10

Perfect Competition

Any firm must recognize two stark realities: no consumer is required to buy the goods that it produces, and even consumers who might want its goods may choose to buy from other firms instead. Firms in a situation of **perfect competition** sell their product in a market with other firms who produce identical or extremely similar products. As a result, if a firm in a perfectly competitive market raises the price of its product by so much as a penny, it will lose all of its sales to competitors. Perfectly competitive firms are sometimes known as **price takers,** because the pressure of other competing firms forces them to accept the prevailing price in the market as given. A related underlying assumption is that a perfectly competitive firm must be a relatively small player in the overall market, so that a perfectly competitive firm can increase or decrease its output without noticeably affecting the overall quantity supplied in the market.

perfect competition: Each firm faces many competitors that sell identical products.

price takers: A firm in a perfectly competitive market that must take the prevailing market price as given.

A perfectly competitive industry is a hypothetical extreme; after all, "perfect" is an extreme word. However, producers in a number of industries do face many competitor firms selling highly similar goods, in which case they must often act as price takers. For example, many farmers grow any given crop and crops grown by different farmers are largely interchangeable. Late in 2006, U.S. corn farmers were receiving about $2.00 per bushel and coffee (of the mild arabica variety) was selling for $1.14 a pound on world markets. A corn farmer who attempted to sell corn for $2.10 per bushel, or a coffee grower who attempted to sell for $1.30 a pound, would not have found any buyers. Similarly, gold was selling for $588 per ounce, which meant that sellers could not negotiate for $600 per ounce. Some manufactured products are similar enough that producers must sell at the price prevailing in the market. In the market for mid-sized color televisions, the quality differences are mostly not perceptible to the untrained eye, so price plays a major role in the decisions of purchasers. Gasoline is very similar between different gas stations. Blank CD-ROM disks are extremely similar regardless of the manufacturer.

This chapter examines how profit-seeking firms make decisions about how much to produce in perfectly competitive industries. Such firms will analyze their costs as discussed in the previous chapter, using the concepts of total cost, fixed cost, variable cost, average cost, marginal cost, and average variable cost. In the short run, the perfectly competitive firm will seek out the quantity of output where profits are highest, or if profits are not possible, where losses are smallest. In the long run, the perfectly competitive firms will react to profits by increasing production further and to losses by reducing production or shutting down. Firms will tailor their decisions about the quantity of inputs like labor and physical capital to purchase according to what they need to produce the profit-maximizing quantity of output at the lowest possible average cost.

Quantity Produced by a Perfectly Competitive Firm

A perfectly competitive setting really has only one major choice to make: namely, what quantity to produce. To understand why this is so, consider a different way of writing out the basic definition of profit:

Profit = Total revenue – Total cost
= (Price) (Quantity produced) – (Average cost) (Quantity produced)

Since a perfectly competitive firm must accept the price for its output as dictated by the forces of demand and supply, it cannot choose the price that it charges. Moreover, the perfectly competitive firm must also pay the market price for inputs to production like labor and physical capital, so it cannot control what it pays for the inputs that make up its costs of production. When the perfectly competitive firm chooses what quantity to produce, then this quantity—along with the prices prevailing in the market for output and inputs—will determine the firm's total revenue, total costs, and profits.

Comparing Total Revenue and Total Cost

As an example of how a perfectly competitive firm decides what quantity to produce, consider the case of the Floppy Production Shoppe, which makes boxes of computer diskettes. Exhibit 10-1 shows total revenue and total costs for the firm. The horizontal axis of the figure shows the quantity of boxes of disks produced; the vertical axis shows both total revenue and total costs, measured in dollars.

A perfectly competitive firm can sell as large a quantity as it wishes, as long as it accepts the prevailing market price. Thus, the total revenue line slopes up at an angle, with the slope of the line determined by the price. In this example, a 10-pack of 3.5-inch floppy computer diskettes costs $3.50. Sales of one box will bring in $3.50, two boxes will be $7, three units will be $10.50, and so on.

The total costs for the Floppy Production Shoppe, broken down into fixed and variable costs, are shown in the table that also appears in Exhibit 10-1. The total cost curve intersects with the vertical axis at a value that shows the level of fixed costs, and then slopes upward. (If it bothers you that these numbers don't seem more realistic, you may wish to imagine that these total cost and quantity figures represent thousands: that is, fixed costs of the plant are $62,000, and the total cost of producing 10,000 boxes of

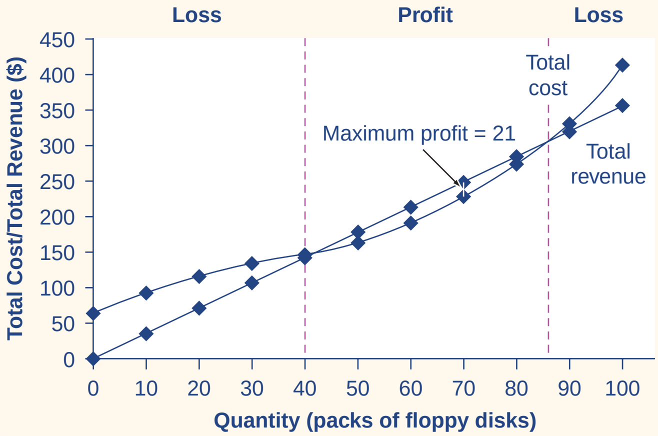

Quantity	Total Cost	Fixed Cost	Variable Cost	Total Revenue	Profit
0	62	62	—	0	–62
10	90	62	28	35	–55
20	110	62	48	70	–40
30	126	62	64	105	–21
40	144	62	82	140	–4
50	166	62	104	175	+9
60	192	62	130	210	+18
70	224	62	162	245	+21
80	264	62	202	280	+16
90	324	62	262	315	–9
100	404	62	342	350	–54

Exhibit 10-1 Total Cost and Total Revenue at the Floppy Production Shoppe

Total revenue for a perfectly competitive firm is a straight line sloping up. The slope is determined by the price of the good. Total cost also slopes up, but with some curvature. Total cost intersects the vertical axis at the level determined by fixed costs. At higher levels of output, total cost begins to slope upward more steeply because of diminishing marginal returns. At levels of output where total cost is higher than total revenue, like 10 or 20, the firm would suffer a loss. At levels of output where total revenue is higher than total costs, like 60 or 70, the firm would earn profits. The maximum profit will occur at the quantity where the gap of total revenue over total cost is largest.

floppy disks is $90,000. The shape of the total revenue and total cost picture does not change if both quantity and cost are expressed with the extra three zeros.)

Based on its total revenue and total cost curves, a perfectly competitive firm like the Floppy Production Shoppe can calculate the quantity of output that will provide the highest level of profit. At any given quantity, total revenue minus total cost will equal profit. On the figure, the vertical gap between total revenue and total cost represents either profit (if total revenues are greater that total costs at a certain quantity) or losses (if total costs are greater that total revenues at a certain quantity). In this example, total costs will exceed total revenues at output levels from 0 to 40, and so over this range of output the firm will be making losses. At output levels from 50 to 80, total revenues exceed total costs, so the firm is earning profits. But then at an output of 90 or 100, total costs again exceed total revenues and the firm is making losses. Total profits appear in the final column of the table. The highest total profits in the table occur at an output of 70, when profits will be $21.

A higher price would mean a steeper slope for the total revenue curve, so that total revenue was higher for every quantity sold; a lower price would mean a flatter slope, so that total revenue was lower for every quantity sold. What happens if the price drops low enough so that the total revenue line is completely below the total cost curve; that is, at every level of output, total costs are higher than total revenues? Then the best that the firm can do is to suffer losses, but a profit-maximizing firm will prefer the quantity of output where total revenues come closest to total costs and thus where the losses are smallest.

Comparing Marginal Revenue and Marginal Costs

Firms often do not have the necessary data they would need to draw a complete total cost curve for all different levels of production. They can't be sure of what total costs would look like if they, say, doubled production or cut production in half, because they haven't tried it. Instead, firms experiment; that is, they produce a slightly greater or lower quantity and observe how profits are affected. In economic terms, this practical approach to maximizing profits means looking at how changes in production affect marginal revenue and marginal cost.

marginal revenue: The additional revenue gained from selling one more unit.

