Markets for Factor Inputs

14

So far we have concentrated on *output markets*: markets for goods and services that firms sell and consumers purchase. In this chapter, we discuss *factor markets*: markets for labor, raw materials, and other inputs to production. Much of our material will be familiar because the same forces that shape supply and demand in output markets also affect factor markets.

We have seen that some output markets are perfectly or almost perfectly competitive, while producers in others have market power. The same is true for factor markets. We will examine three different factor market structures:

1. Perfectly competitive factor markets;
2. Markets in which buyers of factors have monopsony power;
3. Markets in which sellers of factors have monopoly power.

We will also point out instances in which equilibrium in the factor market depends on the extent of market power in *output* markets.

CHAPTER OUTLINE

14.1 Competitive Factor Markets 521
14.2 Equilibrium in a Competitive Factor Market 534
14.3 Factor Markets with Monopsony Power 539
14.4 Factor Markets with Monopoly Power 543

LIST OF EXAMPLES

14.1 The Demand for Jet Fuel 528
14.2 Labor Supply for One- and Two-Earner Households 533
14.3 Pay in the Military 537
14.4 Monopsony Power in the Market for Baseball Players 541
14.5 Teenage Labor Markets and the Minimum Wage 542
14.6 The Decline of Private-Sector Unionism 546
14.7 Wage Inequality—Have Computers Changed the Labor Market? 547

14.1 COMPETITIVE FACTOR MARKETS

A competitive *factor market* is one in which there are a large number of sellers and buyers of a factor of production, such as labor or raw materials. Because no single seller or buyer can affect the price of a given factor, each is a price taker. For example, if individual firms that buy lumber to construct homes purchase a small share of the total volume of lumber available, their purchasing decision will have no effect on price. Likewise, if each supplier of lumber controls only a small share of the market, no individual supplier's decision will affect the price of the lumber that he sells. Instead, the price of lumber (and the total quantity produced) will be determined by the aggregate supply and demand for lumber.

We begin by analyzing the demands for a factor by individual firms. These demands are added to get market demand. We then shift to the supply side of the market and show how market price and input levels are determined.

Demand for a Factor Input When Only One Input Is Variable

Like demand curves for the final goods that result from the production process, demand curves for factors of production are downward sloping. Unlike consumers' demands for goods and services, however, factor demands are **derived demands**: they depend on, and are derived from, the firm's level of output and the costs of inputs. For example, the demand of the Microsoft Corporation for computer programmers is a derived demand that depends not only on the current salaries of programmers, but also on how much software Microsoft expects to sell.

• **derived demand** Demand for an input that depends on, and is derived from, both the firm's level of output and the cost of inputs.

To analyze factor demands, we will use the material from Chapter 7 that shows how a firm chooses its production inputs. We will assume that the firm produces its output using two inputs, capital K and labor L, that can be hired at the prices r (the rental cost of capital) and w (the wage rate), respectively.[1] We will also assume that the firm has its plant and equipment in place (as in a short-run analysis) and must only decide how much labor to hire.

Suppose that the firm has hired a certain number of workers and wants to know whether it is profitable to hire one additional worker. This will be profitable if the additional revenue from the output of the worker's labor is greater than its cost. The additional revenue from an incremental unit of labor, the **marginal revenue product of labor**, is denoted MRP_L. The cost of an incremental unit of labor is the wage rate, w. Thus, it is profitable to hire more labor if the MRP_L is at least as large as the wage rate w.

• **marginal revenue product** Additional revenue resulting from the sale of output created by the use of one additional unit of an input.

How do we measure the MRP_L? It's *the additional output obtained from the additional unit of this labor, multiplied by the additional revenue from an extra unit of output*. The additional output is given by the marginal product of labor MP_L and the additional revenue by the marginal revenue MR.

Recall that in §8.3, marginal revenue is defined to be the increase in revenue resulting from a one-unit increase in output.

