could produce the goods that people wanted to buy.

Russia not only stifles start-ups; it also props up incompetents. It tolerates businesses that cannot pay taxes or wages. They survive because of systems of barter and mutual forbearance of loans and taxes. Suppliers engage in round-robin lending by which everyone owes money to someone and no one ever pays up. That too throws a lifeline to insolvent firms.

Russian factories continue to churn out steel and other products that no one needs. One measure of the deformity is that Russia is littered with factories employing 10,000 or more workers. In the United States, such factories are a rarity. The effect is to keep alive concerns that chew up $1.50 worth of resources in order to turn out a product that is worth only $1 to consumers. Economists call this "negative value added." Ordinary folk call it economic suicide.

SOURCE: *The New York Times*, August 30, 1998, Week in Review, p. 5.

width of the rectangle is Q, the quantity produced. Therefore, the area of the rectangle is $(P - ATC) \times Q$, which is the firm's profit.

Similarly, panel (b) of this figure shows a firm with losses (negative profit). In this case, maximizing profit means minimizing losses, a task accomplished once again by producing the quantity at which price equals marginal cost. Now consider the shaded rectangle. The height of the rectangle is $ATC - P$, and the width is Q. The area is $(ATC - P) \times Q$, which is the firm's loss. Because a firm in this situation is not making enough revenue to cover its average total cost, the firm would choose to exit the market.

QUICK QUIZ: How does the price faced by a profit-maximizing competitive firm compare to its marginal cost? Explain. ◆ When does a profit-maximizing competitive firm decide to shut down?

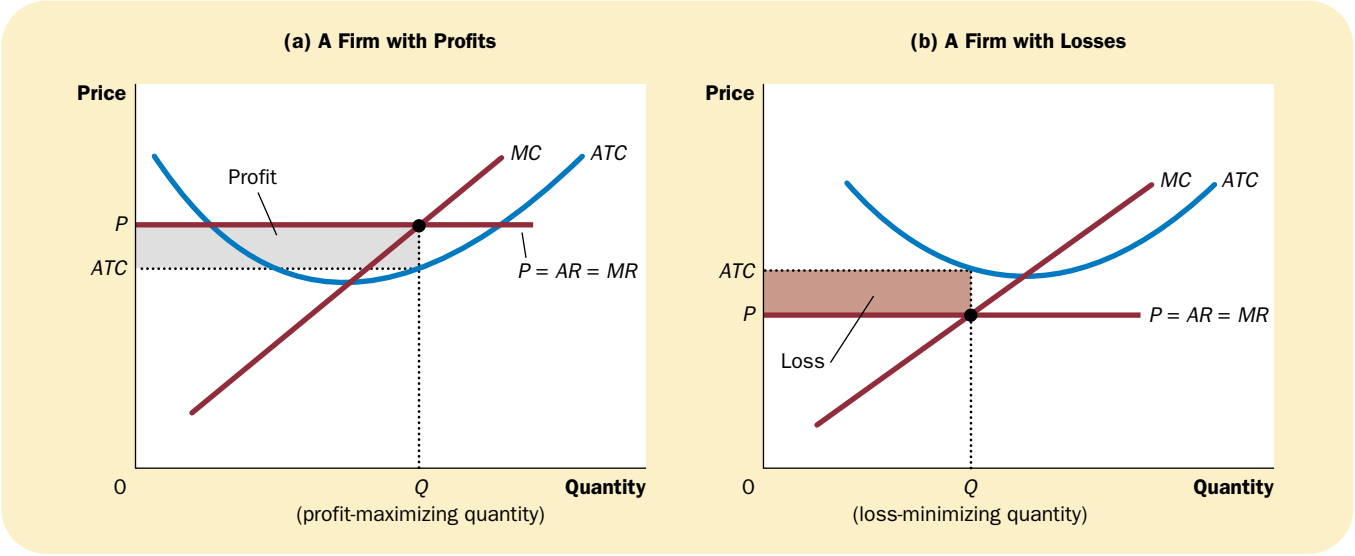

Figure 14-5

PROFIT AS THE AREA BETWEEN PRICE AND AVERAGE TOTAL COST. The area of the shaded box between price and average total cost represents the firm's profit. The height of this box is price minus average total cost ($P - ATC$), and the width of the box is the quantity of output (Q). In panel (a), price is above average total cost, so the firm has positive profit. In panel (b), price is less than average total cost, so the firm has losses.