Exhibit 10-2 presents the marginal revenue and marginal cost curves based on the total revenue and total cost curves in Exhibit 10-1. The **marginal revenue** curve shows the additional revenue gained from selling one more unit. For a perfectly competitive firm,

Exhibit 10-2 Marginal Revenues and Marginal Costs at the Floppy Production Shoppe

For a perfectly competitive firm, the marginal revenue (MR) curve is a horizontal straight line because it is equal to the price of the good. Marginal cost (MC) is sometimes first downward-sloping, if there is a region of increasing marginal returns at low levels of output, but is eventually upward-sloping at higher levels of output as diminishing marginal returns kick in. If the firm is producing at a quantity where MR > MC, like 40 or 50, then it can increase profit by increasing output because the marginal revenue is exceeding the marginal cost. If the firm is producing at a quantity where MC > MR, like 90 or 100, then it can increase profit by reducing output because the reductions in marginal cost will exceed the reductions in marginal revenue. The firm's profit-maximizing choice of output will occur where MR = MC (or at a choice close to that point).

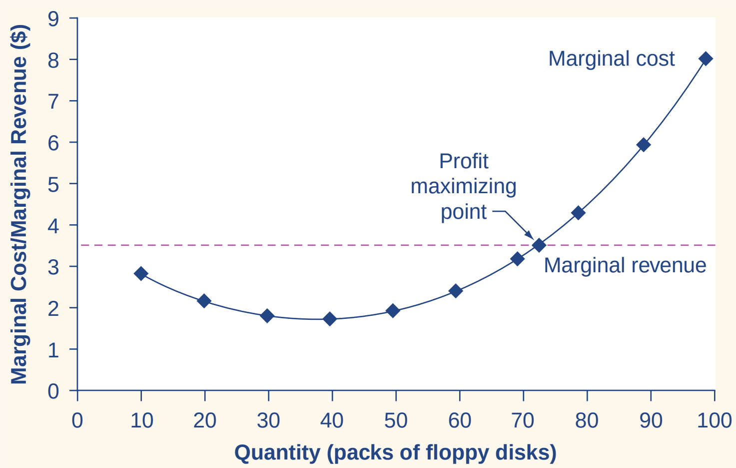

Quantity	Total Cost	Fixed Cost	Variable Cost	Marginal Cost	Total Revenue	Marginal Revenue
0	$62	$62		—	—	—
10	$90	$62	$28	$2.80	$35	$3.50
20	$110	$62	$48	$2.00	$70	$3.50
30	$126	$62	$64	$1.60	$105	$3.50
40	$144	$62	$82	$1.80	$140	$3.50
50	$166	$62	$104	$2.20	$175	$3.50
60	$192	$62	$130	$2.60	$210	$3.50
70	$224	$62	$162	$3.20	$245	$3.50
80	$264	$62	$202	$4.00	$280	$3.50
90	$324	$62	$262	$6.00	$315	$3.50
100	$404	$62	$342	$8.00	$350	$3.50

the marginal revenue curve is drawn as a flat line equal to the price level, which in this example is $3.50 per box of floppy disks. The marginal revenue gained from selling one additional box of floppy disks is just the price of that unit. Since a perfectly competitive firm is a price-taker, it can sell whatever quantity it wishes at the market-determined price. Marginal cost, the cost per additional unit sold, is calculated by dividing the change in total cost by the change in quantity. In this example, marginal cost at first declines as production increases from 10 to 20, which represents the area of increasing marginal returns that is not uncommon at low levels of production. But then marginal costs start to increase, displaying the typical pattern of diminishing marginal returns.

In this example, the marginal revenue and marginal cost curves cross at a price of $3.50 and a quantity between 70 and 80 produced. If the firm started out producing at a level of 60, and then experimented with increasing production to 70, marginal revenues from the increase in production would exceed marginal costs—and so profits would rise. However, if the firm then experimented further with increasing production from 70 to 80, it would find that marginal costs from the increase in production are greater than marginal revenues, and so profits would decline.

The profit-maximizing choice for a perfectly competitive firm will occur at the point where marginal revenue is equal to marginal cost—that is, where MR = MC. A profit-seeking firm should keep expanding production as long as MR > MC. But at the level of output where MR = MC, the firm should recognize that it has achieved the highest possible level of profits and that expanding production into the zone where MR < MC will only reduce profits. Because the marginal revenue received by a perfectly competitive firm is equal to the price P, so that P = MR, the profit-maximizing rule for a perfectly competitive firm can also be written as a recommendation to produce at the quantity where P = MC. (If the firm's choices for quantity of output don't include a choice where MR is exactly equal to MC, then the highest output level where MR > MC will be the profit-maximizing choice.)

Marginal Cost and the Supply Curve

For a perfectly competitive firm, the marginal cost curve is identical to the firm's supply curve. To understand why this perhaps surprising insight holds true, first think about what the supply curve means. A firm checks the market price and then looks at its supply curve to decide what quantity to produce. Now think about what it means to say that a firm will maximize its profits by producing at the quantity where P = MC. This rule means that the firm checks the market price and then looks at its marginal cost curve to determine the quantity to produce. In other words, the marginal cost curve and the firm's supply curve contain the same information—that is, they both tell the firm what quantity to produce, given the market price.

As discussed back in Chapter 4, many of the reasons that supply curves shift relate to underlying changes in costs. For example, a lower price of key inputs or new technologies that reduce production costs cause supply to shift to the right; in contrast, lousy weather or added government regulations can add to costs of certain goods in a way that causes supply to shift to the left. These shifts in the firm's supply curve can also be interpreted as shifts of the marginal cost curve. A shift in costs of production that increases marginal costs at all levels of output—and shifts MC to the left—will cause a perfectly competitive firm to produce less at any given market price. Conversely, a shift in costs of production that decreases marginal costs at all levels of output will shift MC to the right and as a result, a competitive firm will choose to expand its level of output at any given price.

Profits and Losses with the Average Cost Curve

If the price that a firm charges is higher than its average cost of production for that quantity produced, then the firm will earn profits. Conversely, if the price that a firm charges is lower than its average cost of production, the firm will suffer losses. Exhibit 10-3 illustrates three situations: (a) where price intersects marginal cost at a level above the average cost curve, (b) where price intersects marginal cost at a level equal to the average cost curve, and (c) where price intersects marginal cost at a level below the average cost curve.

First consider a situation where the price is equal to $4 for a box of floppy diskettes. The rule for a profit-maximizing perfectly competitive firm is to produce the level of output where P = MR = MC, so the Floppy Production Shoppe will produce a quantity of 80, which is labeled as q in Exhibit 10-3*a*. The firm's total revenue at this price will be shown by the large shaded rectangle, which represents a price of $4 per box and a quantity of 80. Total costs will be the quantity of 80 times the average cost of $3.30, which is shown by the lightly shaded rectangle. Thus, profits will be the heavily shaded rectangle which represents the large rectangle of total revenues minus the rectangle of total costs. It would be calculated as:

Profit = Total revenue – Total cost

$56 = (800) ($4.00) – (80) ($3.30)

Now consider Exhibit 10-3*b*, where the price has fallen to $3.20 for a box of diskettes. Again, the perfectly competitive firm will choose the level of output where P = MR = MC, but in this case, the quantity will be 70. At this price and output level, where the marginal cost curve is crossing the average cost curve, the price received by the firm is exactly equal to its average cost of production. Thus, the total revenue of the firm is exactly equal to the total costs of the firm, so that the firm is making zero profit.

In Exhibit 10-3*c*, the market price has fallen still further to $2.20 for a box of diskettes. Again, the perfectly competitive firm will choose the level of output where P = MR = MC, and in this case the firm will choose to produce a quantity of 50. At this price and output level, the price level received by the firm is below the average cost of production. The total costs for the firm will be $166 (that is, the average cost of $3.32 multiplied by a quantity of 50), which appear as the large rectangle. Total revenue will be $110 (price of $2.20 multiplied by a quantity of 50), which is the smaller shaded rectangle. The firm's losses of $56 are shown by the rectangle created by subtracting the total revenue rectangle from the total cost rectangle.