Formally, the marginal revenue product is $\Delta R/\Delta L$, where L is the number of units of labor input and R is revenue. The additional output per unit of labor, the MP_L, is given by $\Delta Q/\Delta L$, and marginal revenue, MR, is equal to $\Delta R/\Delta Q$. Because $\Delta R/\Delta L = (\Delta R)/(\Delta Q)(\Delta Q/\Delta L)$, it follows that

$$MRP_L = (MR)(MP_L) \tag{14.1}$$

This important result holds for any competitive factor market, whether or not the output market is competitive. However, to examine the characteristics of the MRP_L, let's begin with the case of a perfectly competitive output (and input) market. In a competitive output market, a firm will sell all its output at the market price P. The marginal revenue from the sale of an additional unit of output is then equal to P. In this case, the marginal revenue product of labor is equal to the marginal product of labor times the price of the product:

In §8.2, we explain that because the demand facing each firm in a competitive market is perfectly elastic, each firm will sell its output at a price equal to its average revenue and to its marginal revenue.

$$MRP_L = (MP_L)(P) \tag{14.2}$$

The higher of the two curves in Figure 14.1 represents the MRP_L curve for a firm in a competitive output market. Note that because there are diminishing marginal returns to labor, the marginal product of labor falls as the amount of labor increases. The marginal revenue product curve thus slopes downward, even though the price of the output is constant.

In §6.2, we explain the law of diminishing marginal returns—as the use of an input increases with other inputs fixed, the resulting additions to output will eventually decrease.

[1] We implicitly assume that all inputs to production are identical in quality. Differences in workers' skills and abilities are discussed in Chapter 17.

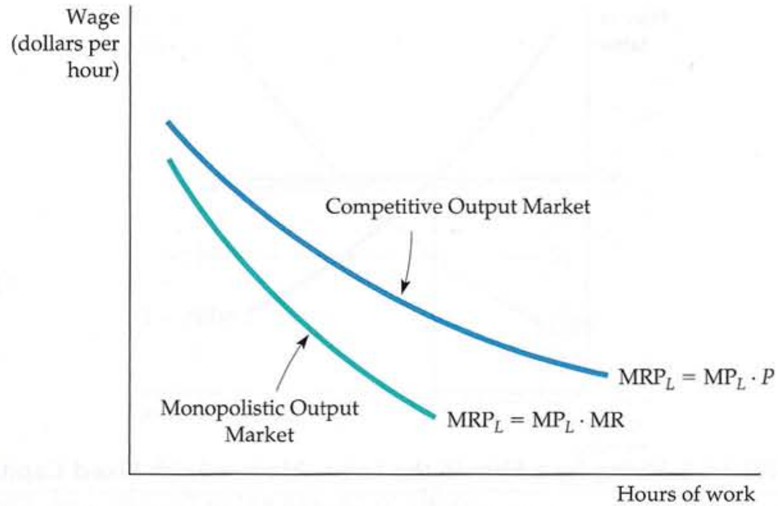

FIGURE 14.1 Marginal Revenue Product

In a competitive factor market in which the producer is a price taker, the buyer's demand for an input is given by the marginal revenue product curve. The MRP curve falls because the marginal product of labor falls as hours of work increase. When the producer of the product has monopoly power, the demand for the input is also given by the MRP curve. In this case, however, the MRP curve falls because both the marginal product of labor and marginal revenue fall.

The lower curve in Figure 14.1 is the MRP_L curve when the firm has monopoly power in the output market. When firms have monopoly power, they face a downward-sloping demand curve and must therefore lower the price of all units of the product in order to sell more of it. As a result, marginal revenue is always less than price ($MR < P$). This explains why the monopolistic curve lies below the competitive curve and why marginal revenue falls as output increases. Thus the marginal revenue product curve slopes downward in this case because the marginal revenue curve *and* the marginal product curve slope downward.

