

CHAPTER 4

Supply Under Perfect Competition



After reading this chapter, you will understand the following:

1. What characteristics define the structure of a market
2. What determines the profit-maximizing output level in the short run for a perfectly competitive firm
3. Under what conditions a firm will continue to operate even if it sustains a loss
4. How a firm's short-run supply curve is related to its cost curves
5. The conditions for long-run equilibrium in a perfectly competitive industry
6. What determines the shape of the long-run supply curve for a perfectly competitive industry
7. The concepts of consumer and producer surplus
8. How efficiently markets perform under perfect competition

Before reading this chapter, make sure you know the meaning of the concepts:

1. Entrepreneurship
2. Efficiency
3. Theories and models
4. Perfectly elastic demand (from the Appendix to Chapter 2)
5. Objectives, constraints, and choices
6. Market performance
7. Monopoly
8. Short- and long-run costs

IN A MARKET economy, competition is everywhere. Competition may take the form of giants such as Ford and Toyota struggling to dominate a market through advertising or introduction of new brands and styles. Sometimes it takes the form of intense rivalry between small firms for purely local markets—for example, two fast-food restaurants on opposite sides of a busy road. Sometimes competition takes the less visible form of small producers selling products like wheat or sugar that are nearly identical for

all firms in the market. In this chapter and the two that follow, we will examine all of these varieties of competition.

Economists refer to the conditions under which competition occurs in a market as **market structure**. Market structure is defined in terms of the number and size of firms, the nature of the product, ease of entry and exit, and availability of information.

In this book we will look at four market structures that have traditionally been emphasized in neoclassical economics. The first, to which this chapter is devoted, is **perfect competition**. The defining characteristics of perfect competition are the presence of many firms, none with a significant share of the market; a product that is homogeneous; easy entry into the industry and exit from it; and equal access to information by buyers and sellers. In saying that no firm has a “significant” share of the market, we mean that each firm is so small that its actions, alone, have no noticeable effect on the market price. By a “homogeneous product,” we mean that the various firms’ products are so nearly alike that they are perfect substitutes in the eyes of buyers. By “ease of entry,” we mean that firms that are just starting to produce the product can do so on an equal footing with existing firms in terms of the prices paid for inputs, availability of technology, access to government permits or licenses, and so on. By “ease of exit,” we mean that firms face no legal barriers to leaving the market and are able to find buyers or other uses for their fixed inputs. Finally, by “equal access to information,” we mean that all buyers and sellers have complete information about the price of the product and of the inputs used to produce it, that buyers know all they need to know about product characteristics, and that all producers have equal knowledge of production techniques.

A second market structure, *monopoly*, is at the opposite extreme from perfect competition. A monopoly is a market in which a single firm accounts for 100 percent of sales of a product that has no close substitutes. Monopoly will be examined in detail in the next chapter.

Perfect competition and monopoly are “ideal type” market structures. Few if any markets exactly fit the definitions, although many approximate them. The next two market structures are more descriptive. **Oligopoly** means a market with only a few firms, at least some of which have a significant share of the market. The product may be either homogeneous or differentiated; there may or may not be significant barriers to entry; and buyers and sellers need not have equal access to all kinds of information. Most familiar markets for branded products, from automobiles to toothpaste, fit in this category. **Monopolistic competition** resembles perfect competition in that there are many small firms and easy entry and exit, but under monopolistic competition the various firms’ products are differentiated from one another. Many sectors of retail trade and small service firms fit this category.

The characteristics of the four market structures are summarized in Table 4.1.

Perfect Competition and Supply in the Short Run

In building a model to fit the market structure of perfect competition, our objectives are first, to show how the profit-maximizing decisions of individual firms determine the quantity they will supply at various prices and, second, to show how individual

Market structure

The key traits of a market, including the number and size of firms, the extent to which the products of various firms are different or similar, ease of entry and exit, and availability of information

Perfect competition

A market structure that is characterized by a large number of small firms, a homogeneous product, freedom of entry and exit, and equal access to information

Oligopoly

A market structure in which there are only a few firms, at least some of which are large in relation to the size of the market

Monopolistic competition

A market structure in which there are many small firms, a differentiated product, and easy entry and exit

TABLE 4.1 MARKET STRUCTURES

	Number and Size of Firms	Nature of Product	Entry and Exit Conditions	Information Availability
Perfect Competition	Many firms, all small	Homogenous	Easy	Equal access to all information
Monopolistic Competition	Many firms, all small	Differentiated	Easy	Some restrictions
Oligopoly	Few firms, at least some of them large	Differentiated or homogeneous	May be some barriers to entry	Some restrictions
Monopoly	One firm	Unique product	Barriers to entry are common	Some restrictions

The structure of a market refers to the conditions under which firms compete in it—the number and size of firms, the nature of the product, the ease of entry and exit, and the availability of information. Perfect competition and monopoly are “ideal” types of structures. Few—if any—markets fit their definitions perfectly. Monopolistic competition and oligopoly are descriptive of most markets in the real world.

firms’ decisions generate market supply curves. We will look first at the short run and then at the long run.

The Constraints

The models used in neoclassical economics assume that, regardless of market structure, the firm’s objective is to maximize profit. That means that differences in the choices that firms make under various market structures must be traced to differences in the constraints they face rather than to differences in objectives.

COST CONSTRAINTS One set of constraints, those imposed by costs, was discussed in the preceding chapter. In the case of perfect competition, we make three special assumptions regarding costs:

1. All firms in the market have access to the same technology and know where to buy inputs at the same prices. These conditions are implied by the assumptions of a homogeneous product and equal access to information by all firms. As a result, all firms have identical long- and short-run cost curves.
2. Economies of scale are exhausted at a small level of output relative to the quantity demanded at the prevailing price. As a result, there is room in the market for many firms producing at the minimum long-run average cost. Without this assumption, it would not be possible to maintain many firms, each small relative to the market.

- There are no sunk costs. Firms that leave the market are able to recover implicit fixed costs by selling their plant and equipment to other firms. This is part of the requirement of free entry and exit.


DEMAND CONSTRAINTS: THE FIRM AS PRICE TAKER Demand is the other principal constraint on the choices made by a profit-maximizing firm. Under perfect competition, the demand constraint has a special form. Because all firms in such an industry are small and have homogeneous products, each firm is a **price taker**. This means that the price at which each firm sells its output is determined by forces beyond the firm's control, namely, supply and demand conditions in the market as a whole. If an individual firm were to raise its price even a fraction above the prevailing market price, it would lose all of its customers. Equally, there would be no point in lowering its price even a fraction below the prevailing market price. If it did so, it would be overwhelmed by more orders than it could possibly fill. In graphical terms, we can say that each firm in a perfectly competitive market faces a horizontal demand curve, as shown in Figure 4.1.

Part (a) of Figure 4.1 shows the supply and demand curves for the market for chicken as given in earlier chapters. The equilibrium price is \$2.00 per pound, and the

Price taker

A firm that sells its output at prices that are determined by forces beyond its control

FIGURE 4.1 MARKET DEMAND AND DEMAND FOR THE PERFECTLY COMPETITIVE FIRM



The perfectly competitive firm is a price taker. It is so small relative to the market as a whole that its decisions do not significantly affect the market price. In this example, the market equilibrium price is \$2.00 per pound. The price will not be much affected if the individual firm shown in part (b) produces 20,000 rather than 40,000 pounds out of the billions of pounds produced in the market as a whole. Because the individual competitive firm is a price taker, the demand curve it faces is perfectly elastic. As a result, marginal revenue equals price for a perfectly competitive firm.

equilibrium quantity is 2 billion pounds per year. Part (b) shows how the market looks from the viewpoint of an individual producer. The range of possible outputs is measured in terms of thousands rather than billions of pounds. The range of choice over which any one firm can vary its output is so small relative to the total quantity demanded that the market price will not be perceptibly affected whether the firm produces 10,000, 20,000, or 40,000 pounds of chicken a year. A 10,000-pound movement is too small even to see on the scale of the market supply and demand curves. As far as the individual firm is concerned, then, the demand curve it faces appears to be horizontal (perfectly elastic) at the market price, even though, when viewed from the perspective of the market as a whole, the demand curve has the usual negative slope.

Previously we introduced the term *marginal cost* to refer to the amount by which total cost changes when output changes by one unit. Now we can introduce a similar term, **marginal revenue**, to refer to the amount by which total revenue changes as a result of a one-unit change in output. Recall that *revenue* means price multiplied by quantity sold. For a firm with a perfectly elastic demand curve, marginal revenue simply equals price. For example, if the price of chicken is \$2 per pound, the firm will receive revenue of \$200 from the sale of 100 pounds of chicken and revenue of \$202 from the sale of 101 pounds. A 1-pound increase in output yields a \$2 increase in revenue, that is, an increase in revenue equal to the product's price. Although marginal revenue and price are equal for a perfectly competitive firm, we will see in the next chapter that this is not the case for market structures in which the firm's demand curve is not perfectly elastic.