If the market price received by a perfectly competitive firm leads it to produce at a quantity where the price is greater than average cost, the firm will earn profits. If the price received by the firm causes it to produce at a quantity where price equals marginal cost, at the point where the MC curve crosses the AC curve, then the firm earns zero profits. Finally, if the price received by the firm leads it to produce at a quantity where the price is less than average cost, the firm will earn losses.

The Shutdown Point

The possibility that a firm may earn losses raises a question: Why can't the firm avoid losses by shutting down and not producing at all? The answer is that shutting down can reduce variable costs to zero, but in the short run, shutting down does not reduce fixed costs. As a result, if the firm produces a quantity of zero, it would still make losses because it would still need to pay for its fixed costs. Thus, when a firm is experiencing losses it must face a question: will its losses increase or decrease if it shuts down?

Exhibit 10-3 Price and Average Cost at the Floppy Production Shoppe

(a) Price is above average cost. Price is at a level where producing at the quantity where P = MC leads to a price that is above average cost. In this case, the firm is earning a profit. Total revenue is the quantity of 80 times the price of $4, or $320, shown by the overall shaded box. Total cost is the quantity of 80 times an average cost of $3.30, or $264, shown by the bottom shaded box. The leftover rectangle where total revenue exceeds total cost is the profit earned.

(b) Price equals average cost. The price is now at a level where producing at the quantity where P = MC leads to a price that is equal to the average cost. Total revenue is now a quantity of 70 times a price of $3.20. Total cost is the same: a quantity of 70 times an average cost of $3.20. Zero profit is being earned in this situation.

(c) Price is lower than average cost. At a price of $2.20, when the firm produces at a quantity where P = MC, the price is below average cost. Here, the firm is suffering losses. Total costs are the large rectangle with a quantity of 50 and a price of $3.32, for total costs of $166. Total revenues are a quantity of 50 and a price of $2.20, or $110, shown by the smaller shaded box. The leftover rectangle on top thus shows the losses; that is, the amount that total cost exceeds total revenue.

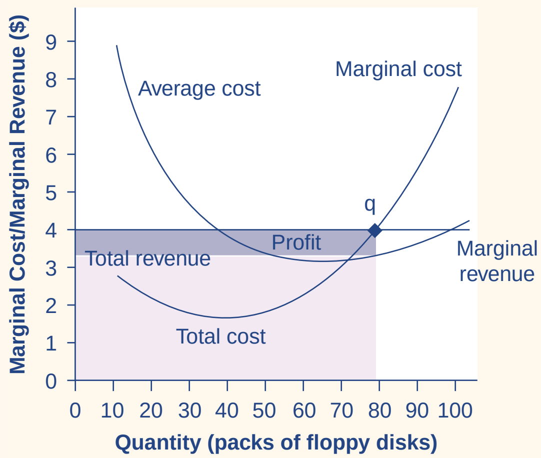

(a)

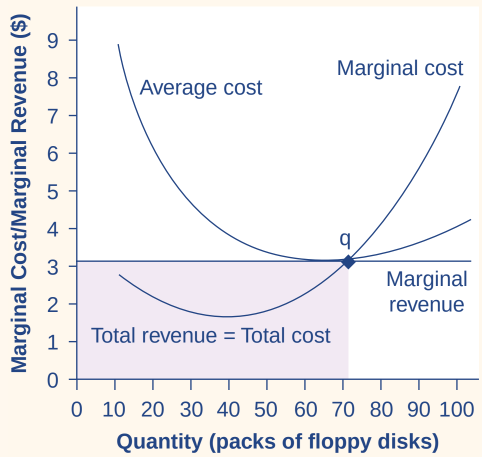

(b)

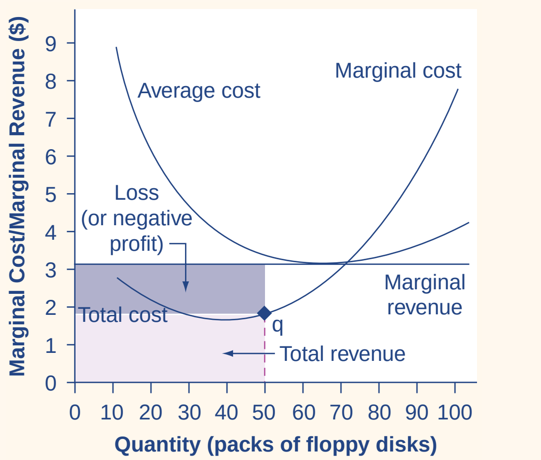

(c)

Quantity	Total Cost	Fixed Cost	Variable Cost	Marginal Cost	Average Cost
0	$62	$62	—	—	—
10	$90	$62	$28	$2.80	$9.00
20	$110	$62	$48	$2.00	$5.50
30	$126	$62	$64	$1.60	$4.20
40	$144	$62	$82	$1.80	$3.60
50	$166	$62	$104	$2.20	$3.32
60	$192	$62	$130	$2.60	$3.20
70	$224	$62	$162	$3.20	$3.20
80	$264	$62	$202	$4.00	$3.30
90	$324	$62	$262	$6.00	$3.60
100	$404	$62	$342	$8.00	$4.04

As an example of this decision, consider the situation of the Yoga Center, which has signed a contract to rent space that costs $10,000 per month. If the firm continues to operate, its marginal costs for hiring yoga teachers is $15,000 for the month. If the firm shuts down, it must still pay the rent, but it would not need to hire labor. Exhibit 10-4

Exhibit 10-4 Should the Yoga Center Shut Down Now or Later?

Scenario 1
The firm shuts down now, so revenues are zero but the firm must pay fixed costs of $10,000.

$$\text{Total revenue} - (\text{fixed cost} + \text{variable cost}) = \text{profit}$$
$$0 - 10{,}000 = -10{,}000$$

Scenario 2
The firm keeps operating, variable costs are $15,000 and revenues are $10,000.

$$\text{Total revenue} - (\text{fixed cost} + \text{variable cost}) = \text{profit}$$
$$10{,}000 - (10{,}000 + 15{,}000) = -15{,}000$$

Conclusion: The firm should shut down now.

Scenario 3
The firm keeps operating, variable costs are $15,000 and revenues are $20,000.

$$\text{Total revenue} - (\text{fixed cost} + \text{variable cost}) = \text{profit}$$
$$20{,}000 - (10{,}000 + 15{,}000) = -5{,}000$$

Conclusion: The firm should shut down later.

shows three possible scenarios. In the first scenario, the Yoga Center shuts down, in which case it faces losses of $10,000. In the second scenario, the Yoga Center remains open, earns revenues of $11,000 for the month, and experiences losses of $14,000. In the third scenario, the Yoga Center remains open and earns revenues of $20,000 for the month, and experiences losses of $5,000. In all three cases, the Yoga Center loses money. In all three cases, when the rental contract expires in the long run, the store should exit this business. But in the short run, in scenario 2, the store's losses increase because it remains open, so the store should shut down immediately. In contrast, in scenario 3 the revenue that the firm can earn is high enough that the store's losses diminish when it remains open, so the store should remain open in the short run.

This example suggests that the key factor is whether the firm can at least cover its variable costs by remaining open. Exhibit 10-5 illustrates this lesson for the example of the Floppy Production Shoppe, by adding the average variable cost curve to the marginal cost and average cost curves. At a price of $2.20 per box, shown in Exhibit 10-5*a*, the firm produces at a level of 50. It is making losses of $56 (as explained earlier), but price is above average variable cost and so the firm continues to operate. However, if the price declined to $1.80 per box, as shown in Exhibit 10-5*b*, if the firm applied its rule of producing where P = MR = MC, it would produce a quantity of 40. This price is below average variable cost for this level of output. At this price and output, total revenues would be $72 (quantity of 40 times price of $1.80) and total cost would be $144, for overall losses of $72. If the firm shuts down its must only pay its fixed costs of $62, so shutting down is preferable to selling at a price of $1.80 per box.

shutdown point: When the revenue a firm receives does not cover its average variable costs, the firm should shut down immediately: the point where the marginal cost curve crosses the average variable cost curve.