Note that the marginal revenue product tells us how much the firm should be willing to pay to hire an additional unit of labor. As long as the MRP_L is greater than the wage rate, the firm should hire more labor. If the marginal revenue product is less than the wage rate, the firm should lay off workers. Only when the marginal revenue product is equal to the wage rate will the firm have hired the profit-maximizing amount of labor. The profit-maximizing condition is therefore

$$MRP_L = w \tag{14.3}$$

Figure 14.2 illustrates this condition. The demand for labor curve D_L is the MRP_L. Note that the quantity of labor demanded increases as the wage rate falls. Because the labor market is perfectly competitive, the firm can hire as many workers as it wants at the market wage w^* and is not able to affect the market wage. The supply of labor curve facing the firm S_L is thus a horizontal line. The profit-maximizing amount of labor that the firm hires, L^*, is at the intersection of the supply and demand curves.

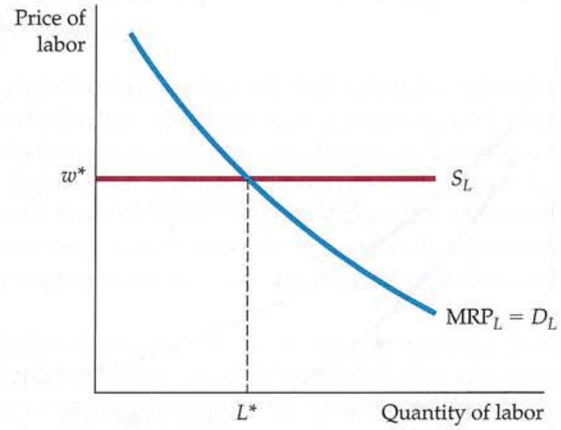

FIGURE 14.2 Hiring by a Firm in the Labor Market (with Fixed Capital)

In a competitive labor market, a firm faces a perfectly elastic supply of labor S_L and can hire as many workers as it wants at a wage rate w^*. The firm's demand for labor D_L is given by its marginal revenue product of labor MRP_L. The profit-maximizing firm will hire L^* units of labor at the point where the marginal revenue product of labor is equal to the wage rate.

Figure 14.3 shows how the quantity of labor demanded changes in response to a drop in the market wage rate from w_1 to w_2. The wage rate might decrease if more people entering the labor force are looking for jobs for the first time (as happened, for example, when the baby boomers came of age). The quantity of labor demanded by the firm is initially L_1, at the intersection of MRP_L and S_1. However, when the supply of labor curve shifts from S_1 to S_2, the wage falls from w_1 to w_2 and the quantity of labor demanded increases from L_1 to L_2.

In §8.3, we explain that a firm maximizes its profit by choosing an output at which marginal revenue equals marginal cost.

Factor markets are similar to output markets in many ways. For example, the factor market profit-maximizing condition that the marginal revenue product of labor be equal to the wage rate is analogous to the output market condition that marginal revenue be equal to marginal cost. To see why this is true, recall that $\text{MRP}_L = (\text{MP}_L)(\text{MR})$ and divide both sides of equation (14.3) by the marginal product of labor. Then,

$$\text{MR} = w/\text{MP}_L \tag{14.4}$$

Because MP_L measures additional output per unit of input, the right-hand side of equation (14.4) measures the marginal cost of an additional unit of output (the wage rate multiplied by the labor needed to produce one unit of output). Equation (14.4) shows that *both the hiring and output choices of the firm follow the same rule: Inputs or outputs are chosen so that marginal revenue (from the sale of output) is equal to marginal cost (from the purchase of inputs).* This principle holds in both competitive and noncompetitive markets.