Marginal revenue

The amount by which total revenue changes as a result of a one-unit increase in quantity sold

Short-Run Profit Maximization for the Firm

How does an individual firm in a perfectly competitive market maximize profits, given the constraints imposed by its cost and demand curves? A simple numerical example will help us answer this question. Our example will be based on the imaginary firm Fieldcom, which was introduced in the last chapter.¹

Part (a) of Figure 4.2 shows short-run cost data for Fieldcom as given in the last chapter. It also shows the revenue Fieldcom earns from the sale of each quantity of output, assuming a market price of \$500 per unit.

Subtracting total cost in column 3 from total revenue in column 2 yields the total profit the firm earns at each output level. The maximum is reached at 19 units per day, where a profit of \$2,000 per day is earned. The profit-maximizing output level is shown graphically in part (b) of Figure 4.2. There the firm's total profit is indicated by the distance between the total revenue and total cost curves. That distance is greatest at 19 units of output.

Instead of comparing total cost and total revenue, we can find the profit-maximizing output level by comparing marginal cost and marginal revenue. Look first at columns 5 and 6 of part (a) of Figure 4.2. Column 5 gives data on marginal cost. (Marginal cost data are printed on lines between the entries in the first four columns to show that marginal cost is the change in cost as output moves from one level to another.) Column 6 shows marginal revenue, which, as explained, is equal to the product's price. Each PDA that Fieldcom sells adds \$500 to its total revenue.

As the table shows, both total cost and total revenue rise as output increases. If the increase in revenue exceeds the increase in cost (that is, if marginal revenue is greater

than marginal cost), boosting output by one unit increases total profit. If the increase in cost exceeds the increase in revenue (that is, if marginal cost is greater than marginal revenue), raising output by one unit reduces total profit. Therefore, to maximize profit a firm should expand its output as long as marginal revenue exceeds marginal cost and should stop as soon as rising marginal cost begins to exceed marginal revenue. A comparison of columns 5 and 6 of Figure 4.2 shows that for Fieldcom this means producing 19 units of output per day—the same number we arrive at when we compare total cost and total revenue.

FIGURE 4.2 SHORT-RUN PROFIT MAXIMIZATION UNDER PERFECT COMPETITION

(a)

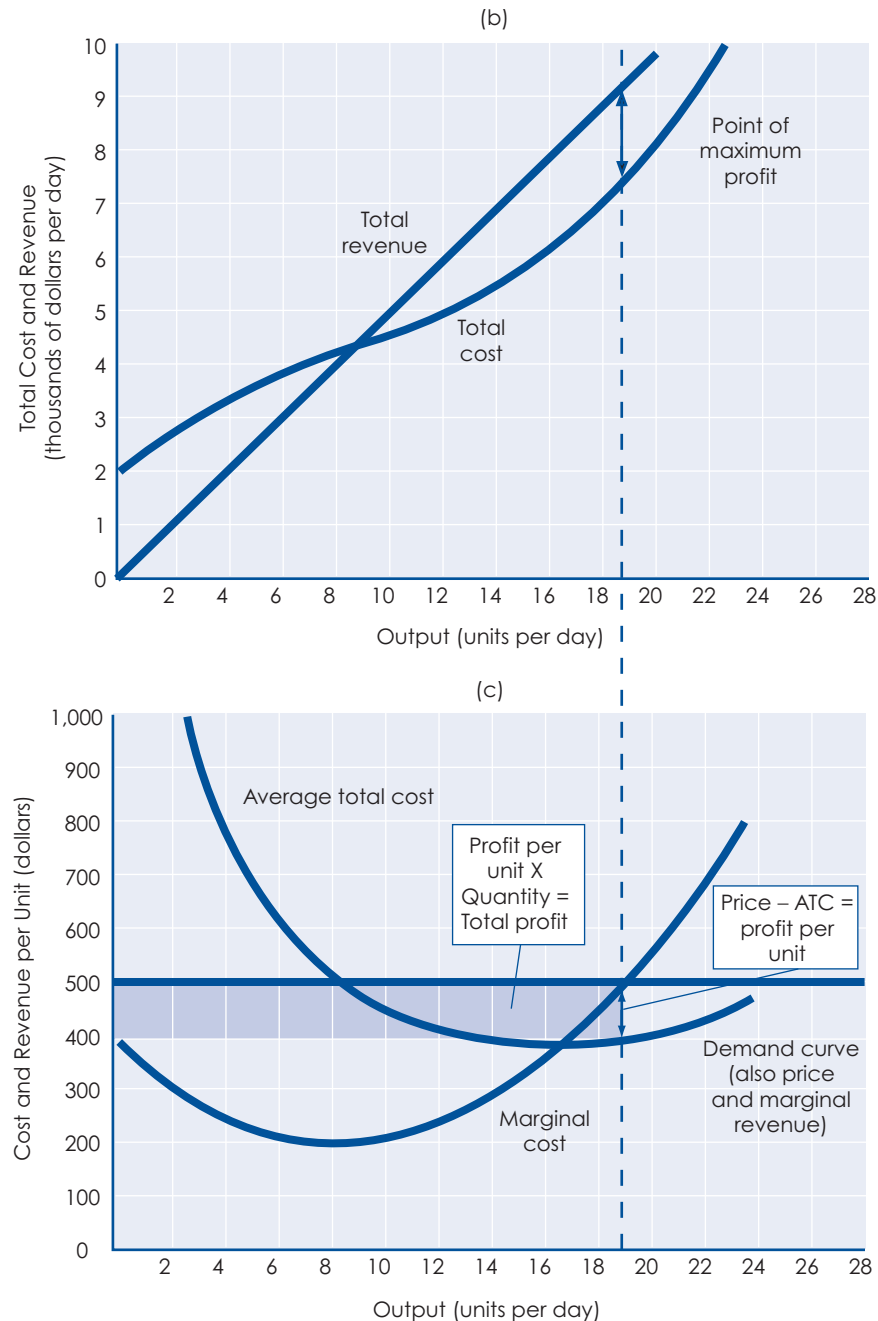
Quantity of Output (1)	Total Revenue (2)	Total Cost (3)	Total Profit (2) – (3) (4)	Marginal Cost (5)	Marginal Revenue (6)
0	\$ 0	\$2,000	-\$2,000		
1	500	2,380	-1,880	\$380	\$500
2	1,000	2,720	-1,720	340	500
3	1,500	3,025	-1,525	305	500
4	2,000	3,300	-1,300	275	500
5	2,500	3,550	-1,000	250	500
6	3,000	3,780	-780	230	500
7	3,500	3,955	-495	215	500
8	4,000	4,200	-200	205	500
9	4,500	4,400	100	200	500
10	5,000	4,605	395	205	500
11	5,500	4,820	680	215	500
12	6,000	5,050	950	230	500
13	6,500	5,300	1,200	250	500
14	7,000	5,575	1,425	275	500
15	7,500	5,880	1,620	305	500
16	8,000	6,220	1,780	340	500
17	8,500	6,600	1,900	380	500
18	9,000	7,025	1,975	425	500
19	9,500	7,500	2,000	475	500
20	10,000	8,030	1,970	530	500
21	10,500	8,620	1,880	590	500
22	11,000	9,275	1,725	655	500
23	11,500	10,000	1,500	725	500
24	12,000	10,800	1,200	800	500

(continues)

This figure shows the profit-maximizing level of output chosen by a perfectly competitive firm, Fieldcom, Inc., given a market price of \$500 per unit. That level of output can be found by comparing total cost and total revenue, as shown in parts (a) and (b). It can also be found by comparing marginal cost and marginal revenue. (Because the firm is a price taker, marginal revenue is equal to price.) Profit increases up to the point at which rising marginal cost begins to exceed marginal revenue; after that point, it declines. Regardless of the approach used, the profit-maximizing output is 19 units per day and the maximum profit per day is \$2,000.

The marginal approach to short-run profit maximization is shown graphically in part (c) of Figure 4.2. Up to about 19 units of output, the marginal cost curve lies below the marginal revenue curve, so each added unit of output increases profit. (The graph, unlike the table, pictures output as a continuous quantity so that profit maximization need not occur exactly at an even number of units.) Beyond that point, the marginal cost

FIGURE 4.2 SHORT-RUN PROFIT MAXIMIZATION UNDER PERFECT COMPETITION, CONTINUED



curve rises above the marginal revenue curve so each added unit of output reduces profit. The point of profit maximization—the point at which the rising section of the marginal cost curve intersects the marginal revenue curve—matches the point in part (b) at which the spread between total revenue and total cost is greatest.