THE SUPPLY CURVE IN A COMPETITIVE MARKET

Now that we have examined the supply decision of a single firm, we can discuss the supply curve for a market. There are two cases to consider. First, we examine a market with a fixed number of firms. Second, we examine a market in which the number of firms can change as old firms exit the market and new firms enter. Both cases are important, for each applies over a specific time horizon. Over short periods of time, it is often difficult for firms to enter and exit, so the assumption of a fixed number of firms is appropriate. But over long periods of time, the number of firms can adjust to changing market conditions.

THE SHORT RUN: MARKET SUPPLY WITH A FIXED NUMBER OF FIRMS

Consider first a market with 1,000 identical firms. For any given price, each firm supplies a quantity of output so that its marginal cost equals the price, as shown in panel (a) of Figure 14-6. That is, as long as price is above average variable cost, each firm's marginal-cost curve is its supply curve. The quantity of output supplied to the market equals the sum of the quantities supplied by each of the 1,000 individual firms. Thus, to derive the market supply curve, we add the quantity supplied by each firm in the market. As panel (b) of Figure 14-6 shows, because the

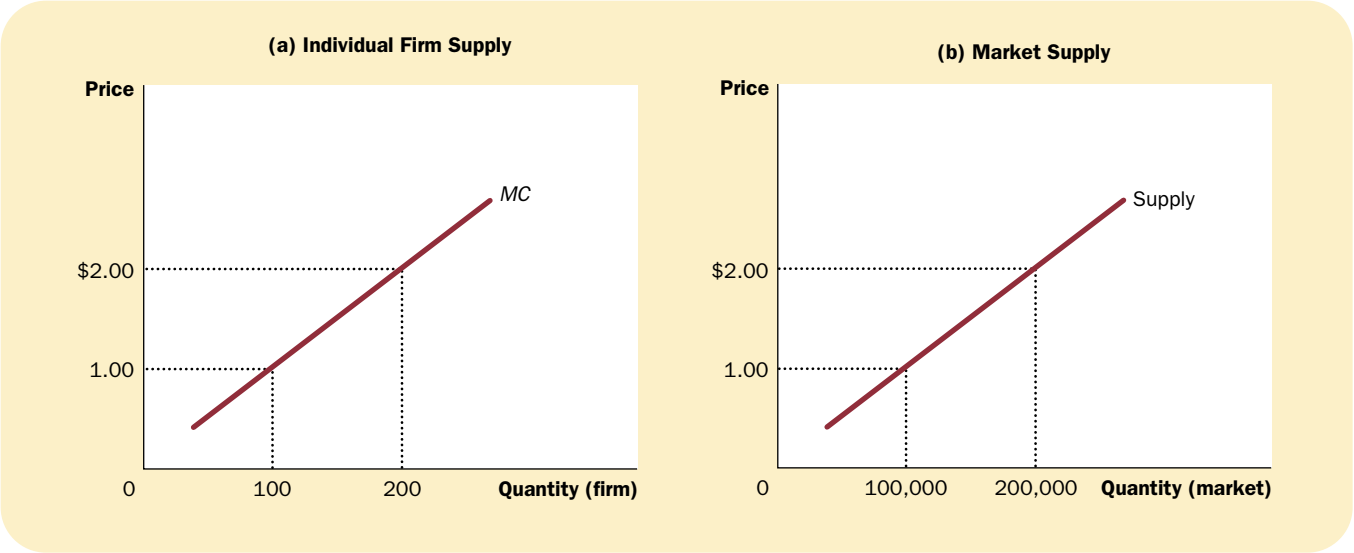

MARKET SUPPLY WITH A FIXED NUMBER OF FIRMS. When the number of firms in the market is fixed, the market supply curve, shown in panel (b), reflects the individual firms' marginal-cost curves, shown in panel (a). Here, in a market of 1,000 firms, the quantity of output supplied to the market is 1,000 times the quantity supplied by each firm.

Figure 14-6

firms are identical, the quantity supplied to the market is 1,000 times the quantity supplied by each firm.

THE LONG RUN: MARKET SUPPLY WITH ENTRY AND EXIT

Now consider what happens if firms are able to enter or exit the market. Let's suppose that everyone has access to the same technology for producing the good and access to the same markets to buy the inputs into production. Therefore, all firms and all potential firms have the same cost curves.