Demand for a Factor Input When Several Inputs Are Variable

When the firm simultaneously chooses quantities of two or more variable inputs, the hiring problem becomes more difficult because a change in the price of one input will change the demand for others. Suppose, for example, that both

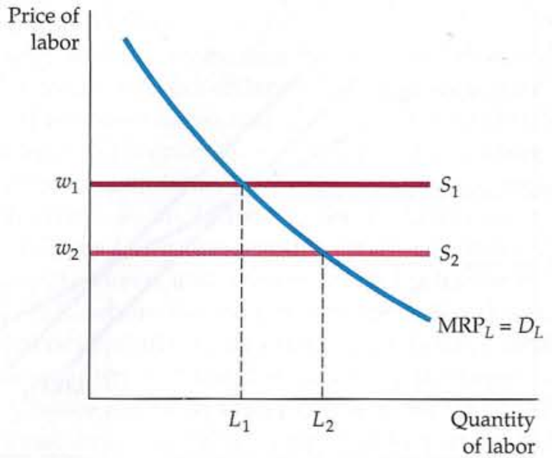

FIGURE 14.3 A Shift in the Supply of Labor

When the supply of labor facing the firms is S_1, the firm hires L_1 units of labor at wage w_1. But when the market wage rate decreases and the supply of labor shifts to S_2, the firm maximizes its profit by moving along the demand for labor curve until the new wage rate w_2 is equal to the marginal revenue product of labor. As a result, L_2 units of labor are hired.

labor and assembly-line machinery are variable inputs for producing farm equipment. Let's say that we wish to determine the firm's demand for labor curve. As the wage rate falls, more labor will be demanded even if the firm's investment in machinery is unchanged. But as labor becomes less expensive, the marginal cost of producing the farm equipment falls. Consequently, it is profitable for the firm to increase its output. In that case, the firm is likely to invest in additional machinery to expand production capacity. Expanding the use of machinery causes the marginal revenue product of labor curve to shift to the right; in turn, the quantity of labor demanded increases.

Figure 14.4 illustrates this. Suppose that when the wage rate is \$20 per hour, the firm hires 100 worker-hours, as shown by point *A* on the MRP_{L1} curve. Now consider what happens when the wage rate falls to \$15 per hour. Because the marginal revenue product of labor is now greater than the wage rate, the firm will demand more labor. But the MRP_{L1} curve describes the demand for labor when the use of machinery is fixed. In fact, a greater amount of labor causes the marginal product of *capital* to rise, which encourages the firm to rent more machinery as well as hire more labor. Because there is more machinery, the marginal product of labor will increase. (With more machinery, workers can be more productive.) The marginal revenue product curve will therefore shift to the right (to MRP_{L2}). Thus, when the wage rate falls, the firm will use 140 hours of labor. This is shown by a new point on the demand curve, *C*, rather than 120 hours as given by *B*. *A* and *C* are both on the firm's demand for labor curve (with machinery variable) D_L; *B* is not.

Note that as constructed, the demand for labor curve is more elastic than either of the two marginal product of labor curves (which presume no change in the amount of machinery). Thus, when capital inputs are variable in the long run, there is a greater elasticity of demand because firms can substitute capital for labor in the production process.

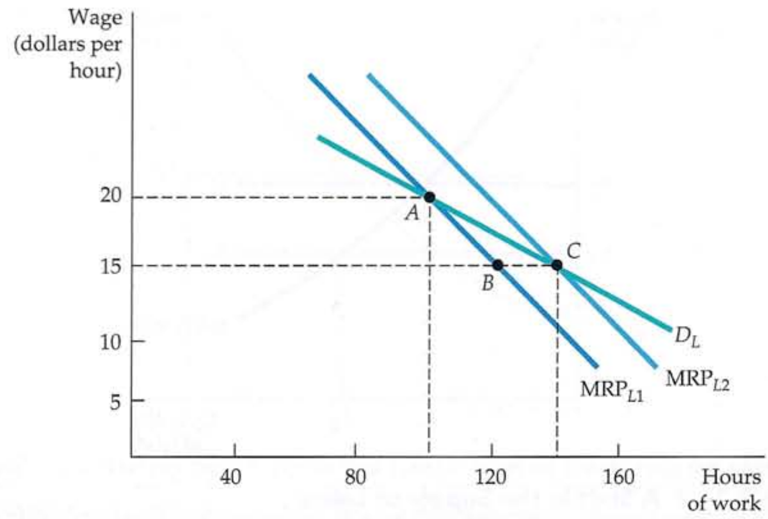

FIGURE 14.4 Firm's Demand Curve for Labor (with Variable Capital)