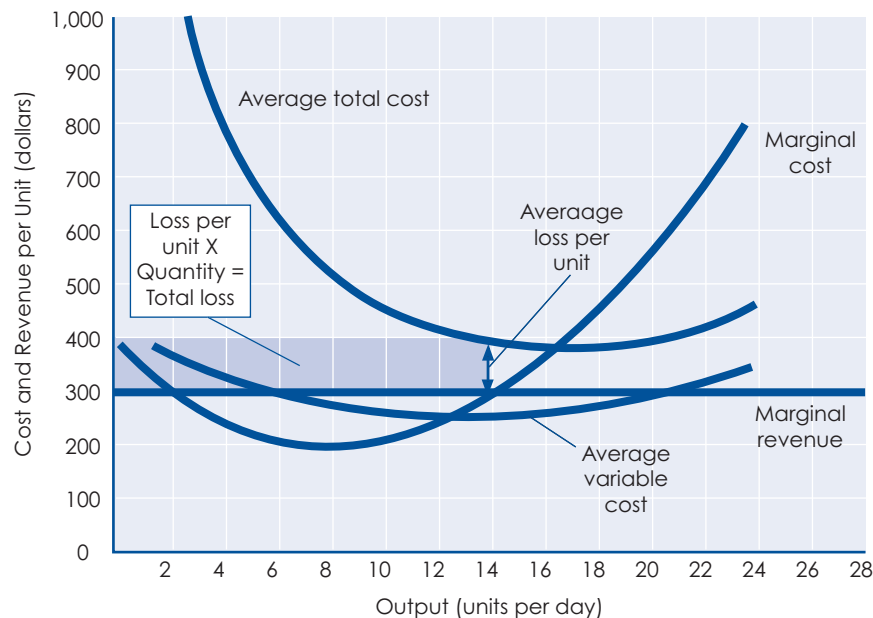
In part (c), the vertical distance between the demand curve, which shows price, and the average total cost curve represents the profit per unit. Profit per unit multiplied by the number of units gives total profit. Thus, from the standpoint of part (c), total profit equals the area of the shaded rectangle.

Minimizing Short-Run Losses

In the example just given, Fieldcom was able to make a profit at a price of \$500. However, market conditions might not always be so favorable. Suppose, for example, that the market price drops to \$300. A lower market price means a downward shift in the firm's perfectly elastic demand curve. Being a price taker, the firm can do nothing about the price and will have to adjust its output as best it can to meet the new situation. The required adjustments are shown in Figure 4.3.

There is no output level at which the firm can earn a profit given a price of \$300. Unable to earn a profit, the firm must focus on keeping its losses to a minimum. With

FIGURE 4.3 MINIMIZING SHORT-RUN LOSSES UNDER PERFECT COMPETITION



If the product's market price is too low to permit earning a profit, the firm must try to keep its losses to a minimum. For Fieldcom, Inc., given a price of \$300 per unit, the point of minimum loss is 14 units of output per day. The marginal cost curve intersects the marginal revenue curve at a point higher than average variable cost but lower than average total cost. Each unit of output sold earns more than its share of variable cost but not enough to pay for its share of total cost when its share of fixed cost is included.

a price of \$300 per unit, the minimum loss occurs at 14 units of output. As in the previous case, that is the output level beyond which marginal cost begins to exceed the product's price.

In graphical terms, we note that the point at which the rising section of the marginal cost curve intersects the marginal revenue curve lies between the average total cost and average variable cost curves.² Because the demand curve is below the average total cost curve, there is a loss on each unit sold. The total loss is equal to the shaded rectangle (loss per unit times quantity of output). However, the demand curve lies above the average variable cost curve. This means that revenue per unit is more than enough to cover variable cost and, hence, that each unit sold makes at least some contribution to covering fixed cost. Thus, losses are smaller than they would be if no output were produced, assuming that fixed costs must be paid even when output drops to zero.


As an aid to understanding the logic of the loss-minimizing decision, suppose for a moment that wages are the firm's only variable cost and that rent on its building is its only fixed cost. At the point shown, the firm is bringing in more than enough revenue to pay its wage bill (variable costs); the remainder will help pay the rent. If the firm shuts down temporarily, it will have to pay the rent with no help at all from current revenue. That would mean a loss equal to fixed cost—even more of a loss than at 14 units of output per day.

The logic of continuing operations in order to minimize losses applies only in the short run, when the costs of fixed inputs must be borne regardless of how much output is produced. A firm would not continue to operate indefinitely with the price below average total cost as shown in Figure 4.3. In the long run, a firm can free itself of fixed costs by selling its equipment, allowing long-term leases to expire, and so on. We will return to the conditions under which firms will leave the industry later in the chapter.

Shutting Down to Cut Short-Run Losses

What would happen if the price of PDAs dropped even lower than \$300? Would it then still be worthwhile for the firm to keep making them even though it was losing money? The answer, as shown in Figure 4.4, is no.

The figure assumes a price of \$225 per unit. With such a low price the firm cannot make a profit at any output level. But this time the best thing for the firm to do in the short run is to shut down. As illustrated by *Economics in the News 4.1*, temporary shutdowns are a normal way of adapting to changing supply and demand conditions. The example illustrates the point that shutting down is not at all the same as going out of business. Provided that the outlook for the future is good, it makes sense for a firm to keep its plant intact, pay its rent, and even continue some benefits for employees to ensure that they will be ready to come back when called. The firm therefore does not escape its fixed costs. When market conditions improve, inventories are brought into line with demand, and as the market price rises again, the firm can resume operations. Only if market conditions are never expected to improve will the firm consider winding up its affairs and going out of business.

FIGURE 4.4 SHUTTING DOWN TO MINIMIZE SHORT-RUN LOSS

The price of a firm's output may drop so low that the firm must shut down in order to keep short-run losses to a minimum. As illustrated here, such a situation occurs for Fieldcom at a price of \$225 per unit. Marginal cost rises above marginal revenue at about 11 units of output. That output yields a smaller loss (\$2,345) than those slightly greater or lower. However, the loss can be reduced to just \$2,000 a day if the firm shuts down. The marginal cost curve in this case intersects the marginal revenue curve at a point below average variable cost. That is the signal to shut down.

In the case of a temporary shutdown, it can be misleading to follow the rule of expanding output until marginal cost begins to exceed marginal revenue. With the price at \$225, such a point is reached at about 11 units of output per day. That output level does give the firm a lower loss than a level slightly higher or slightly lower; but in this case the firm incurs an even smaller loss by not producing at all.

The reason 11 units of output does not minimize loss is that the demand curve lies below the average variable cost curve at that point. Suppose again that wages are the firm's only variable cost and rent is its only fixed cost. At 11 units of output, revenue is not enough even to meet the firm's payroll. The firm will do better to send its workers home and save the cost of wages, even though when it does this the owners will have to pay the entire rent from reserves, without any help from current sales revenue.

The Firm's Short-Run Supply Curve

The examples just given provide the information needed to draw a short-run supply curve for a perfectly competitive firm. Let's work through an example like the one



Economics in the News 4.1 CHANGING WITH THE SEASONS

Croatia is a small country that must make the most of its resources as it hurries to catch up to the living standards of the rest of Europe. One asset that will help it do so is a stunning Adriatic coastline with brilliant sun, sparkling clear water, and hundreds of islands, perfect for get-away weekends by Parisians or Berliners.

In the opinion of many, the crown jewel among the Croatian islands is Hvar. The tiny port of Hvar Town is a tourist paradise of ancient red-roofed houses cascading down a steep hill from an old fort at the top to a quaint fishing harbor below. Taking advantage of some fine, clear weather in mid-March, a recent visitor found the streets of Hvar Town lined with hotels, restaurants, and souvenir shops—all of them closed! Although it has facilities to serve thousands of tourists in the summer, all but one hotel and all but three or four restaurants close for the season each winter, to open again in the late spring.

At first it seems like such a waste. All that natural beauty is still there, and the water is just as clear even if the air is a few degrees cooler. Why don't the hotels and



Ancient street on the island of Hvar, Croatia

restaurants just offer low, off-season rates in order to stay busy all winter? With a little thought, though, the economic logic of the decision to shut down becomes clear. Hotel and restaurant owners have to take a close look at their costs when deciding whether to operate on a year-round or seasonal basis. Some costs, like property taxes and interest on bank loans can't be avoided by shutting down, but others, especially labor costs for clerks and kitchen staff, can be eliminated when the establishment is closed. To a certain point, it pays to stay open by offering lower, off-season rates, but when those rates fall so low that they don't even cover the wages of the cooks and cleaners, it is time to put up the shutters and wait for spring.