Decisions about entry and exit in a market of this type depend on the incentives facing the owners of existing firms and the entrepreneurs who could start new firms. If firms already in the market are profitable, then new firms will have an incentive to enter the market. This entry will expand the number of firms, increase the quantity of the good supplied, and drive down prices and profits. Conversely, if firms in the market are making losses, then some existing firms will exit the market. Their exit will reduce the number of firms, decrease the quantity of the good supplied, and drive up prices and profits. *At the end of this process of entry and exit, firms that remain in the market must be making zero economic profit*. Recall that we can write a firm's profits as

$$\text{Profit} = (P - ATC) \times Q.$$

This equation shows that an operating firm has zero profit if and only if the price of the good equals the average total cost of producing that good. If price is above average total cost, profit is positive, which encourages new firms to enter. If price

is less than average total cost, profit is negative, which encourages some firms to exit. *The process of entry and exit ends only when price and average total cost are driven to equality.*

This analysis has a surprising implication. We noted earlier in the chapter that competitive firms produce so that price equals marginal cost. We just noted that free entry and exit forces price to equal average total cost. But if price is to equal both marginal cost and average total cost, these two measures of cost must equal each other. Marginal cost and average total cost are equal, however, only when the firm is operating at the minimum of average total cost. Therefore, *the long-run equilibrium of a competitive market with free entry and exit must have firms operating at their efficient scale.*

Panel (a) of Figure 14-7 shows a firm in such a long-run equilibrium. In this figure, price P equals marginal cost MC, so the firm is profit-maximizing. Price also equals average total cost ATC, so profits are zero. New firms have no incentive to enter the market, and existing firms have no incentive to leave the market.

From this analysis of firm behavior, we can determine the long-run supply curve for the market. In a market with free entry and exit, there is only one price consistent with zero profit—the minimum of average total cost. As a result, the long-run market supply curve must be horizontal at this price, as in panel (b) of Figure 14-7. Any price above this level would generate profit, leading to entry and an increase in the total quantity supplied. Any price below this level would generate losses, leading to exit and a decrease in the total quantity supplied. Eventually, the number of firms in the market adjusts so that price equals the minimum of

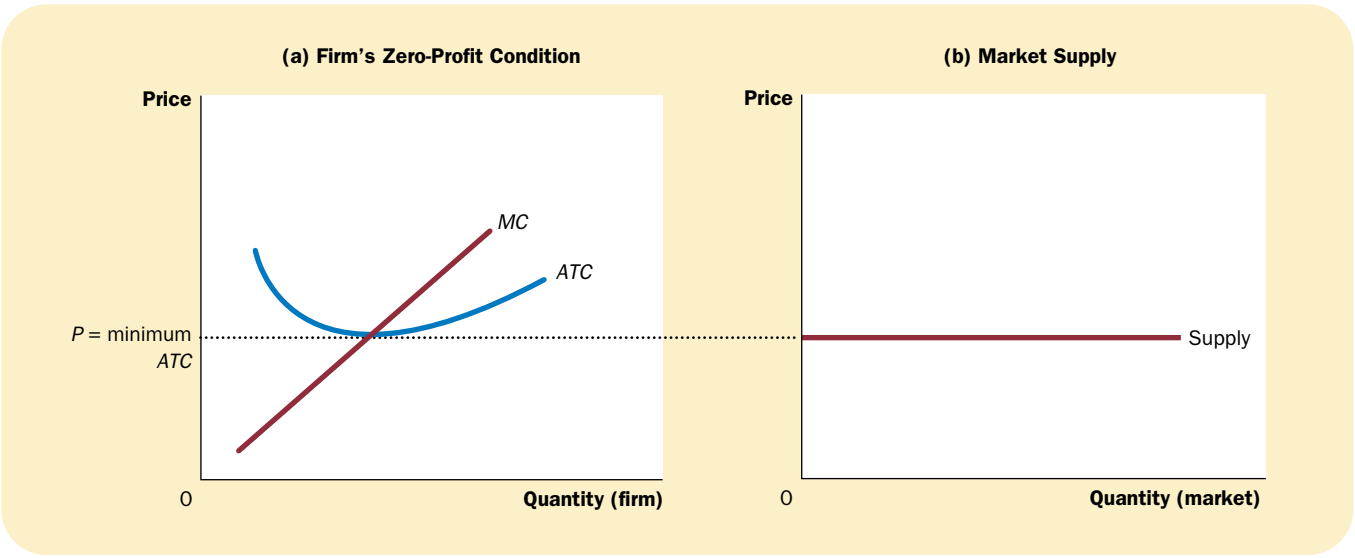

Figure 14-7

MARKET SUPPLY WITH ENTRY AND EXIT. Firms will enter or exit the market until profit is driven to zero. Thus, in the long run, price equals the minimum of average total cost, as shown in panel (a). The number of firms adjusts to ensure that all demand is satisfied at this price. The long-run market supply curve is horizontal at this price, as shown in panel (b).

average total cost, and there are enough firms to satisfy all the demand at this price.

WHY DO COMPETITIVE FIRMS STAY IN BUSINESS IF THEY MAKE ZERO PROFIT?