When two or more inputs are variable, a firm's demand for one input depends on the marginal revenue product of both inputs. When the wage rate is \$20, *A* represents one point on the firm's demand for labor curve. When the wage rate falls to \$15, the marginal product of capital rises, encouraging the firm to rent more machinery and hire more labor. As a result, the MRP curve shifts from MRP_{L1} to MRP_{L2}, generating a new point *C* on the firm's demand for labor curve. Thus *A* and *C* are on the demand for labor curve, but *B* is not.

The Market Demand Curve

Recall from §4.3 that the market demand curve for a product shows how much of the product consumers are willing to buy as the price of the product changes.

When we aggregated the individual demand curves of consumers to obtain the market demand curve for a product, we were concerned with a single industry. However, a factor input such as skilled labor is demanded by firms in many different industries. Moreover, as we move from industry to industry, we are likely to find that firms' demands for labor (which are derived in part from the demands for the firms' output) vary substantially. Therefore, to obtain the total market demand for labor curve, we must first determine each industry's demand for labor, and then add the industry demand curves horizontally. The second step is straightforward. Adding industry demand curves for labor to obtain a market demand curve for labor is just like adding individual product demand curves to obtain the market demand curve for that product. So let's concentrate our attention on the more difficult first step.

Determining Industry Demand The first step—determining industry demand—takes into account the fact that both the level of output produced by the firm and its product price change as the prices of the inputs to production change. It is easiest to determine market demand when there is a single producer. In that case, the marginal revenue product curve is the industry demand curve for the input. When there are many firms, however, the analysis is more complex because of the possible interaction among the firms. Consider, for instance, the demand for labor when output markets are perfectly competitive. Then, the marginal revenue product of labor is the product of the price of the good and the marginal product of labor (see equation 14.2), as shown by the curve MRP_{L1} in Figure 14.5(a).

Suppose initially that the wage rate for labor is \$15 per hour and that the firm demands 100 worker-hours of labor. Now the wage rate for this firm falls to \$10 per hour. If no other firms could hire workers at the lower wage, then our firm would hire 150 worker-hours of labor (by finding the point on the MRP_{L1} curve that corresponds to the \$10-per-hour wage rate). But if the wage rate falls for all firms in an industry, the industry as a whole will hire more labor. This will lead to more output from the industry, a shift to the right of the industry supply curve, and a lower market price for its product.

In Figure 14.5(a), when the product price falls, the original marginal revenue product curve shifts downward, from MRP_{L1} to MRP_{L2}. This shift results in a lower quantity of labor demanded by the firm—120 worker-hours rather than 150. Consequently, industry demand for labor will be lower than it would be if only one firm were able to hire workers at the lower wage. Figure 14.5(b) illustrates this. The lighter line shows the horizontal sum of the individual firms' demands for labor that would result if product price did not change as the wage falls. The darker line shows the industry demand curve for labor, which takes into account the fact that product price will fall as all firms expand their output in response to the lower wage rate. When the wage rate is \$15 per hour, industry demand for labor is L_0 worker-hours. When it falls to \$10 per hour, industry demand increases to L_1. Note that this is a smaller increase than L_2, which would occur if the product price were fixed. The aggregation of industry demand curves into the market demand curve for labor is the final step: To complete it, we simply add the labor demanded in all industries.