The intersection of the average variable cost curve and the marginal cost curve, which shows the price where firm lacks enough revenue to cover its variable costs, is called the **shutdown point.** If the perfectly competitive firm can charge a price above the shutdown point, then the firm is at least covering its average variable costs, and it

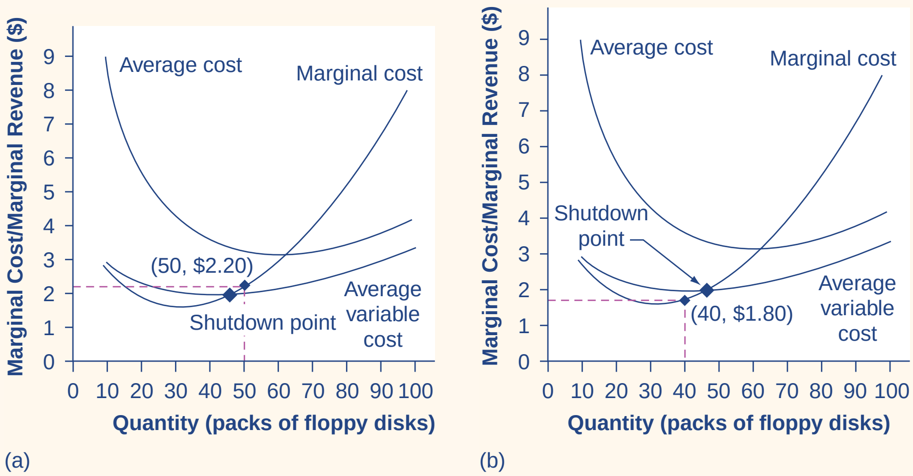

Quantity	Total Cost	Fixed Cost	Variable Cost	Marginal Cost	Average Cost	Average Variable Cost
0	$62	$62	—	—	—	—
10	$90	$62	$28	$2.80	$9.00	$2.80
20	$110	$62	$48	$2.00	$5.50	$2.40
30	$126	$62	$64	$1.60	$4.20	$2.13
40	$144	$62	$82	$1.80	$3.60	$2.05
50	$166	$62	$104	$2.20	$3.32	$2.08
60	$192	$62	$130	$2.60	$3.20	$2.16
70	$224	$62	$162	$3.20	$3.20	$2.31
80	$264	$62	$202	$4.00	$3.30	$2.52
90	$324	$62	$262	$6.00	$3.60	$2.91
100	$404	$62	$342	$8.00	$4.04	$3.42

Exhibit 10-5 The Shutdown Point for the Floppy Production Shoppe

(a) The price is $2.20. Producing at the quantity where P = MC leads to a situation where price is below average cost but above average variable cost. Thus, the firm is suffering losses, but because it is covering its variable costs, it will continue to operate.

(b) Price drops to $1.80, which leads to a quantity of output where price is less than average variable cost. Because the firm is now not even covering its variable costs, it will shut down immediately. The point where MC crosses AVC is called the shutdown point.

should limp ahead even if it is making losses in the short run, since at least those losses will be smaller than if the firm shuts down immediately. However, if the firm is receiving a price below the price at the shutdown point, then the firm is not even covering its variable costs. In this case, staying open is making the firm's losses larger, and it should shut down immediately.

Short-Run Outcomes for Perfectly Competitive Firms

The average cost and average variable cost curves divide the marginal cost curve into three segments, as shown in Exhibit 10-6. At the market price, which the perfectly competitive firm accepts as given, the profit-maximizing firm chooses the output level where price or marginal revenue, which are the same thing for a perfectly competitive firm, is equal to marginal cost: P = MR = MC.

First consider the upper zone, where prices are above the level where marginal cost MC crosses AC at the zero profit point. At any price above that level, the firm will earn profits in the short run. If the price falls exactly on the zero profit point where MC and AC curves cross, then the firm earns zero profits. Next consider a price which falls into the zone between the zero profit point where MC crosses AC and the shutdown point

Exhibit 10-6 Profit, Loss, Shutdown

The marginal cost curve can be divided into three zones, based on where it is crossed by the average cost and average variable cost curves. The point where MC crosses AC is called the zero-profit point. If the market price is at a level higher than the price at the zero-profit point, then price will be above the average cost curve and the firm is earning profits. If the price is exactly at the zero-profit point, then the firm is making zero profits. If the price below the price at the shutdown point, then price will fall below the average cost curve and the firm will be making losses. The point where MC crosses AVC is called the shutdown point. If price falls in the zone between the shutdown point and the zero-profit point, then the firm is making losses but will continue to operate in the short run, since it is covering its variable costs. However, if price falls below the price at the shutdown point, then the firm will shut down immediately, since it is not even covering its variable costs.

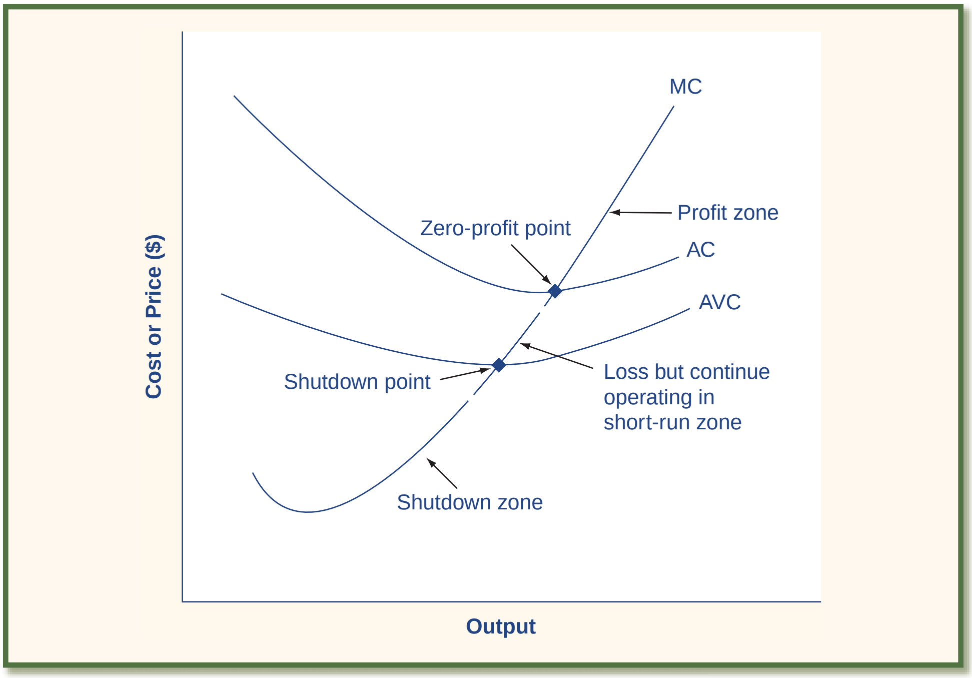

where MC crosses AVC. In this case, the firm will be making losses in the short run—but since the firm is more than covering its variable costs, the losses are smaller than if the firm shut down immediately. Finally, consider a price below the shutdown point where MC crosses AVC. At any price like this one, the firm will shut down immediately, because it cannot even cover its variable costs by operating.

Entry and Exit in the Long Run

The line between the short run and the long run can't be defined precisely with a stopwatch, or even with a calendar. It varies according to the specific business. But in the long run, the firm can adjust all factors of production.

In a competitive market, profits are a red cape that incites businesses to charge. If a business is making a profit in the short run, it has an incentive to expand existing factories or to build new ones. New firms may start production, as well. **Entry** is the long-run process of expanding production in response to a sustained pattern of profit opportunities.

entry: The long-run process of firms beginning and expanding production in response to a sustained pattern of profits.