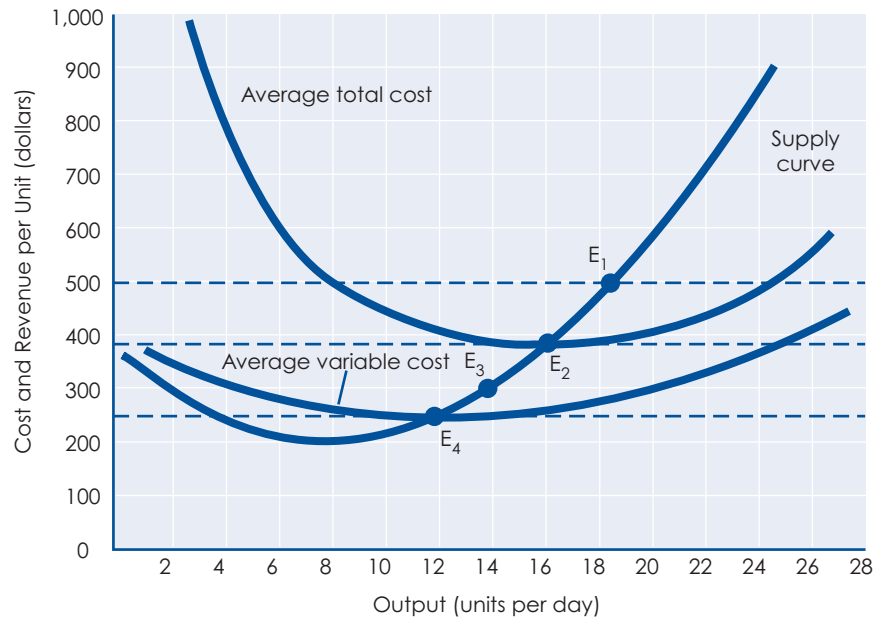
And what about the one hotel that stays open in the winter? Is the owner just ignorant of economics? No, there's an economic logic behind this hold-out

strategy, too. If just one hotel stays open while all others shut, the few visitors to the town are enough to push its revenues above the break-even point. It's all an example of competition in action.

shown in Figure 4.5 starting with a price of \$500. As we saw earlier, Fieldcom will turn out 19 devices a day at that price. Point E_1 of the firm's short-run marginal cost curve thus is a point on its supply curve.

Now suppose that the demand for PDAs slackens and the market price begins to fall. As it does so, the point at which marginal revenue equals marginal cost moves downward along the firm's marginal cost curve. Soon point E_2 is reached—the point at which marginal cost and average total cost are equal. This occurs at an output of about 17 units and a price of about \$385. At that price, the best the firm can do is break even; either a greater or a smaller output will result in a loss.

If the price falls still lower, the firm's objective becomes one of keeping its loss to a minimum. At a price of \$300, for example, the firm minimizes its loss by making 14 units (point E_3). In the range of prices between minimum average total cost and minimum average variable cost, the supply curve continues to follow the marginal cost curve.

FIGURE 4.5 DERIVATION OF THE SHORT-RUN SUPPLY CURVE

This graph shows how a short-run supply curve for Fieldcom, Inc. can be derived from its cost curves. When the price and marginal revenue is \$500, the firm will produce at point E_1 . As the price falls the firm will move downward along its short run marginal cost curve as shown by points E_2 and E_3 . The firm will continue to produce at the point at which price equals marginal cost until marginal cost falls below average variable cost. E_4 thus is the lowest point on the firm's supply curve. Below that price the firm will shut down.


At about \$250 the price reaches the lowest point on the average variable cost curve. There the firm is just on the edge of shutting down—it is covering its variable costs with nothing to spare. Its loss is equal to its fixed costs. At any lower price the firm will minimize its losses by shutting down. Thus, point E_4 is the lowest point on the marginal cost curve that can be considered part of the firm's supply curve.

The preceding discussion of the firm's short-run supply decision can be summed up as follows: *The short-run supply curve of a profit-maximizing firm operating in a perfectly competitive market coincides with the upward sloping part of the marginal cost curve lying above its intersection with the average variable cost curve.*

The Industry's Short-Run Supply Curve

Once we have a supply curve for each firm in an industry, we can add them together to construct a supply curve for the industry as a whole. Figure 4.6 shows how this can be done, beginning with the supply curves for three firms. To get the total supply of the three firms at each price, the quantities supplied by each firm are added together. In graphical terms this means adding the supply curves horizontally. To generalize the process to an industry with many firms, the individual supply curves of the remaining firms in the industry would be added to the three shown.

FIGURE 4.6 DERIVATION OF A SHORT-RUN INDUSTRY SUPPLY CURVE



A short-run industry supply curve can be obtained by summing the supply curves of individual firms. Here this method is shown for the first three firms in an industry. The supply curves of additional firms would be added in the same way. If the prices of inputs change as industry output varies, the industry supply curve will need to be adjusted.

In adding the firms' supply curves together, we assumed that input prices did not change as output expanded. For a small firm in a perfectly competitive industry, this is a realistic assumption. However, if all firms in an industry try to grow at the same time, the assumption may not hold. In fact, input prices will rise unless a greater quantity of inputs can be purchased without paying higher prices, that is, unless the short-run supply curves for inputs to the industry are perfectly elastic. If input prices rise as the industry's total output grows, each firm's cost curves will shift upward as the output of all firms increases. That will make the short-run industry supply curve somewhat steeper than the sum of the individual supply curves.

Long-Run Equilibrium Under Perfect Competition

Up to this point we have considered changes in industry output that result from firms' decisions to produce more or less as the market price changes. In doing so, however, we have neglected an important part of a competitive industry's response to changes in demand: the processes of entry and exit.

Consideration of entry and exit moves us from the short run to the long run. In the last chapter we distinguished between the long run, when all inputs can be varied, and the short run, when some inputs are fixed. The ability to vary all inputs in the long run—even durable ones such as land, structures, and major pieces of equipment—allows firms to enter a market for the first time, starting with a new plant and work force. It also means that they can leave a market for good, releasing all their employees and selling their plant and equipment. (Sometimes firms leave voluntarily, with the owners selling the firm's assets and dividing up the proceeds. Other times they leave the market only when forced to do so, such as when creditors resort to a bankruptcy court to force a sale of the firm's assets in order to pay its debts.) Typically, as an industry expands and contracts, many firms enter and leave it.

Free entry and exit of firms is one of the basic traits of a perfectly competitive market. Free entry does not mean that firms can enter at no cost. They may have to pay a great deal to purchase equipment, hire key employees, and so on. Free entry simply means that if they are willing to make the necessary investment, new firms are free to compete with existing ones on a level playing field. They are not kept out by patents or licensing requirements, trade secrets, collusion by firms already in the industry, or lack of access to raw materials. Likewise, free exit means that firms face no legal barriers to shutting down or moving if they find that they cannot make a profit. Strictly interpreted, free exit also means that firms have no sunk costs. When they leave the industry, they can put fixed assets to other uses or find buyers for them.

Free entry and exit did not play a direct role in our discussion of a firm's short-run supply decision. However, as we will now see, it is crucial to understanding how a competitive market works in the long run.

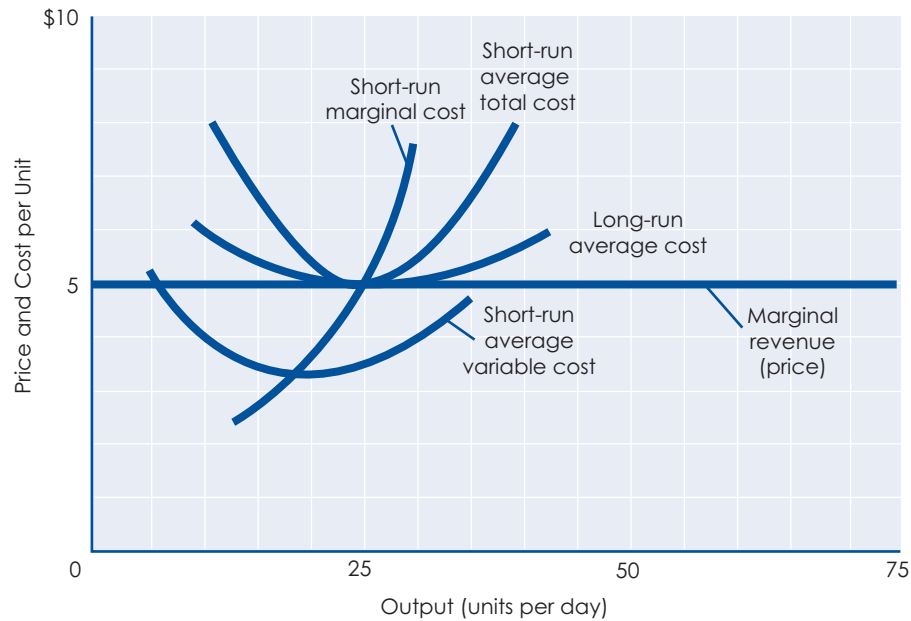
Long-Run Equilibrium for a Competitive Firm

At numerous points we have used the term *equilibrium* to refer to a state of affairs in which economic decision makers have no incentive to change their plans. Three conditions are required for a perfectly competitive firm to be in equilibrium in the long run:

1. The firm must have no incentive to produce a larger or smaller output given the size of its plant (that is, the amount of fixed inputs it uses). That requires that short-run marginal cost be equal to short-run marginal revenue, which in turn means that the short-run equilibrium condition is also a condition for long-run equilibrium.
2. Each firm must have no incentive to change the size of its current plant (that is, the amount of fixed inputs it uses).
3. There must be no incentive for new firms to enter the industry or for existing firms to leave it.

Figure 4.7 shows a perfectly competitive firm for which these three requirements are met. First, short-run marginal cost equals price at 25 units of output per day, which is the level of output the firm will choose in order to make the maximum profit. Second, the firm has a plant that is just the right size to make short-run average total cost equal to the lowest possible long-run average cost at the chosen output level. The short-run average total cost curve for a plant of any other size would give a higher average total cost for the chosen output. Third, both long-run average cost and short-run average total cost are equal to price at the equilibrium level of output. This guarantees that there is no incentive for entry or exit. As always, average total cost comprises both explicit and implicit costs, including the opportunity cost of capital, or “normal profit.” When price equals average total cost, then, firms are earning zero economic profit. Any positive economic profit would attract new firms into the industry, whereas negative economic profits (economic losses) would cause firms to leave the industry.