At first, it might seem odd that competitive firms earn zero profit in the long run. After all, people start businesses to make a profit. If entry eventually drives profit to zero, there might seem to be little reason to stay in business.

To understand the zero-profit condition more fully, recall that profit equals total revenue minus total cost, and that total cost includes all the opportunity costs of the firm. In particular, total cost includes the opportunity cost of the time and money that the firm owners devote to the business. In the zero-profit equilibrium, the firm's revenue must compensate the owners for the time and money that they expend to keep their business going.

Consider an example. Suppose that a farmer had to invest $1 million to open his farm, which otherwise he could have deposited in a bank to earn $50,000 a year in interest. In addition, he had to give up another job that would have paid him $30,000 a year. Then the farmer's opportunity cost of farming includes both the interest he could have earned and the forgone wages—a total of $80,000. Even if his profit is driven to zero, his revenue from farming compensates him for these opportunity costs.

Keep in mind that accountants and economists measure costs differently. As we discussed in Chapter 13, accountants keep track of explicit costs but usually miss implicit costs. That is, they measure costs that require an outflow of money from the firm, but they fail to include opportunity costs of production that do not involve an outflow of money. As a result, in the zero-profit equilibrium, economic profit is zero, but accounting profit is positive. Our farmer's accountant, for instance, would conclude that the farmer earned an accounting profit of $80,000, which is enough to keep the farmer in business.

"We're a nonprofit organization—we don't intend to be, but we are!"

A SHIFT IN DEMAND IN THE SHORT RUN AND LONG RUN

Because firms can enter and exit a market in the long run but not in the short run, the response of a market to a change in demand depends on the time horizon. To see this, let's trace the effects of a shift in demand. This analysis will show how a market responds over time, and it will show how entry and exit drive a market to its long-run equilibrium.

Suppose the market for milk begins in long-run equilibrium. Firms are earning zero profit, so price equals the minimum of average total cost. Panel (a) of Figure 14-8 shows the situation. The long-run equilibrium is point A, the quantity sold in the market is Q_1, and the price is P_1.

Now suppose scientists discover that milk has miraculous health benefits. As a result, the demand curve for milk shifts outward from D_1 to D_2, as in panel (b). The short-run equilibrium moves from point A to point B; as a result, the quantity rises from Q_1 to Q_2, and the price rises from P_1 to P_2. All of the existing firms respond to the higher price by raising the amount produced. Because each firm's supply curve reflects its marginal-cost curve, how much they each increase production is determined by the marginal-cost curve. In the new, short-run equilibrium, the price of milk exceeds average total cost, so the firms are making positive profit.

Over time, the profit in this market encourages new firms to enter. Some farmers may switch to milk from other farm products, for example. As the number of firms grows, the short-run supply curve shifts to the right from S_1 to S_2, as in panel (c), and this shift causes the price of milk to fall. Eventually, the price is driven back down to the minimum of average total cost, profits are zero, and firms stop entering. Thus, the market reaches a new long-run equilibrium, point C. The price of milk has returned to P_1, but the quantity produced has risen to Q_3. Each firm is again producing at its efficient scale, but because more firms are in the dairy business, the quantity of milk produced and sold is higher.

WHY THE LONG-RUN SUPPLY CURVE MIGHT SLOPE UPWARD

So far we have seen that entry and exit can cause the long-run market supply curve to be horizontal. The essence of our analysis is that there are a large number of potential entrants, each of which faces the same costs. As a result, the long-run market supply curve is horizontal at the minimum of average total cost. When the demand for the good increases, the long-run result is an increase in the number of firms and in the total quantity supplied, without any change in the price.

There are, however, two reasons that the long-run market supply curve might slope upward. The first is that some resource used in production may be available only in limited quantities. For example, consider the market for farm products. Anyone can choose to buy land and start a farm, but the quantity of land is limited. As more people become farmers, the price of farmland is bid up, which raises the costs of all farmers in the market. Thus, an increase in demand for farm products cannot induce an increase in quantity supplied without also inducing a rise in farmers' costs, which in turn means a rise in price. The result is a long-run market supply curve that is upward sloping, even with free entry into farming.