Losses are the black thundercloud that causes businesses to flee. If a business is making losses in the short run, it will either keep limping along or just shut down, depending on whether its revenues are covering its variable costs. But in the long run, firms that are facing losses will shut down at least some of their output, and some firms will cease production altogether. **Exit** is the long-run process of reducing production in response to a sustained pattern of losses.

exit: The long-run process of firms reducing production and shutting down in response to a sustained pattern of losses.

How Entry and Exit Lead to Zero Profits

No perfectly competitive firm acting alone can affect the market price. However, the combination of many firms entering or exiting the market will affect overall supply in the

Business Failure

In the model of perfectly competitive firms, those that consistently can't make money will "exit," which is a nice bloodless word for a more painful process. When a business fails, after all, workers lose their jobs, investors lose their money, and owners and managers can lose their dreams, too. Many businesses fail. From 2000–2001, for example, 585,000 firms were born in the United States, but 553,000 firms died. About 95% of these business births and deaths involved small firms with fewer than 20 employees.

Sometimes a business fails because of poor management or workers, or because of tough domestic or foreign competition. Businesses also fail from a variety of causes that might best be summarized as bad luck. Conditions of demand and supply in the market shift in an unexpected way, so that the prices that can be charged for outputs fall or the prices that need to be paid for inputs rise. With millions of businesses in the U.S. economy, even a small fraction of them failing will affect many people—and business failures can be very hard on the workers and managers directly involved. But from the standpoint of the overall economic system, business deaths are sometimes a necessary evil if a market-oriented system is going to offer a flexible mechanism for satisfying customers, keeping costs low, and inventing new products.

market. In turn, a shift in supply for the market as a whole will affect the market price. Entry and exit are the driving forces behind a process that, in the long run, pushes the price in a perfectly competitive market to be at the zero-profit point where the marginal cost curve crosses the average cost curve.

To understand how short-run profits for a perfectly competitive firm will evaporate in the long run, imagine that the market price starts at level where, when the perfectly competitive firm applies the rule P = MR = MC, it produces a quantity where price is above the average cost curve and the firm is earning profits. However, these profits attract entry. Entry of many firms causes the market supply curve to shift to the right, so that a greater quantity is produced at every market price. The shift of the supply curve to the right pushes down the market price. Entry will continue to shift supply to the right until the price is driven down to the zero-profit level, where no firm is earning profits.

Short-run losses will fade away by reversing this process. Now imagine that the market price starts at level where, when the perfectly competitive firm applies the rule P = MR = MC, it produces a quantity where price is below the average cost curve and the firm is experiencing losses. These losses lead some firms to exit. The process of exit by a number of firms causes the market supply curve to shift to the left, so that a lesser quantity is produced at every market price, and the market price rises. The process of exit will continue as long as firms are making losses, until the price rises to the zero-profit level, where firms are no longer suffering losses.

Thus, while a perfectly competitive firm can earn profits in the short run, in the long run the process of entry will push down prices until they reach the zero profit level. Conversely, while a perfectly competitive firm may earn losses in the short run, firms will not continually lose money. In the long run, firms making losses are able to escape from their fixed costs, and their exit from the market will push the price back up to the zero profit level. In the long run, this process of entry and exit will drive the price in perfectly competitive markets to the zero profit point at the bottom of the AC curve, where marginal cost crosses average cost.

Economic Profit vs. Accounting Profit

This talk of "zero profit" in the long run must worry anyone who is thinking about starting a business—or making a financial investment in one. But the conclusion that perfectly competitive firms end up with zero profits in the long run is less dire than it may sound, because profit means something different to economists than it does to businesses. The difference arises because economists insist on considering all the costs of production, including opportunity costs.

To understand the difference between the economic view of profit and the business or accounting view of profit, think about a business where the owner plunks down $1 million to pay for fixed costs, and at the end of a year, after total revenues and total costs are tallied, the owner has a business that is worth $1 million and has earned a return of $50,000, or 5% of the original investment. How large a profit has the owner earned? To a business owner and the accountant, the answer is clear: the profit is the 5% or $50,000 that was earned. To put it another way, the **accounting profit** is total revenues minus the firm's expenditures on costs.

accounting profit: Total revenues minus the firm's costs, without taking opportunity cost into account.

To an economist, the question is more difficult, because the economist wants to take all costs, including opportunity cost, into account. Say that the business owner could have invested the money in another business, run by someone else, and it would have paid an 8% return. From the economic point of view, this opportunity cost of how the financial capital could have been invested must be included as one of the firm's costs. Thus, **economic profit** is measured by total revenues minus all of the firm's costs, including both expenditures and opportunity costs of financial capital. From this perspective, the business owner has not experienced a gain of 5% for the year, but a loss of 3% (that is, 5% accounting profit minus 8% opportunity cost).

economic profit: Total revenues minus all of the firm's costs, including opportunity costs.

For the economy as a whole, financial investors earn an average or "normal" level of profit. That normal level of profit determines the opportunity cost of financial investments. Thus, an economist who says that perfectly competitive firms will earn "zero profit" in the long run really means "zero profit after opportunity cost has been taken into account," or to put it differently, that the firm will earn a normal level of profit. When an economist says that a firm will earn positive profits, the economist really means that the profits will be above the normal level—what an accountant would refer to as extraordinary or high profits. Conversely, losses to an economist include both returns that are negative in an accounting sense and also returns that are slightly positive in an accounting sense, but lower than the opportunity cost of financial capital.

Although perfectly competitive firms will earn zero economic profits—or a normal rate of profit—in the long run, it's important to remember that not all industries are perfectly competitive, and the long run, by definition, takes some time to arrive. In the real world, businesses will often earn economic profits or losses in the short run.

The Economic Function of Profits

In popular discussions, the term "profit" is sometimes pronounced as if the speaker is sucking on a lemon. But profits are not a dirty word, any more than "high prices" or "low prices" are dirty words. Profits have an economic rationale, which is built upon the incentives that they provide for entry and exit.

High profits are a messenger. They transmit the message that the willingness of people to pay for that good is higher than the average cost of production, and so society as a whole will benefit if more resources are allocated to production of that certain good or service. Losses are a messenger, too. They carry the message that fewer resources should be allocated to production of a certain good or service, because the benefits people are receiving from that production—as measured by their willingness to pay for it—is below the cost of the resources used in the production process.

Profits also serve as a messenger to businesses to seek out the most cost-effective methods of production. As profits lead to entry or losses lead to exit, perfectly competitive firms are pressured to seek out the minimum point of the average cost curve, which is the point where economic profits are zero and also the point where output is being produced at the lowest possible average cost.

Factors of Production in Perfectly Competitive Markets

A firm in a perfectly competitive market must accept the market price for its output, and must also pay the prevailing market price for its inputs like labor and physical capital. However, a perfectly competitive firm can decide what quantity of those inputs to use. The earlier discussion of how perfectly competitive firms make decisions about the quantity of output to produce was built on a distinction between short-run decisions, where certain inputs costs like physical capital investment were fixed, and long-run decisions where all inputs to production could vary. In the same spirit, let's first consider how a perfectly competitive firm will determine the quantity that it uses of a variable input like labor in the short run, and then how it will determine its quantity of physical capital and its overall method of production in the long run.

The Derived Demand for Labor

When a profit-maximizing firm considers how much labor to hire in the short run, it may first calculate the profit-maximizing quantity of output, and then figure out the quantity of labor needed to produce that quantity of output. For this reason, the demand for labor (and other inputs) is sometimes called a "derived demand"; that is, the demand for labor is derived from the firm's decision about how much to produce.

For an example of this decision process operates in the short run, consider the Cheep-Cheep Company in Exhibit 10-7, a firm in a perfectly competitive industry that hires workers to assemble birdhouses. Assume that the cost of hiring a worker is $100 per day and that fixed costs of the birdhouse workshop are $7,000. The quantity of output for a profit-seeking perfectly competitive firm will be determined by applying the rule that production up to the point where P = MC. Thus, if the price for a birdhouse is $6, then the profit-maximizing level of output (from the choices in the table) is 85, and the derived demand for labor to produce that desired level of output is 5 workers. If a surge in demand for birdhouses drives the market price up to $10, then the optimal level of output would become 96, and the derived demand for labor rises to six workers.