FIGURE 4.7 A PERFECTLY COMPETITIVE FIRM IN LONG-RUN EQUILIBRIUM



Long-run equilibrium in a perfectly competitive industry requires that the typical firm (1) have no short-run incentive to change the level of its output; (2) have no long-run incentive to change the size of the plant used to produce its output; and (3) have no long-run incentive to enter or leave the industry. This requires that price, short-run marginal cost, short-run average total cost, and long-run average cost all have the same value in equilibrium as shown here.

The three conditions for long-run equilibrium are summarized in the following equation.


$$\text{Price} = \text{Marginal cost} = \text{Short-run average total cost} = \text{Long-run average cost}$$

If any part of this equation does not hold, firms will have a reason to change their plans. If price does not equal short-run marginal cost, they will have an incentive to change their output levels by changing the quantity of variable inputs used, even if they cannot, in the short run, change the size of their plants. If short-run average total cost does not equal long-run average cost, their current plant is too large or too small to produce their current level of output at the least possible cost. They will want to change the size of the plants they are using, so their plant will be the ideal size to produce their current output. If price is lower than long-run average cost, firms in the industry will want to leave it; if price is above long-run average total cost, firms outside the industry will want to enter it.

Industry Adjustment to Falling Demand

A state of long-run equilibrium, such as that shown in Figure 4.7, exists only as long as outside conditions do not change. Suppose, though, that those conditions do change—for example, there is a long-run decrease in the market demand for the firm's product. Figure 4.8 shows what will happen.

FIGURE 4.8 LONG-RUN ADJUSTMENT TO DECLINING DEMAND



Part (a) represents a typical firm in a perfectly competitive industry; part (b) represents the industry as a whole. At first, both the firm and the industry are in long-run equilibrium at a price of \$5. Then something happens to shift the market demand curve leftward from D_1 to D_2 . In the short run, the price falls to \$4 at the intersection of D_2 and S_1 . The firm's short-run response is to move downward along its marginal cost curve. Because the price is still above average variable cost, the firm does not shut down. After a while, some firms (not the one shown) get tired of taking losses and leave the industry. This causes the market supply curve to shift toward S_2 and the market price to recover. The typical firm returns to the break-even point. The market has traced out part of its long-run supply curve as shown by the large arrow.

Part (a) of Figure 4.8 shows a set of cost curves for a typical firm. Part (b) is a supply-and-demand diagram for the market as a whole. The short-run industry supply curves shown are built up from those of the individual firms in the market (see Figure 4.6). The demand curves in part (b) are market demand curves.

Suppose that initially the short-run market supply and demand curves are in the positions S_1 and D_1 . The equilibrium price is \$5. Each firm takes this price as given and adjusts its output on that basis, producing 25 units. At that price and output, a typical firm would just break even. (Remember, though, that "breaking even" in the economic sense means earning enough to cover all costs, including the opportunity cost of capital.)

Now something happens—say, a change in consumer tastes or incomes—that shifts the demand curve to a new position, D_2 . The short-run result is a drop in the market price, to \$4. Each firm, being a price taker, will view the decline in price as beyond its control and will adjust to it as best it can. As shown in part (a) of Figure 4.8, this means cutting back output slightly in order to keep losses to a minimum, but not shutting down completely. Each firm's movement downward along its short-run marginal cost curve is what causes the movement of the market as a whole downward and to the left along the short-run supply curve.

However, the new situation cannot be a long-run equilibrium because each firm is operating at a loss. The firms' owners are not earning a normal rate of return, that is, they are not earning enough to cover the opportunity costs of keeping their capital invested in the industry. If the market demand curve shows no hope of shifting back to the right, some owners will pull their capital out of the industry. They may go bankrupt, abandoning their fixed assets to their creditors. They may sell their plant and equipment and get out while they can, or they may keep their firms running but convert their plants to make goods for other, more profitable markets.³

For the sake of the example, suppose that the typical firm shown in Figure 4.8 is not one of the first to leave. As some other firms withdraw, industry output falls by the amount of their output. The short-run market supply curve, which now comprises fewer individual supply curves, shifts to the left toward S_2 . As it does so, the market price begins to move upward along demand curve D_2 . When the price gets all the way back to \$5, the firms remaining in the industry will no longer be losing money. Firms will stop leaving the industry, and the market will have reached a new long-run equilibrium. At the new equilibrium price, short-run marginal cost, short-run average total cost, and long-run average cost will once again be equal.

This sequence of events has traced out a portion of the industry's *long-run supply curve*, as shown by the large horizontal arrow. A long-run supply curve for an industry shows the path along which equilibrium price and quantity move when there is a lasting change in demand. Movement along this curve requires enough time for firms to adjust the sizes of their plants or enter or leave the market.

Industry Adjustment to Rising Demand


When there is a long-run increase in demand, freedom of entry plays the same role that freedom of exit plays when demand falls. Such a case is shown in Figure 4.9. The starting position in this figure is the same as that in Figure 4.8. Short-run supply curve S_1 and demand curve D_1 result in an equilibrium price of \$5. The individual firm breaks even at an output of 25 units. Now watch what happens as the demand curve shifts to the right, to D_2 . The short-run result is an increase in the market price, to \$6. The typical firm adjusts to the new price by moving up along its short-run marginal cost curve. As all firms do this, the market moves up and to the right along short-run supply curve S_1 .

Again, however, the short-run position is not the new long-run equilibrium, because now all firms are making an economic profit. Entrepreneurs will soon spot this healthy, growing market as a good one in which to invest. Some of them may start new firms in this market; others may shift plant and equipment from making something else to making goods for this industry. Whether the entry is by new firms or by existing ones that are producing for this market for the first time, new entries will cause the supply curve to shift to the right, toward S_2 .

As the short-run market supply curve shifts to the right, the price falls. It does not fall far enough to drive the new entrants out of the market, but it does fall far enough to drive pure economic profits back to zero. Entry of firms into the market will stop, and the market will reach a new long-run equilibrium at the intersection of S_2 and D_2 .

Once again a portion of the long-run supply curve for the industry has been traced out, as shown by the large horizontal arrow in Figure 4.9. This long-run supply

FIGURE 4.9 LONG-RUN ADJUSTMENT TO AN INCREASE IN DEMAND



In this figure, both the firm and the industry start out in equilibrium at a price of \$5. Then something happens to shift the market demand curve rightward to D_2 . In the short run, the price rises to \$6 at the intersection of D_2 and S_1 . The firm's short-run response is to move upward along its marginal cost curve, earning better-than-normal profits. After a while, the high profits attract new firms into the industry. As those firms enter, the market supply curve shifts toward S_2 . Profits for the typical firm return to zero, and new firms stop entering the industry. Again the market has traced out part of its long-run supply curve as shown by the large arrow.

curve again is perfectly elastic. A rightward shift in the demand curve has, in the long run, produced an increase in quantity supplied but no rise in price.

As a final detail, note the importance of the assumption that there are no sunk costs in the industry. If entering the industry required specialized investments that could not be recovered later, firms would view them as opportunity costs when deciding whether to enter the market. They would not enter unless the price was high enough (and was expected to stay high enough) to give them a normal rate of return on the nonrecoverable investments. Once in the industry, however, those sunk costs would no longer affect decisions, according to the “bygones are bygones” principle. They would not count as part of the fixed (but not sunk) costs that must be covered by revenue for continued operation to be worthwhile. Thus, existing firms may stay in business indefinitely even if the price falls somewhat below what would be needed to attract new firms. When sunk costs are present, then, the industry supply curve is no longer a two-way street. Such an industry would, in effect, follow one supply curve when expanding and a different, lower one when contracting.

Although the theoretical model of perfect competition does not allow for sunk costs, such costs are common in the real world. Consider the history of that uniquely American entertainment establishment, the drive-in theater. In the early years after World War II, drive-in theaters were a growing business. With demand

high, many entrepreneurs entered the industry. Later, demand for this form of entertainment decreased. Even when market demand dropped well below the level needed to make it worthwhile to construct new drive-ins, existing operators stayed in business. They did so even though they were no longer earning enough to cover the original sunk cost of their screens and projection houses because those facilities could neither be moved nor converted to any other use. Only when demand fell still lower, so that revenues no longer covered recoverable fixed costs (such as the cost of land) and variable costs (such as wages, electricity, and film rentals) did drive-in theater operators finally leave the market.

The Elasticity of Long-Run Supply

The long-run industry supply curve in Figures 4.8 and 4.9 is perfectly elastic (See Appendix to Chapter 2 for a discussion on elasticities). Given such a curve, a change in demand affects only the equilibrium quantity, not the price, in the long run. However, that is not the only possible case. Supply curves that are positively sloped, negatively sloped, and U-shaped are also possible.