A second reason for an upward-sloping supply curve is that firms may have different costs. For example, consider the market for painters. Anyone can enter

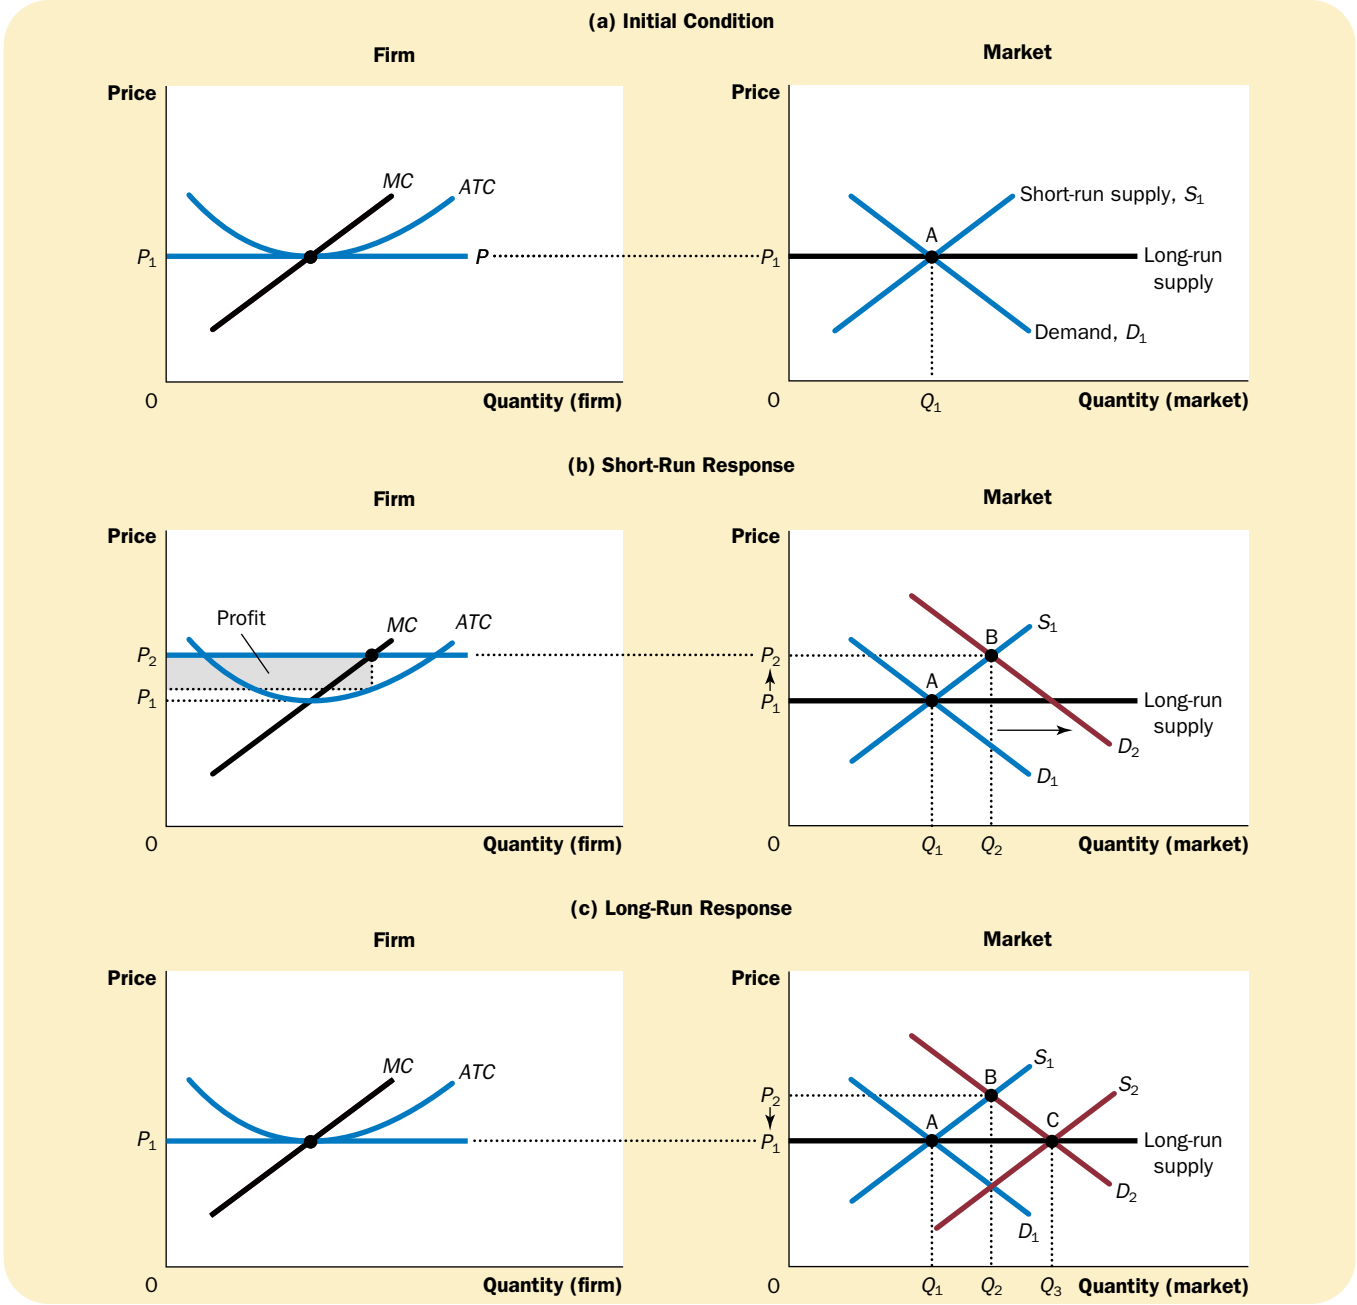

Figure 14-8

AN INCREASE IN DEMAND IN THE SHORT RUN AND LONG RUN. The market starts in a long-run equilibrium, shown as point A in panel (a). In this equilibrium, each firm makes zero profit, and the price equals the minimum average total cost. Panel (b) shows what happens in the short run when demand rises from D_1 to D_2. The equilibrium goes from point A to point B, price rises from P_1 to P_2, and the quantity sold in the market rises from Q_1 to Q_2. Because price now exceeds average total cost, firms make profits, which over time encourage new firms to enter the market. This entry shifts the short-run supply curve to the right from S_1 to S_2, as shown in panel (c). In the new long-run equilibrium, point C, price has returned to P_1 but the quantity sold has increased to Q_3. Profits are again zero, price is back to the minimum of average total cost, but the market has more firms to satisfy the greater demand.

IN THE NEWS

Entry or Overinvestment?

IN COMPETITIVE MARKETS, STRONG DEMAND leads to high prices and high profits, which then lead to increased entry, falling prices, and falling profits. To economists, these market forces are one reflection of the invisible hand at work. To the business managers, however, new entry and falling profits can seem like a "problem of overinvestment."

In Some Industries, Executives Foresee Tough Times Ahead; A Key Culprit: High Profits

BY BERNARD WYSOCKI, JR.

MONTEREY, CALIF.—About 20 executives are huddled in a conference room with a team of management consultants, and the mood is surprisingly somber.

It's a fine summer day, the stock market is booming, the U.S. economy is in great shape, and some of the companies represented here are posting stronger-than-expected profits. Best of all, perhaps, these lucky executives are just a chip shot away from the famed Pebble Beach golf course. They ought to be euphoric.

Instead, an undertone of concern is evident among these executives from Mobil Corp., Union Carbide Corp. and other capital-intensive companies. In between golf, fine meals and cigars, they hear a sobering message from their hosts.