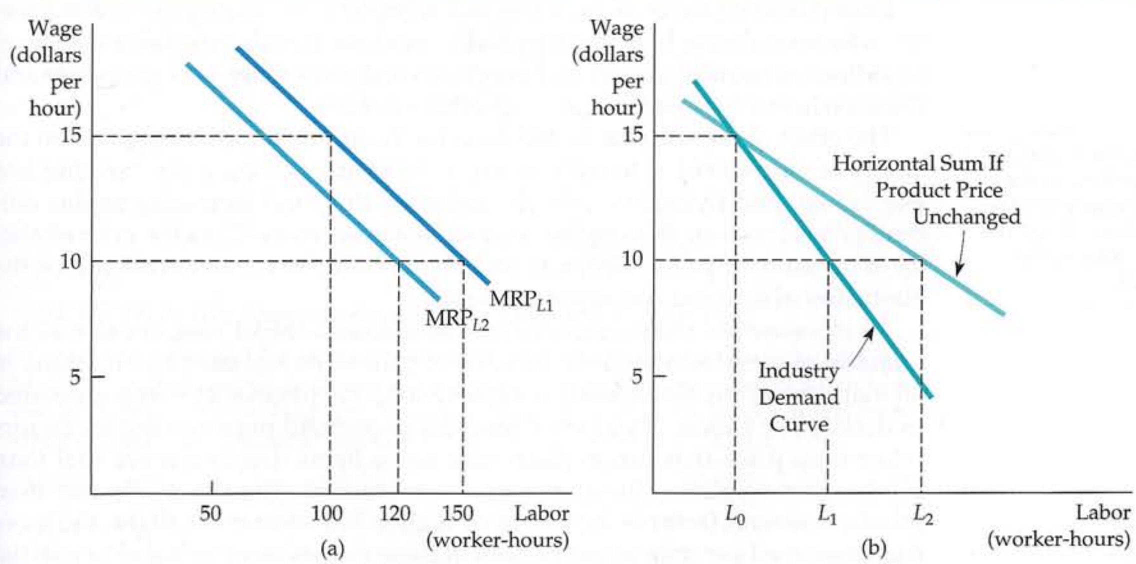

FIGURE 14.5 The Industry Demand for Labor

The demand curve for labor of a competitive firm, MRP_{L1} in (a), takes the product price as given. But as the wage rate falls from \$15 to \$10 per hour, the product price also falls. Thus the firm's demand curve shifts downward to MRP_{L2}. As a result, the industry demand curve, shown in (b), is more inelastic than the demand curve that would be obtained if the product price were assumed to be unchanged.

The derivation of the market demand curve for labor (or for any other input) is essentially the same when the output market is noncompetitive. The only difference is that it is more difficult to predict the change in product price in response to a change in the wage rate because each firm in the market is likely to be pricing strategically rather than taking price as given.

EXAMPLE 14.1 The Demand for Jet Fuel

As discussed in Example 9.3 on the airline industry (page 321), there have been several periods during the past few decades when fuel costs for U.S. airlines increased rapidly, in tandem with rising world oil prices. For example, whereas fuel costs made up 12.4 percent of total operating costs in 1971, that number rose to about 30 percent in 1980. As we would expect, the amount of jet fuel used by airlines during this period fell as its price rose. Thus the output of the airline industry, as measured by the number of ton-miles, rose by 29.6 percent, while the amount of jet fuel consumed increased by only 8.8 percent. (One ton-mile is short for one ton of passengers, baggage, or freight transported one mile.) Fuel prices fell substantially during the mid-1980s and, relative to 1980 levels, remained low (in real terms) until about 2005, at which point they again increased dramatically. Overall, the cost of jet fuel, as a share of operating costs, remains the second-highest expense for airlines (after labor), averaging about 20 percent of total operating costs while fluctuating between 10 and 30 percent.

Understanding the demand for jet fuel is important to managers of oil refineries, who must decide how much jet fuel to produce. It is also crucial to managers of airlines, who must project fuel purchases and costs when fuel prices rise and decide whether to invest in more fuel-efficient planes.[2]

In §2.4, we define the price elasticity of demand as the percentage change in quantity demanded resulting from a 1-percent change in the price of a good.

The effect of the increase in fuel costs on the airline industry depends on the ability of airlines either to cut fuel usage by reducing weight (by carrying less excess fuel) and flying more slowly (reducing drag and increasing engine efficiency) or to pass on their higher costs in customer prices. Thus the price elasticity of demand for jet fuel depends both on the ability to conserve fuel and on the elasticities of demand and supply of travel.