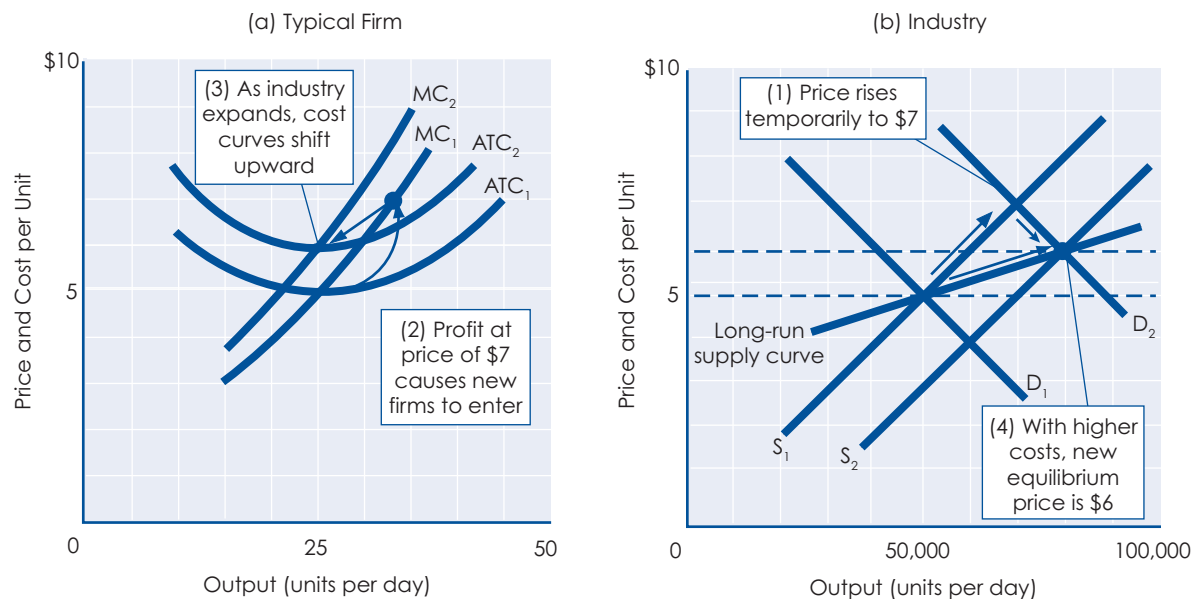
The shape of the long-run industry supply curve depends mainly on what happens to the industry's input prices in the long run as output expands. If the long-run supply curves for all inputs to the industry are perfectly elastic, the prices of those inputs will not change as the quantities of them demanded by the industry increase. It may also be that the industry uses such a small part of the total supply of each input that any change in input prices that does occur will be slight. For example, cookie stores use such a small part of the total supply of flour and eggs that expansion or contraction of such stores will have no perceptible effect on the market prices of those inputs. Industry output can therefore expand without affecting the costs of the individual firms, and the supply curve will be perfectly elastic.

Suppose, however, that the industry is a heavy user of relatively specialized inputs whose outputs can be boosted only at an increasing cost. An example is the home construction business, which uses a substantial portion of all lumber produced. An expansion of the construction industry will cause lumber suppliers to exhaust the lowest-cost stands of trees and begin harvesting higher-cost timber. The home construction industry also employs a significant proportion of all carpenters in the country. If the industry expands, carpenters' wages may have to rise relative to those of, say, auto mechanics in order to attract additional workers into the occupation.

Figure 4.10 shows what happens in such an industry as a permanent increase in demand causes output to expand. As in the preceding case, the shift in demand first pushes up price along the short-run supply curve. New firms enter the market. However, the expansion of the industry raises input prices. Each firm's short-run marginal cost and average total cost curves are shifted upward from MC_1 to MC_2 and from ATC_1 to ATC_2 as shown. As a result, the new long-run equilibrium is at a higher price than the initial equilibrium. The long-run industry supply curve, drawn through the two points of short-run equilibrium, therefore has a positive slope.

It is also possible for the price of an input to decrease as the industry's total output increases. For example, as sales of electronic equipment expand, the firms that make

FIGURE 4.10 A POSITIVELY SLOPED LONG-RUN INDUSTRY SUPPLY CURVE



In Figures 4.8 and 4.9, it was assumed that input prices do not change as industry output expands. This pair of diagrams shows what happens if industry expansion causes input prices to rise. As output expands, rising input prices push up the firm's marginal cost curve from MC_1 to MC_2 and its average total cost from ATC_1 to ATC_2 . The result is a new long-run equilibrium price that is higher than the initial price. The long-run industry supply curve thus has a positive slope.

components for such equipment may be able to use cheaper production methods. If that occurs, the short-run cost curves for all firms will drift downward as new firms enter the industry. The long-run supply curve then will be negatively sloped.

Finally, it is possible for these various forces to operate together. At first long-run supply is influenced by the falling price of one special input, but beyond a certain point some other special input becomes a limiting factor that causes the long-run supply curve to bend upward. The long-run industry supply curve then becomes U-shaped.

As we have seen, many variations are possible. Only through direct observation of the industry in question can we tell which possibility applies.

Market Performance Under Perfect Competition

Perfectly competitive markets have long earned high marks for several aspects of performance. In this section we look at market performance under perfect competition as it relates to the concepts of consumer and producer surplus, as well as the questions of *what* should be produced and *how* it should be produced.

Consumer Surplus

As we have shown earlier, a demand curve shows how much of a good consumers wish to purchase at any given price. We now turn to a different question to which the demand

curve also can provide an answer: How much will consumers be willing to pay for an additional unit of a good, given the quantity they already have?

The Demand Curve as Willingness to Pay

Let's return for a brief time to the individual consumer's demand curve. Figure 4.11 shows a demand curve for apples for a college student, Hannah Lee. Lee stocks up on snack foods at a local supermarket and often includes apples in her purchases. The demand curve given in the figure shows that the number of apples she eats each month depends on their price. Currently the price of an apple is \$.40. At this price, she buys ten per month. On other days she substitutes an orange or a banana.

The demand curve indicates that \$.40 is the maximum that Lee would be willing to pay for the tenth apple. If the price rose to \$.45, she would substitute some other fruit for the tenth apple. However, she would not cut out apples altogether. Although \$.40 is the maximum she is willing to pay for the tenth apple, she would not give up the ninth apple unless the price rose above \$.45. Similarly, she would be willing to pay up to \$.50 before giving up the eighth apple, up to \$.55 before giving up the seventh, and so on. The height of the demand curve at each point (emphasized here by a vertical bar) shows the maximum that she would willingly pay for each unit consumed. That maximum decreases as the quantity consumed increases, in accordance with the principle of diminishing marginal utility. In other words, as she consumes more apples, each additional apple offers less additional satisfaction and she therefore is not willing to pay as much as she did for previous apples within that time period.

Measuring the Surplus

Figure 4.11 shows the maximum that Lee is willing to pay for various quantities of apples, but it also shows that she need not actually pay this amount. At the going price of \$.40, she pays only a total of \$4.00 for the ten apples she buys each month. Except in the case of the last unit purchased, she gets each unit for less than what she would willingly have paid for it. The difference between what she would willingly have paid for each unit and the amount actually paid at the market price is called the **consumer surplus** for that unit. The consumer surplus on each unit is shown by the shaded portion of the corresponding vertical bar. For example, the surplus on the first apple for which Lee would have willingly paid \$.85, if necessary, is \$.45 because she actually paid only \$.40. The total consumer surplus on all units purchased is shown by the sum of the shaded portions of the bars. The area of the triangle between the demand curve and the market price is an approximate measure of consumer surplus.

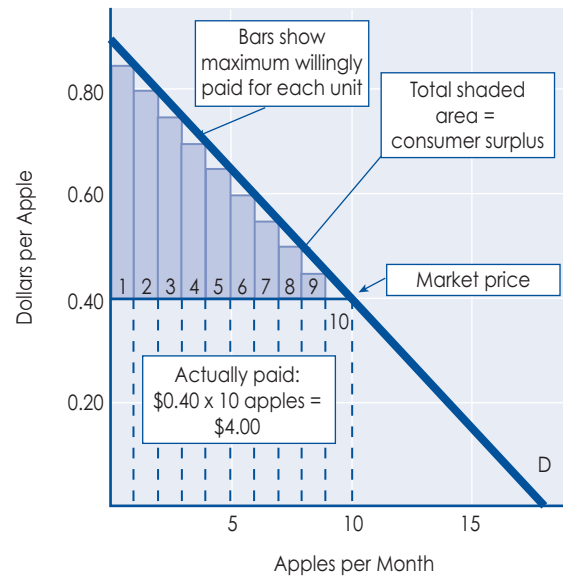
Consumer surplus

The difference between the maximum that a consumer would be willing to pay for a unit of a good and the amount that he or she actually pays

Consumer Surplus, Producer Surplus, and Gains from Exchange

The reasoning behind the notion of consumer surplus can be extended to the producers' side of the market as well. Consider Figure 4.12, which shows a typical market operating

FIGURE 4.11 CONSUMER SURPLUS



The height of a demand curve shows the maximum that this consumer would be willing to pay for an additional unit of a good. For example, she would be willing to pay up to \$.85 for the first apple bought each month but only \$.55 for the seventh. The maximum she would willingly pay for each unit is shown by a vertical bar. In this case, the market price is \$.40; thus, she buys 10 apples a month, paying a total of \$4.00. The difference between what she actually pays at the market price and the maximum she would have been willing to pay, shown by the shaded area, is called consumer surplus.

according to principles of supply and demand. The equilibrium market price is established at the point where the supply and demand curves cross. The demand curve, as we have seen, measures the maximum amount that consumers would be willing to pay for each unit sold; for example, they will pay no more than \$1.50 for the one thousandth. Consumer surplus is a measure of the difference between the maximum that consumers would have been willing to pay and what they actually pay at the market price.