"I feel like the prophet of doom" is the welcoming line of R. Duane Dickson, a director of Mercer Management Consulting and host of the meeting. "It's our belief that the downturn has started. I can't tell you how far it's going to go. But it could be a very ugly one."

For two days, the executives and their advisers discuss what they expect in their industries between now and 2000: growing overcapacity, world-wide product gluts, price wars, shakeouts, and consolidations. . . .

One man who attended the Pebble Beach meeting, Joseph Soviero, a Union Carbide vice president, cites an odd but basic problem in chemicals: the strong profits of the past few years. "The profitability that the industry sees during the good times has always led to overinvesting, and it has this time," Mr. Soviero says. He adds that the chemicals business cycle is alive and has peaked. At Union Carbide, he says, "we always talk about the cycle" and try to manage it.

So far, demand isn't a big problem. In many industries, it is still growing steadily, though slowly. What is developing is too much supply, stemming from the recurring problem of overinvestment. . . . The next few years will bring fierce competition and falling prices.

SOURCE: *The Wall Street Journal*, August 7, 1997, p. A1.

the market for painting services, but not everyone has the same costs. Costs vary in part because some people work faster than others and in part because some people have better alternative uses of their time than others. For any given price, those with lower costs are more likely to enter than those with higher costs. To increase the quantity of painting services supplied, additional entrants must be encouraged to enter the market. Because these new entrants have higher costs, the price must rise to make entry profitable for them. Thus, the market supply curve for painting services slopes upward even with free entry into the market.

Notice that if firms have different costs, some firms earn profit even in the long run. In this case, the price in the market reflects the average total cost of the *marginal firm*—the firm that would exit the market if the price were any lower. This firm earns zero profit, but firms with lower costs earn positive profit. Entry does not eliminate this profit because would-be entrants have higher costs than firms already in the market. Higher-cost firms will enter only if the price rises, making the market profitable for them.