Labor	Output	Fixed Cost	Variable Cost	Total Cost	Marginal Cost
0	0	$7,000	$0	$7,000	—
1	12	$7,000	$100	$7,100	$8.33
2	28	$7,000	$200	$7,200	$6.25
3	48	$7,000	$300	$7,300	$5.00
4	68	$7,000	$400	$7,400	$5.00
5	85	$7,000	$500	$7,500	$5.88
6	96	$7,000	$600	$7,600	$9.09
7	105	$7,000	$700	$7,700	$11.11

Exhibit 10-7 The Derived Demand for Labor: The Cheep-Cheep Birdhouse Company

Thus, one way to determine what quantity of a variable input to hire is to first determine the profit maximizing quantity to produce, and then to work backwards and determine what quantity of the input is needed.

The Marginal Revenue Product of Labor

Instead of calculating the amount of labor to hire based on the profit-maximizing quantity, an alternative approach is to compare the cost of hiring a worker to the value of what a worker produces.

The cost of hiring a worker is straightforward: it's the wage or compensation paid to the worker. To calculate the value of what a worker produces, find the quantity of production attributable to the marginal worker and multiply it by the price of that output. Exhibit 10-8 takes a different look at the basic cost information from Exhibit 10-7, focusing on the output of workers for workers at the Cheep Cheep Birdhouse Company. The first column repeats information from Exhibit 10-7, showing different numbers of workers and what they produce. The third column of the figure shows **marginal physical product,** which is the quantity of goods produced by the marginal worker. The fourth column shows the **marginal revenue product (MRP),** which reveals how much revenue a firm could receive from hiring an additional worker and selling the output of that worker. The MRP is calculated by multiplying the marginal physical product by the market price of $6. A second calculation of marginal revenue product, this one based on a market price for birdhouses of $10, appears in the final column. The two different marginal revenue products, one based on a price for birdhouses of $6 and one based on a price of $10, are also shown in Exhibit 10-8. The marginal revenue product of hiring additional workers gradually declines because of diminishing marginal returns.

marginal physical product: The quantity of goods produced by the addition of an additional input (like an additional worker).

marginal revenue product (MRP): Reveals how much revenue a firm could receive from hiring an additional worker and selling the output of that worker.

Exhibit 10-8 Marginal Revenue Product: The Cheep Cheep Example Continues

The wage is shown as a horizontal line, because it does not vary according to how many workers the firm hires. Marginal revenue product first slopes up, because of the increasing marginal returns at low levels of output in this example, and then slopes down because of diminishing marginal returns at higher levels of output. If a worker brings in more than enough marginal revenue to cover the wage of that worker, it will make sense to hire the worker. If the price of the birdhouse is $6, then the firm will hire five workers, because the marginal revenue product of the fifth worker is $102. If the price rises to $10 per birdhouse, then the value of what workers produce increases, and the firm will hire six workers. Finally, if the wage rises—in the diagram, imagine the horizontal line going up—the firm will hire fewer workers, because fewer workers will have a high enough marginal revenue product to justify paying the higher wage.

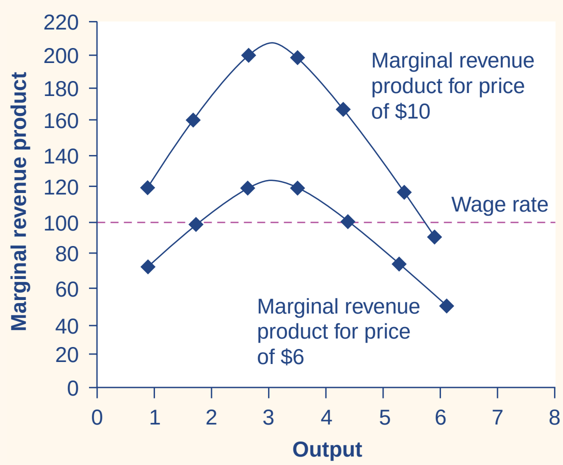

Labor	Output	Marginal Physical Product	Marginal Revenue Product (price $6)	Marginal Revenue Product (price $10)
0	0	—		—
1	12	12	$72	$120
2	28	16	$96	$160
3	48	20	$120	$200
4	68	20	$120	$200
5	85	17	$102	$170
6	96	11	$66	$110
7	105	9	$54	$90

When a profit-seeking, price-taking perfectly competitive firm is thinking about the quantity of labor to hire, it will compare the input price with the marginal revenue product for that input. For example, say that the cost for Cheep Cheep of hiring a birdhouse maker is $100. At a market price of $6, Cheep Cheep will notice that hiring a third worker and the fourth worker will have a marginal revenue product of $120, but only cost $100. So it will clearly make sense to hire these workers. In fact, it makes sense to hire the fifth worker as well, since the marginal revenue product of that worker still exceeds the wage. But at a market price for birdhouses of $6, it won't make sense to hire the sixth or seventh worker, because the marginal revenue product of hiring a worker will be less than $100.

A higher market price for the output of birdhouses will encourage more hiring. If the market price for birdhouses increased to $10, then Cheep Cheep will wish to hire a sixth worker, with a marginal revenue product of $110, but not a seventh worker, since the marginal revenue product of that final worker would only be $90. The optimal quantity of labor to hire will be where the price of the input is equal to the marginal revenue product of that input.

The concept of the marginal revenue product of labor hired explains why the demand curve for labor slopes down. When the wage is lower, more workers will have a marginal revenue product above that wage, so that firms will be willing to hire them. Conversely, when the wage is higher, fewer workers will have a marginal revenue product above that wage, so firms will demand a lower quantity of workers.

The derived demand approach and the marginal revenue product approach to figuring out the quantity of labor to hire are just two different ways of using the same basic information on production functions, input costs, and output prices. The two approaches should always obtain the same answer; in fact, the two methods can be used to check on each other.

Are Workers Paid as Much as They Deserve?

Many workers strongly suspect that they are not being paid as much as they deserve. Is this complaint nothing more than self-interested people who desire some additional income? Or is there reason to believe that in a competitive labor market, workers are systematically paid too little? The argument here suggests that workers are paid what they deserve in the sense that they are paid according their marginal revenue product—that is, according to the value of what the marginal worker produces.

In the real economy, calculating the marginal revenue product may be nearly impossible for many workers. Imagine a company making refrigerators where workers perform many jobs: some work on the assembly line, some maintain and repair equipment, some do paperwork for paying workers and suppliers, some plan future investment in physical capital, some do sales and marketing, and some work on the shipping dock. How could any accountant disentangle exactly what each of these very different worker contributes to making a refrigerator?

However, even when firms cannot calculate marginal revenue product directly, the pressures of competition in the labor market will push pay toward the marginal revenue product of labor. After all, if a foolish employer consistently paid its workers more than the marginal revenue product of what the workers produced, then that employer would make losses, and such pay could not be sustained. On the other hand, if an employer tried to pay workers less than what they produce as measured by their marginal revenue product, then other employers will see an opportunity to hire away some of the marginal workers by offering them a modest increase over their current pay. With these competitive pressures in the labor market, wages are pushed in the direction of marginal revenue product. Of course, just as a market price is not always exactly at equilibrium, but only tending toward equilibrium, it will also be true that wages are not always exactly at the marginal revenue product, although they tend toward that level. Studies of labor markets over time, and between countries, show that differences in how much workers are paid are strongly linked to the value of what workers produce.

If workers find it difficult or costly to switch between employers, and firms know that workers may be unable or unwilling to move to other jobs, then the employer will be able to pay those workers somewhat less than their marginal revenue product. For example, a worker in a rural area may have relatively few employment options that do not involve high monetary and emotional costs of moving. A recent immigrant may lack the information and the social network that would provide information and connections to other jobs. A society with widespread discrimination on the basis of race or gender may block people in certain groups from taking certain jobs. In any of these cases, employers may be able to pay certain workers less than the marginal revenue product of their labor.