To measure the short-run elasticity of demand for jet fuel, we use as the quantity of fuel demanded the number of gallons of fuel used by an airline in all markets within its domestic route network. The price of jet fuel is measured in dollars per gallon. A statistical analysis of demand must control for factors other than price that can explain why some firms demand more fuel than others. Some airlines, for example, use more fuel-efficient jet aircraft than others. A second factor is the length of flights: The shorter the flight, the more fuel consumed per mile of travel. Both of these factors were included in a statistical analysis that relates the quantity of fuel demanded to its price. Table 14.1 shows some short-run price elasticities. (They do not account for the introduction of new types of aircraft.)

[2]This example is drawn in part from Joseph M. Cigliano, "The Demand for Jet Fuel by the U.S. Domestic Trunk Airlines," *Business Economics* (September 1982): 32–36.

TABLE 14.1 Short-Run Price Elasticity of Demand for Jet Fuel

Airline	Elasticity	Airline	Elasticity
American	–.06	Delta	–.15
Continental	–.09	United	–.10
Northwest	–.07		

The jet fuel price elasticities for the airlines range in value from –.06 (for American) to –.15 (for Delta). Overall, the results show that the demand for jet fuel as an input to the production of airline flight-miles is very inelastic. This finding is not surprising: In the short run, there is no good substitute for jet fuel. The long-run elasticity of demand is higher, however, because airlines can eventually introduce more energy-efficient airplanes.

Figure 14.6 shows the short- and long-run demands for jet fuel. The short-run demand curve, MRP_{SR}, is much less elastic than the long-run demand curve because it takes time to substitute newer, more fuel-efficient airplanes when the price of fuel goes up.

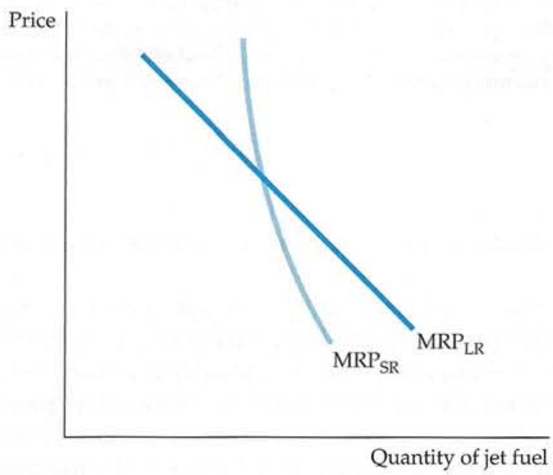

FIGURE 14.6 The Short- and Long-Run Demand for Jet Fuel

The short-run demand for jet fuel MRP_{SR} is more inelastic than the long-run demand MRP_{LR}. In the short run, airlines cannot reduce fuel consumption much when fuel prices increase. In the long run, however, they can switch to longer, more fuel-efficient routes and put more fuel-efficient planes into service.

The Supply of Inputs to a Firm

When the market for a factor input is perfectly competitive, a firm can purchase as much of that input as it wants at a fixed market price, which is determined by the intersection of the market demand and supply curves, as shown in Figure 14.7(a). The input supply curve facing a firm is then perfectly elastic. Thus, in Figure 14.7(b), a firm is buying fabric at $10 per yard to sew into clothing. Because the