Thus, for these two reasons, the long-run supply curve in a market may be upward sloping rather than horizontal, indicating that a higher price is necessary to induce a larger quantity supplied. Nonetheless, the basic lesson about entry and exit remains true. *Because firms can enter and exit more easily in the long run than in the short run, the long-run supply curve is typically more elastic than the short-run supply curve.*

QUICK QUIZ: In the long run with free entry and exit, is the price in a market equal to marginal cost, average total cost, both, or neither? Explain with a diagram.

CONCLUSION: BEHIND THE SUPPLY CURVE

We have been discussing the behavior of competitive profit-maximizing firms. You may recall from Chapter 1 that one of the *Ten Principles of Economics* is that rational people think at the margin. This chapter has applied this idea to the competitive firm. Marginal analysis has given us a theory of the supply curve in a competitive market and, as a result, a deeper understanding of market outcomes.

We have learned that when you buy a good from a firm in a competitive market, you can be assured that the price you pay is close to the cost of producing that good. In particular, if firms are competitive and profit-maximizing, the price of a good equals the marginal cost of making that good. In addition, if firms can freely enter and exit the market, the price also equals the lowest possible average total cost of production.

Although we have assumed throughout this chapter that firms are price takers, many of the tools developed here are also useful for studying firms in less competitive markets. In the next three chapters we will examine the behavior of firms with market power. Marginal analysis will again be useful in analyzing these firms, but it will have quite different implications.

Summary

- Because a competitive firm is a price taker, its revenue is proportional to the amount of output it produces. The price of the good equals both the firm's average revenue and its marginal revenue.
- To maximize profit, a firm chooses a quantity of output such that marginal revenue equals marginal cost. Because marginal revenue for a competitive firm equals the market price, the firm chooses quantity so that price equals marginal cost. Thus, the firm's marginal cost curve is its supply curve.
- In the short run when a firm cannot recover its fixed costs, the firm will choose to shut down temporarily if the price of the good is less than average variable cost. In the long run when the firm can recover both fixed and variable costs, it will choose to exit if the price is less than average total cost.
- In a market with free entry and exit, profits are driven to zero in the long run. In this long-run equilibrium, all firms produce at the efficient scale, price equals the minimum of average total cost, and the number of firms adjusts to satisfy the quantity demanded at this price.
- Changes in demand have different effects over different time horizons. In the short run, an increase in demand raises prices and leads to profits, and a decrease in

demand lowers prices and leads to losses. But if firms can freely enter and exit the market, then in the long run the number of firms adjusts to drive the market back to the zero-profit equilibrium.

Key Concepts

competitive market, p. 292
average revenue, p. 294
marginal revenue, p. 294
sunk cost, p. 298

Questions for Review

1. What is meant by a competitive firm?
2. Draw the cost curves for a typical firm. For a given price, explain how the firm chooses the level of output that maximizes profit.
3. Under what conditions will a firm shut down temporarily? Explain.
4. Under what conditions will a firm exit a market? Explain.
5. Does a firm's price equal marginal cost in the short run, in the long run, or both? Explain.
6. Does a firm's price equal the minimum of average total cost in the short run, in the long run, or both? Explain.
7. Are market supply curves typically more elastic in the short run or in the long run? Explain.

Problems and Applications

1. What are the characteristics of a competitive market? Which of the following drinks do you think is best described by these characteristics? Why aren't the others?
   a. tap water
   b. bottled water
   c. cola
   d. beer
2. Your roommate's long hours in Chem lab finally paid off—she discovered a secret formula that lets people do an hour's worth of studying in 5 minutes. So far, she's sold 200 doses, and faces the following average-total-cost schedule:

Q	Average Total Cost
199	$199
200	200
201	201

   If a new customer offers to pay your roommate $300 for one dose, should she make one more? Explain.
3. The licorice industry is competitive. Each firm produces 2 million strings of licorice per year. The strings have an average total cost of $0.20 each, and they sell for $0.30.
   a. What is the marginal cost of a string?
   b. Is this industry in long-run equilibrium? Why or why not?
4. You go out to the best restaurant in town and order a lobster dinner for $40. After eating half of the lobster, you realize that you are quite full. Your date wants you to finish your dinner, because you can't take it home and because "you've already paid for it." What should you do? Relate your answer to the material in this chapter.
5. Bob's lawn-mowing service is a profit-maximizing, competitive firm. Bob mows lawns for $27 each. His total cost each day is $280, of which $30 is a fixed cost. He mows 10 lawns a day. What can you say about Bob's short-run decision regarding shut down and his long-run decision regarding exit?
6. Consider total cost and total revenue given in the table below:

	Quantity							
	0	1	2	3	4	5	6	7
Total cost	$8	$9	$10	$11	$13	$19	$27	$37
Total revenue	0	8	16	24	32	40	48	56

   a. Calculate profit for each quantity. How much should the firm produce to maximize profit?

b. Calculate marginal revenue and marginal cost for each quantity. Graph them. (Hint: Put the points between whole numbers. For example, the marginal cost between 2 and 3 should be graphed at 2 1/2.) At what quantity do these curves cross? How does this relate to your answer to part (a)?
c. Can you tell whether this firm is in a competitive industry? If so, can you tell whether the industry is in a long-run equilibrium?