Physical Capital Investment and the Hurdle Rate

In the long term, a perfectly competitive firm can adjust all of its inputs, including its investments in physical capital like buildings and machinery. Such investments in physical capital, by their nature, involve thinking about a rate of return over time; that is, the firm purchases the inputs of physical capital in the present and then uses them for production over a period of time in the future.

Consider a firm making plans for investments in physical capital. Managers from all over the company submit proposals. Each proposal includes either an estimate of the marginal revenue product of this capital investment—that is, how much additional revenue will this investment generate, or alternatively, by how much will this investment reduce the costs of the current level of output. Higher-level managers collect these proposals and check them over, estimating the rate of return that can be realistically expected from each proposal. The rates of return for a hypothetical list of 25 projects, ranked from top to bottom, are illustrated by the bar graph in Exhibit 10-9. In deciding which project to carry out, the profit-seeking firm will obviously prefer projects with high rates of return rather than the project with low rates of return (as long as the risk of the projects is similar). But to determine how many investment projects the firm should undertake, the managers of the firm must compare the rates of return to the cost of financial capital.

Financial capital always comes at a cost. Sometimes the firm must pay this cost directly; for example, the firm may borrow financial capital from a bank, and thus need to pay an interest rate. However, even if the firm can use funds that it has earned from past profits, spending this money on an investment in physical capital still has an opportunity cost. After all, the firm's financial capital could have been invested in some other way that would have paid a rate of return. Thus, the cost of financial capital to the firm can be conceived of either as a rate of return or interest rate paid to others, or as an opportunity cost of a rate of return that was not received. The specific ways in which firms can raise money for investment—including reinvesting the profits that they earn or using bank loans, bonds, and stocks—will be discussed in Chapter 19.

Firms must also remember that any physical capital investment has some degree of risk. The business world is full of horror stories of factories that were constructed months or years late and cost much more than expected. In other cases, the costs may

Exhibit 10-9 Considering Investment Projects

When a firm makes plans for investments in physical capital, it compares the projected rates of return on these investments with the cost of financial capital to the firm. Often, this comparison will use a hurdle rate as the rate of return that must be exceeded to make the project worth undertaking. In this example, projects A to S do not exceed the hurdle rate. A lower cost of capital will result in more investment projects being undertaken (imagine the horizontal hurdle rate line moving down, so that more projects cross the hurdle). Conversely, a higher cost of financial capital will result in fewer physical capital investment projects being undertaken.

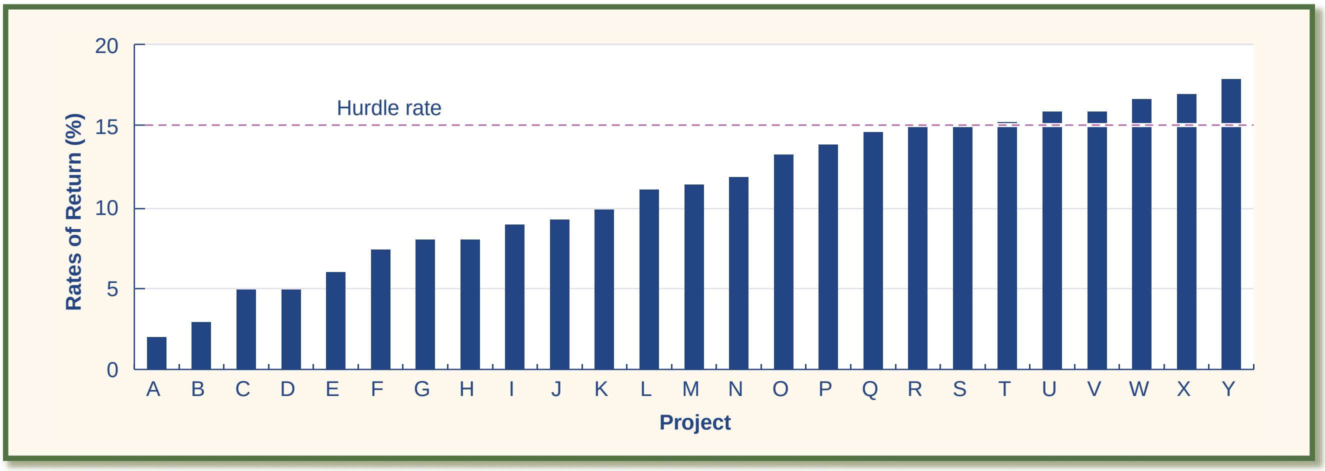

be estimated accurately, but the returns turn out to be lower than expected. When a company drills for oil, or decides to make a major investment in upgrading its Internet capabilities, it can't be certain of the results in advance.

Firms often take risk into account by using a "hurdle rate," which is an interest rate that is somewhat higher than the actual cost of financial capital to the firm. The intuitive meaning of the hurdle rate is that the rate of return that proposed investment projects must cross before being approved. Estimates of typical hurdle rates, based on studies of the investment decisions that firms make, are often about 5–8 percentage points higher than the actual cost of financial capital, whether it is measured as the borrowing interest rate or the opportunity cost of capital. This higher hurdle interest rate can be thought of as a risk premium; that is, it's an adjustment for the risks involved and it gives the firm a margin for errors of bad luck and overly optimistic forecasts.

If a firm chooses a 15% hurdle rate for its physical capital investment, then such projects must pay an expected return of at least 15% to be worthwhile. A 15% rate of return is represented by the dashed line across Exhibit 10-9. If the cost of financial capital for the firm was higher, then fewer investment projects will make economic sense. Conversely, if the cost of financial capital declines, then a larger number of investment projects will make sense. This logic explains why the demand for financial capital is a downward-sloping curve; the downward slope means that a higher interest rate leads to a lower demand for financial capital, because fewer physical capital investments appear desirable.

Physical Capital Investment and Long-Run Average Cost

In a perfectly competitive market, firms seek out the combination of variable inputs like labor and fixed inputs like physical capital investment inputs that will allow them to produce at the minimum of the long-run average cost curve. If a firm does not produce at the lowest possible average cost, then firms with a lower cost of production will be able to sell at a lower price. Thus, the pressure of competition between firms will shape both the specific kinds of machinery and equipment and the overall quantity of investment in physical capital.

Efficiency in Perfectly Competitive Markets

When profit-maximizing firms in perfectly competitive markets combine with utility-maximizing consumers, something remarkable happens: the resulting quantities of outputs of goods and services demonstrate both productive and allocative efficiency (terms that were first introduced back in Chapter 2).

Productive efficiency means producing without waste, so that the choice is on the frontier of the production possibility frontier. In the long run in a perfectly competitive market, because of the process of entry and exit, the price in the market is equal to the minimum of the long-run average cost curve. In other words, goods are being produced and sold at the lowest possible average cost. When the average cost of production is as low as possible, the economy will be operating on the production possibility frontier.

Allocative efficiency means that among the points on the production possibility frontier, the point that is chosen is socially preferred in a particular and specific sense. In a perfectly competitive market, price will be equal to the marginal cost of production. Think about the price that is paid for a good as a measure of the social benefit received for that good; after all, the willingness to pay conveys what the good was worth to a buyer. Then think about the marginal cost of producing the good as representing not just the cost for the firm, but more broadly as the social cost of producing that good. When perfectly competitive firms follow the rule that profits are maximized by producing at the quantity where price is equal to marginal cost, they are thus ensuring that the social benefits received from producing a good are in line with the social costs of production.