Now turn to the supply curve. The height of the supply curve at any point represents the minimum that producers would willingly accept for the unit. For example, producers would be unwilling to accept less than \$.75 for the one thousandth unit sold. If they could not get at least that much, the producers would divert their resources to an alternative use rather than produce the one thousandth unit.

However, as the figure is drawn, producers receive the market price of \$1.00 per unit for all units sold, including the one thousandth. On that unit, they earn a producer surplus of twenty-five cents. The **producer surplus** earned on each unit is the difference between the market price and the minimum that the producers would have been willing to accept in exchange for that unit—the difference between \$1.75 for the one thousandth

Producer surplus

The difference between what producers receive for a unit of a good and the minimum they would be willing to accept

FIGURE 4.12 GAINS FROM EXCHANGE



This figure shows that both consumers and producers gain from exchange. Here the equilibrium market price is \$1.00 per unit. The demand curve shows the maximum that consumers would willingly pay for each unit. Consumers' gain from exchange takes the form of consumer surplus, shown by the area between the demand curve and the market price. The supply curve shows the minimum that producers would willingly accept rather than put their resources to work elsewhere. Producers earn a surplus equal to the difference between what they actually receive at the market price and the minimum they would have been willing to accept. The producer surplus is shown by the area between the supply curve and the market price. Assuming equilibrium is reached at the point of intersection of the two curves, total gains from exchange are thus the entire area between them up to the intersection.

unit in our example. The total producer surplus earned on all units is shown by the area between the supply curve and the market price.

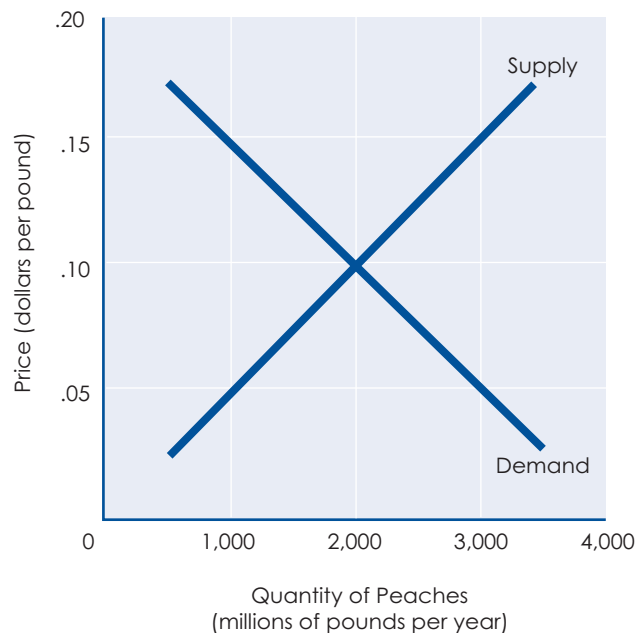
We see, then, that the concept of surplus in the market is symmetrical. Consumers buy the goods, except for the very last unit, for less than the maximum amount they would have been willing to pay; and producers sell the goods, except for the very last unit, for more than the minimum amount they would have been willing to accept. Thus, *both buyers and sellers gain from exchange*. That is why markets exist. As long as participation is voluntary, they make everyone who buys and sells in them better off than they would be if they did not participate. Assuming an equilibrium at the intersection of the supply and demand curves, as in Figure 4.12, the total of the mutual gains from exchange—consumer surplus plus producer surplus—is equal to the entire shaded triangle between the supply and demand curves to the left of their intersection point. Perfectly competitive markets result in total surplus being as large as possible. Economists refer to this as being allocatively efficient.

What Should Be Produced

When the concept of market performance was introduced, we used a diagram similar to Figure 4.13 to show the quantity of a good (peaches, in this case) that must be produced for a market to perform efficiently. The demand curve, we said, represents the amount consumers are willing to pay for an additional pound of peaches, given any level of output. That amount reflects the benefit of the marginal pound as perceived by consumers. The supply curve represents the amount suppliers require if they are to produce an additional pound of peaches. That amount corresponds to the opportunity cost to producers of supplying the marginal pound. As long as the demand curve is higher than the supply curve, trade at a price between the two curves can potentially benefit both parties. Accordingly, opportunities for mutually beneficial trades are exhausted (and efficiency is achieved) only if production is carried out to the point of intersection between the supply and demand curves, but not beyond that point.

In this chapter we have shown that the supply curve in a perfectly competitive market is the summation of the marginal cost curves of the individual firms. At any given market price, producers will supply the quantity that makes marginal cost equal to marginal revenue, which in turn is equal to the market price. Thus,

FIGURE 4.13 EFFICIENT OUTPUT UNDER PERFECT COMPETITION



This figure shows supply and demand curves for a perfectly competitive market for peaches. Under perfect competition, each firm's efforts to keep marginal cost equal to marginal revenue ensure that the industry will produce at some point on the supply curve. Equilibrium for the market as a whole can occur only at the point where the supply and demand curves intersect. That represents an efficient resolution to the question, What should be produced?

the price-quantity combination at which transactions take place in a competitive market will be some point on the supply curve.

For the market to end up at the right point on the supply curve, the market price must correspond to the intersection of the supply curve with the demand curve. To see why this will happen in a competitive market, we can put Chapter 2's analysis of market equilibrium together with this chapter's conclusions regarding perfect competition. From Chapter 2, we know that a price that is higher than the intersection of the supply and demand curves results in a surplus. Unplanned accumulation of inventory causes the price to fall. As the price falls, firms move down along the supply curve so as to keep marginal cost equal to marginal revenue. Similarly, a price that is lower than the intersection of the supply and demand curves results in a shortage. Depletion of inventories causes the market price to rise. Firms move up along the supply curve to keep marginal cost and marginal revenue equal. Thus, in a perfectly competitive market the equilibrium point will correspond to the intersection of the supply and demand curves. That is the efficient outcome.

Generalizing from these conclusions, we can see that in an economy in which all markets were perfectly competitive and in which there were no externalities the efficient quantity of every good would be produced. That would represent an efficient solution to the overall question of what should be produced—how many peaches, apples, tomatoes, and so on. Beginning from a situation in which the competitive markets for all products were in equilibrium, it would not be possible to substitute any one good for another (say, by producing more peaches at the expense of using fewer resources to produce apples) in a way that would make any person better off without making at least one other person worse off.

How to Produce

The preceding conclusion about what should be produced holds in both short- and long-run equilibrium. In addition, in the long run only, perfectly competitive markets ensure that each good is produced at the lowest possible cost—a key aspect of how goods should be produced.

To understand why goods are produced at the lowest cost in a situation of long-run competitive equilibrium, review Figures 4.8 and 4.9. In those figures the point of long-run equilibrium is shown to occur at the point where the typical firm operates at the minimum point of both the short-run average total cost curve and the long-run average cost curve.

Starting from such a point, a decrease in demand causes the market price to fall. In response, each firm reduces its output to the point where short-run marginal cost equals the new price. Although short-run marginal cost is lower at that output, short-run average total cost is higher because the firm moves up and to the left along the average total cost curve as output falls. Thus, at this point in the adjustment to falling demand, the given level of output is not being produced at the lowest possible cost. That is inefficient.

However, the inefficient situation does not last. Because short-run average total cost exceeds the market price, the firms suffer an economic loss, and some of them will

leave the industry. Assuming no further change in demand, as the number of firms in the industry decreases, each firm is able to increase its output and move back toward its point of least-cost production. Similar reasoning applies to the expansion of industry output in response to an increase in demand.

Under perfect competition, firms are led not only to produce at the lowest possible short-run average total cost, given the size of their plants, but also to select the correct plant size to minimize average cost in the long run. To see why, suppose that one firm had a plant that was not the optimal size. As we saw in the last chapter, the short-run average total cost curve for such a firm would be tangent to its long-run average cost curve at a point above and to the right of the point of long-run minimum cost (if the plant were too large) or above it and to the left (if the plant were too small). The firm with the wrong size plant would thus be at a cost disadvantage relative to its competitors. As competition drove the market price toward a level equal to minimum long-run average cost, the firm would either adjust the size of its plant to the cost-minimizing level or leave the industry because of economic losses.