7. From *The Wall Street Journal* (July 23, 1991): "Since peaking in 1976, per capita beef consumption in the United States has fallen by 28.6 percent . . . [and] the size of the U.S. cattle herd has shrunk to a 30-year low."
   a. Using firm and industry diagrams, show the short-run effect of declining demand for beef. Label the diagram carefully and write out in words all of the changes you can identify.
   b. On a new diagram, show the long-run effect of declining demand for beef. Explain in words.
8. "High prices traditionally cause expansion in an industry, eventually bringing an end to high prices and manufacturers' prosperity." Explain, using appropriate diagrams.
9. Suppose the book-printing industry is competitive and begins in a long-run equilibrium.
   a. Draw a diagram describing the typical firm in the industry.
   b. Hi-Tech Printing Company invents a new process that sharply reduces the cost of printing books. What happens to Hi-Tech's profits and the price of books in the short run when Hi-Tech's patent prevents other firms from using the new technology?
   c. What happens in the long run when the patent expires and other firms are free to use the technology?
10. Many small boats are made of fiberglass, which is derived from crude oil. Suppose that the price of oil rises.
    a. Using diagrams, show what happens to the cost curves of an individual boat-making firm and to the market supply curve.
    b. What happens to the profits of boat makers in the short run? What happens to the number of boat makers in the long run?
11. Suppose that the U.S. textile industry is competitive, and there is no international trade in textiles. In long-run equilibrium, the price per unit of cloth is $30.
    a. Describe the equilibrium using graphs for the entire market and for an individual producer.

    Now suppose that textile producers in other countries are willing to sell large quantities of cloth in the United States for only $25 per unit.
    b. Assuming that U.S. textile producers have large fixed costs, what is the short-run effect of these imports on the quantity produced by an individual producer? What is the short-run effect on profits? Illustrate your answer with a graph.
    c. What is the long-run effect on the number of U.S. firms in the industry?
12. Suppose there are 1,000 hot pretzel stands operating in New York City. Each stand has the usual U-shaped average-total-cost curve. The market demand curve for pretzels slopes downward, and the market for pretzels is in long-run competitive equilibrium.
    a. Draw the current equilibrium, using graphs for the entire market and for an individual pretzel stand.
    b. Now the city decides to restrict the number of pretzel-stand licenses, reducing the number of stands to only 800. What effect will this action have on the market and on an individual stand that is still operating? Use graphs to illustrate your answer.
    c. Suppose that the city decides to charge a license fee for the 800 licenses. How will this affect the number of pretzels sold by an individual stand, and the stand's profit? The city wants to raise as much revenue as possible and also wants to ensure that 800 pretzel stands remain in the city. By how much should the city increase the license fee? Show the answer on your graph.
13. Assume that the gold-mining industry is competitive.
    a. Illustrate a long-run equilibrium using diagrams for the gold market and for a representative gold mine.
    b. Suppose that an increase in jewelry demand induces a surge in the demand for gold. Using your diagrams, show what happens in the short run to the gold market and to each existing gold mine.
    c. If the demand for gold remains high, what would happen to the price over time? Specifically, would the new long-run equilibrium price be above, below, or equal to the short-run equilibrium price in part (b)? Is it possible for the new long-run equilibrium price to be above the original long-run equilibrium price? Explain.
14. (This problem is challenging.) *The New York Times* (July 1, 1994) reported on a Clinton administration proposal to lift the ban on exporting oil from the North Slope of Alaska. According to the article, the administration said that "the chief effect of the ban has

been to provide California refiners with crude oil cheaper than oil on the world market. . . . The ban created a subsidy for California refiners that had not been passed on to consumers." Let's use our analysis of firm behavior to analyze these claims.

a. Draw the cost curves for a California refiner and for a refiner in another part of the world. Assume that the California refiners have access to inexpensive Alaskan crude oil and that other refiners must buy more expensive crude oil from the Middle East.

b. All of the refiners produce gasoline for the world gasoline market, which has a single price. In the long-run equilibrium, will this price depend on the costs faced by California producers or the costs faced by other producers? Explain. (Hint: California cannot itself supply the entire world market.) Draw new graphs that illustrate the profits earned by a California refiner and another refiner.

c. In this model, is there a subsidy to California refiners? Is it passed on to consumers?