It's useful to explore what is meant by allocative efficiency by walking through an example of what allocative inefficiency would imply. Begin by assuming that the market for ordinary television sets is perfectly competitive, and so $P = MC$. Now consider what it would mean if firms in that market produced a lesser quantity of television sets. At a lesser quantity, marginal costs will not yet have increased as much, so that price will exceed marginal cost; that is, $P > MC$. In that situation, the benefit to society as a whole of producing additional goods, as measured by the willingness of consumers to pay for marginal units of a good, would be higher than the cost of the inputs of labor and physical capital needed to produce the marginal good. In other words, the gains to society as a whole from producing additional marginal units will be greater than the costs. Conversely, consider what it would mean if, compared to the level of output at the allocatively efficient choice when $P = MC$, firms produced a greater quantity of television sets. At a greater quantity, marginal costs of production will have increased so that $P < MC$. In that case, the marginal costs of producing additional television sets are greater than the benefit to society as measured by what people are willing to pay. For society as a whole, since the costs are outstripping the benefits, it will make sense to produce a lower quantity of such goods.

When perfectly competitive firms maximize their profits by producing the quantity where $P = MC$, they also assure that the benefits to consumers of what they are buying, as measured by the price they are willing to pay, equal to the costs to society of producing the marginal units, as measured by the marginal costs the firm must pay—and thus that allocative efficiency holds.

The statements that a perfectly competitive market in the long run will feature both productive and allocative efficiency do need to be taken with a few grains of salt. Remember, economists are using the concept of "efficiency" in a particular and specific sense, not as a synonym for "desirable in every way." Perfect competition in the long run is a hypothetical extreme. In other competitive situations discussed in the next two chapters, like monopoly, monopolistic competition, and oligopoly, firms will not always produce at the minimum of average cost, nor will they always set price equal to marginal cost. Thus, these other competitive situations will not produce productive and allocative efficiency. Moreover, real-world markets include many issues that are assumed away in the model of perfect competition, including pollution, inventions of new technology, poverty which may make some people unable to pay for basic necessities of life, government programs like national defense or education, discrimination in labor markets, and buyers and sellers who must deal with imperfect and unclear information. These issues are explored in later chapters of this book. However, the theoretical efficiency of perfect competition does provide a useful benchmark for comparing what issues arise from these real-world problems.

Conclusion

Extreme assumptions often provide a useful starting point for analysis. In a physics class, it may be useful to analyze gravity by assuming that a ball which is dropped from a tower experiences zero friction as it falls through the air. This assumption of zero friction is not literally true. But if the amount of friction from the air is relatively small compared to the pull of gravity, then assuming zero friction may not make much difference to the results. If friction is significant, then calculations for the effects of friction can be added, once the basic no-friction example is understood.

The assumption of perfect competition that firms are perfect price-takers, unable to raise their price by so much as a penny without losing all of their sales, is not literally true. But in a market with many competitors who have highly similar products, the assumption of perfect competition may provide a reasonable start for the analysis. In a highly competitive market—even if it is not strictly speaking perfectly competitive—firms will not lose all of their sales if they raise prices by a penny, but they might lose a large enough proportion of their sales to cripple the firm if they raise prices by 5%. In a highly competitive real world market, not all firms will be producing at the absolute bottom of the average cost curve at all times—but the firms will still face considerable pressure to keep average costs as low as possible. Firms in a highly competitive market may earn profits in the short run, but the entry of new firms or the expansion of existing firms means that such profits will not persist in the long run. Truly perfect competition is a hypothetical extreme, but many real-world producers face highly competitive markets every day.

Key Concepts and Summary

1. A **perfectly competitive** firm is a price taker, which means that it must accept the prices at which its sell goods and the prices at which it purchase inputs as determined in the market. If a perfectly competitive firm attempts to charge even a tiny amount more than the market price, it will be unable to make any sales.

2. As a perfectly competitive firm produces a greater quantity of output, its **total revenue** steadily increases at a constant rate determined by the given market price.

3. Profits will be highest (or losses will be smallest) at the quantity of output where total revenues exceed total costs by the greatest amount (or where total revenues fall short of total costs by the smallest amount). Alternatively, profits will be highest where **marginal revenue,** which is price for a perfectly competitive firm, is equal to marginal cost.

4. If the market price faced by a perfectly competitive firm is above average cost at the profit-maximizing quantity of output, then the firm is making profits. If the market price is below average cost at the profit-maximizing quantity of output, then the firm is making losses. If the market price is at average cost, at the profit-maximizing level of output, then the firm is making zero profits. The point where the marginal cost curve crosses the average cost curve, at the minimum of the average cost curve, is called the "zero profit point."

5. If the market price faced by a perfectly competitive firm is below average variable cost at the profit-maximizing quantity of output, then the firm should shut down operations immediately. If the market price faced by a perfectly competitive firm is above average variable cost, but below average cost, then the firm should continue producing in the short run, but exit in the long run. The point where the marginal cost curve crosses the average variable cost curve is called the **shutdown point.**

6. In the long run, firms will respond to profits through a process of **entry,** where existing firms expand output and new firms enter the market. Conversely, firms will react to losses in the long run through a process of **exit,** in which existing firms reduce output or cease production altogether.

7. Through the process of entry in response to profits and exit in response to losses, the price level in a perfectly competitive market will move toward the zero-profit point where the marginal cost curve crosses the AC curve, at the minimum of the average cost curve.

8. **Accounting profit** is measured by taking total revenues and subtracting expenditures. Economic profit is measured after taking total revenue, subtracting all expenditures, and also subtracting the opportunity cost of financial capital. Thus, zero **economic profit** actually means a normal accounting rate of profit.

9. A profit-maximizing firm will determine its optimal quantity of output, and then hire the amount of the variable input, like labor, needed to produce that quantity of output.

10. The **marginal revenue product,** the value of what is produced by an additional worker, is calculated by taking the quantity of output produced by the marginal worker and multiplying it by the price of that output. A profit-maximizing firm will hire workers as long as the marginal revenue product of what they produce exceeds the wage, up to the point where the marginal revenue product is equal to the wage.

11. If an employer paid wages that were higher than marginal revenue product—that is, higher than the value of what workers produced—the employer will tend to make

losses and such wages cannot be continued. If one employer paid wages that were lower than the marginal revenue product, then other employers will see a profit opportunity in hiring those workers and paying them more. As long as workers can move freely between comparable alternative jobs, the pressures of the labor market will tend to push wages toward marginal revenue product.

12. Firms will invest as long as the expected returns of the investment exceed the cost of capital, where the cost of capital typically includes a "hurdle rate" to take account of the risk involved.

13. When perfectly competitive firms produce at the minimum of the long-run average cost curve, the market will show productive efficiency, since they are producing at the lowest possible average cost. When perfectly competitive firms produce at the quantity where P = MC, the market will illustrate allocative efficiency, since each good is being produced up to the quantity where the amount that the good benefits society, as measured by the price people are willing to pay, is equal to the cost to society, as measured by the marginal cost of production.

Review Questions

1. How does a perfectly competitive firm decide what price to charge?
2. What prevents a perfectly competitive firm from seeking higher profits by increasing the price that it charges?
3. How does a perfectly competitive firm calculate total revenue?
4. Briefly explain the reason for the shape of a marginal revenue curve for a perfectly competitive firm.
5. What rule does a perfectly competitive firm apply to determine its profit-maximizing quantity of output?
6. How does the average cost curve help to show whether a firm is making profits or losses?
7. What two lines on a cost curve diagram intersect at the zero-profit point?
8. Should a firm shut down immediately if it is making losses?
9. How does the average variable cost curve help a firm know whether it should shut down immediately?
10. What two lines on a cost curve diagram intersect at the shutdown point?
11. Why does entry occur?
12. Why does exit occur?
13. Do entry and exit occur in the short run, the long run, both, or neither?
14. What price will a perfectly competitive firm end up charging in the long run? Why?
15. What's the difference between accounting profit and economic profit? Which one is higher?
16. Why is labor called a "derived demand"?
17. How can a perfectly competitive firm use derived demand to decide how much labor to hire?
18. How is the marginal revenue product of labor calculated?
19. How can a perfectly competitive firm use the marginal revenue product of labor to determine what quantity of labor to hire?
20. What is the "hurdle rate"?
21. How does a firm use the hurdle rate to decide which investments in physical capital to undertake?
22. Will a perfectly competitive industry display productive efficiency? Why or why not?
23. Will a perfectly competitive industry display allocative efficiency? Why or why not?