Other Aspects of Market Performance

The tendency of perfectly competitive markets to produce the efficient quantity of each good and to produce those quantities at the lowest cost are important strengths of this market structure. In these respects, long-run equilibrium in perfectly competitive markets sets a standard against which the performance of other market structures can be judged.

Nevertheless, it would be claiming far too much to say that perfect competition has the best possible market performance under all conditions, and it would be premature to condemn every feature of real-world markets that differs from the structural characteristics of perfect competition. Before we write off all markets that are not made up exclusively of small firms, all markets in which products are not homogeneous, all markets in which newly entering firms encounter entry barriers or incur sunk costs, or all markets in which some participants know things that others do not, many questions must be asked. Among them are the following:

- Is it possible that other market structures equal or at least approximate the efficiency of perfect competition?
- How do alternative market structures perform when attention is focused on innovation and entrepreneurship rather than on equilibrium under conditions where technology and product characteristics are assumed to be unchanging?
- When markets fail to perform efficiently, what public-policy options are available? How should the dangers of government failure be weighed against the dangers of market failure?

Only when these additional aspects of the problem have been explored will we be in a position to make a balanced judgment of market performance under various market structures.



Summary

- 1. What characteristics define the structure of a market?** A *market structure* is defined in terms of the number and size of firms in the market, the nature of the product, ease of entry and exit, and availability of information. A *perfectly competitive market* has the following traits: (1) There are many buyers and sellers, each of which is small compared with the market as a whole; (2) the product is homogeneous; (3) it is easy to enter or leave the market; and (4) all buyers and sellers have equal access to information. Other market structures to be studied in this course include *monopoly, oligopoly, and monopolistic competition*.
- 2. What determines the profit-maximizing output level in the short run for a perfectly competitive firm?** In the short run the relationship between marginal cost and *marginal revenue* (price) determines the profit-maximizing output level for a perfectly competitive firm. The firm should expand output up to, but not beyond, the point at which marginal cost rises to the level of marginal revenue, provided that marginal revenue is at least equal to average variable cost at that point.
- 3. Under what conditions will a firm continue to operate even if it sustains a loss?** If marginal revenue is below average total cost at the point at which marginal cost and marginal revenue are equal, the firm cannot earn a profit. It will minimize loss in the short run by staying open if marginal revenue is above average variable cost. If marginal revenue is below average variable cost at the same point, the firm will minimize loss by shutting down.
- 4. How is a firm's short-run supply curve related to its cost curves?** The short-run supply curve for a perfectly competitive firm is the upward-sloping part of the marginal cost curve lying above its intersection with the average variable cost curve.
- 5. What are the conditions for long-run equilibrium in a perfectly competitive industry?**

Long-run equilibrium in a perfectly competitive industry requires (1) that price be equal to short-run marginal cost so that each firm is content with the level of output it is producing; (2) that short-run average total cost be equal to long-run average cost so that firms are satisfied with the size of their plants, given their output rate; and (3) that price be equal to long-run average cost so that there is no incentive for new firms to enter the industry or for existing firms to leave it.

- 6. What determines the shape of the long-run supply curve for a perfectly competitive industry?** A perfectly competitive industry adjusts to long-run changes in demand through exit of firms (in the case of a drop in market demand) or entry of new firms (in the case of a rise in market demand). If input prices do not change as the industry's output changes, the industry's long-run supply curve will be perfectly elastic. If input prices rise, the long-run supply curve will have a positive slope; if they fall, it will have a negative slope.
- 7. How efficiently do markets perform under perfect competition?** Under conditions of equilibrium, a perfectly competitive market produces a quantity of output that corresponds to the intersection of the market's supply and demand curves. This equilibrium generates the maximum total surplus possible for the market. In an economy in which all markets are in perfectly competitive equilibrium and there are no externalities, the question of what to produce is thus resolved efficiently. Also, in a situation of long-run equilibrium the output of a perfectly competitive market is produced at the lowest possible cost. This means that the question of how to produce is also resolved efficiently.

Key Terms

	Page #
Oligopoly	134
Monopolistic competition	134

Market structure	134
Perfect competition	134
Price taker	136
Marginal revenue	137
Consumer surplus	153
Producer surplus	154

Problems and Topics for Discussion

- Market structures** Give examples (other than those presented in the text) of industries that fit, or approximate, the market structures of perfect competition, monopoly, oligopoly, and monopolistic competition.
- Buyers as price takers** The concept of price taking can apply to buyers as well as to sellers. A price-taking buyer cannot influence prices by changing the amount purchased. Are you a price taker for the goods you buy? Can you give an example of a firm that might not be a price taker in the market in which it buys one or more of its inputs?
- Changes in fixed cost and the supply curve** Fieldcom buys some automated equipment to speed up production of its PDAs. The equipment adds \$500 per day to the firm's fixed costs, but it saves \$50 per unit in variable costs. Rework the graph in Figure 4.5 to show how the new equipment affects Fieldcom's supply curve. (You may want to rework part (a) of Figure 4.2 as a basis for the new supply curve.) What is the minimum price the firm must now charge to continue operating in the short run? What is the lowest price at which it can break even?
- Long- and short-run supply** Using the theory of perfect competition as presented in this chapter, explain why supply will be more responsive to changes in market prices in the long-run than in the short-run.
- Long-run supply with falling input prices** Figure 4.10 shows the long-run adjustment of a competitive industry to an increase in demand in the case in which input prices rise as industry output increases. Assume instead that input prices fall as output rises. Draw a new set of diagrams to show how a typical firm and the industry as a whole would respond to an increase in demand.

Case for Discussion

Independent Truckers as a Perfectly Competitive Industry

The next time you are out on the highway, take a look at the trucks that are passing you. You will see many that belong to large firms, such as UPS, that haul large numbers of small shipments all over the country on regular schedules. You will also see trucks that bear the names of companies such as Sears or Sun Oil, for which transportation of their own products is a small part of their total operations.

If you look closely, though, you will see that about one truck in four looks a little different. The tractors, many of which are brightly painted and highly chromed, often have sleepers attached to them. The trailers, often refrigerated, are likely to be filled with farm produce moving to market. These are the trucks of independent owner-operators, who move much of the nation's output of farm goods and some manufactured goods.

Each firm in this market consists of a person who owns and drives just one truck. There are tens of thousands of owner operators—just how many is hard to count.

From the shipper's point of view, one refrigerated truck is about as good as another as long as it is headed in the right direction. And most independent truckers will go wherever their loads take them.

Entry into the market is easy. Some people go into business with a used truck and as little as \$5,000. Most operators buy their trucks on credit. Exit is also easy—too easy, some say. Many independent truckers

go broke every year, and the number of firms rises and falls with the state of the economy.

Information is the lifeblood of the owner-operator. Truckers cannot make money unless they can find loads. Empty return runs after making a delivery are deadly to profitability. In the past, truck stop pay-phones and gossip over a cup of coffee were the main information channels, but the Internet has revolutionized the flow of information among shippers and truckers. Several competing web sites, some of which are free, others of which charge brokerage fees, match truckers to loads.

People who run the giant trucking companies that haul manufactured goods often look down their noses at the independent truckers with their loads of apples and potatoes. They call them gypsies or worse. But this system succeeds in putting fresh produce on dinner tables in every town every day.

QUESTIONS

1. In what ways does the independent trucking industry approximate the requirements of perfect competition? Are there any ways in which it does not meet those requirements?
2. On the average, the firms in a perfectly competitive industry earn no pure economic profits. However, average conditions do not always apply. What would you expect to happen to the profits of independent truckers when the economy enters a recession? When it enters a period of prosperity? What do you think would happen to the number of firms in the industry at such times?
3. Diesel fuel is a major input for independent truckers. What would you expect to happen to the profits of independent truckers and the number of firms in the industry as the price of fuel rises and falls? Outline the sequence of events in each case. (Drawing a graph may help.)

End Notes

1. Bear in mind that in the real world, it is hard to find firms that exactly fit the ideal type of perfect competition. For example, the description of Fieldcom in Chapter 3 implied that the firm's PDAs had special qualities of "ruggedization" that set them apart from other computers. This would violate the assumption of product homogeneity. For the purposes of this chapter, we will imagine that the idea of "ruggedized" PDAs did not work out, but that the Martins instead have discovered that they can survive by making "generic" PDAs that function just like others on the market.
2. This graph shows why we emphasize that profit maximization occurs where the *rising* section of the marginal cost curve intersects the marginal revenue curve. There is sometimes also an intersection of the *falling* section of the marginal cost curve with the \$300-per-unit marginal revenue curve, as is the case at about 2 units of output in Figure 4.3. That intersection is *not* a point of profit maximization, but rather, one of loss maximization.
3. The discussion of exit from a perfectly competitive market seems to pose a paradox: If all firms are *exactly* alike, why don't they all stay in the market as long as conditions are favorable, and then all leave the market at the same instant when conditions become unfavorable? However, real-world markets only approximate the conditions of perfect competition. In such markets, small differences in firms' circumstances of cost or demand, or in the temperaments of their owners, will cause some to leave the market before others do.