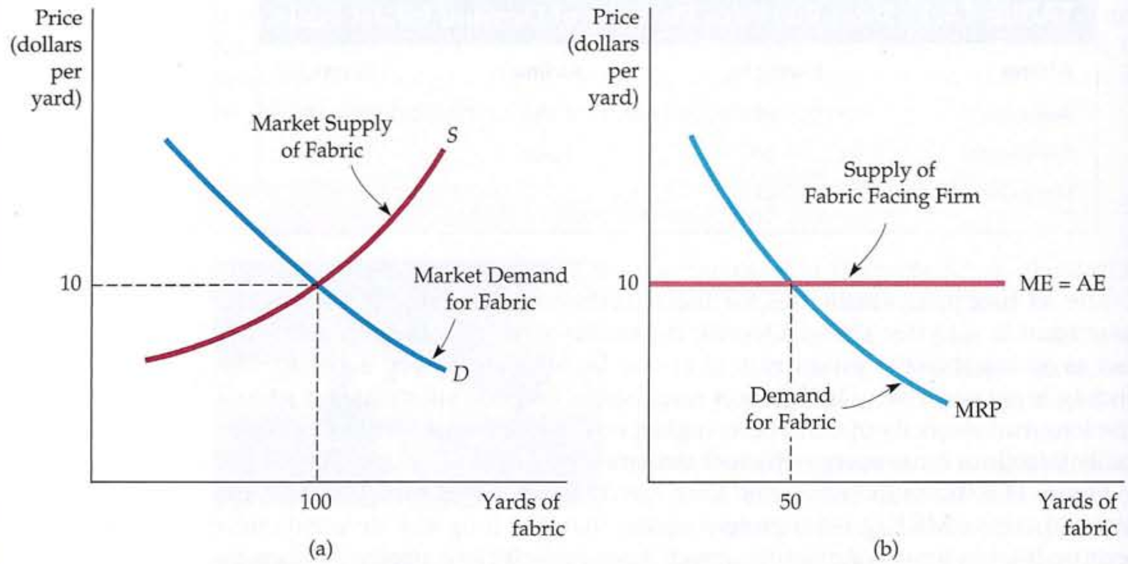

FIGURE 14.7 A Firm's Input Supply in a Competitive Factor Market

In a competitive factor market, a firm can buy any amount of the input it wants without affecting the price. Therefore, the firm faces a perfectly elastic supply curve for that input. As a result, the quantity of the input purchased by the producer of the product is determined by the intersection of the input demand and supply curves. In (a), the industry quantity demanded and quantity supplied of fabric are equated at a price of \$10 per yard. In (b), the firm faces a horizontal marginal expenditure curve at a price of \$10 per yard of fabric and chooses to buy 50 yards.

• **average expenditure curve** Supply curve representing the price per unit that a firm pays for a good.

• **marginal expenditure curve** Curve describing the additional cost of purchasing one additional unit of a good.

firm is only a small part of the fabric market, it can buy all it wants without affecting the price.

In Section 10.5 we explained that the supply curve AE facing the firm in Figure 14.7(b) is its **average expenditure curve** (just as the demand curve facing a firm is its *average revenue curve*), because it represents the price per unit that the firm pays for the good. On the other hand, the **marginal expenditure curve** represents the firm's expenditure on an *additional unit* that it buys. (The marginal expenditure curve in a factor market is analogous to the marginal revenue curve in the output market.) The marginal expenditure depends on whether you are a competitive buyer or a buyer with monopsony power. If you are a competitive buyer, the cost of each unit is the same no matter how many units you purchase; it is the market price of the good. The price paid is the average expenditure per unit, and the marginal expenditure is equal to the average. Consequently, when the factor market is competitive, the average expenditure and marginal expenditure curves are identical horizontal lines, just as the marginal and average revenue curves are identical (and horizontal) for a competitive firm in the output market.

How much of the input should a firm facing a competitive factor market purchase? As long as the marginal revenue product curve lies above the marginal expenditure curve, profit can be increased by purchasing more of the input because the benefit of an additional unit (MRP) exceeds the cost (ME). However, when the marginal revenue product curve lies below the marginal expenditure curve, some units yield benefits that are less than cost. Therefore, profit maximization requires that *marginal revenue product be equal to marginal expenditure*:

$$\text{ME} = \text{MRP} \tag{14.5}$$

When we considered the special case of a competitive output market, we saw that the firm bought inputs, such as labor, up to the point at which the marginal revenue product was equal to the price of the input ω, as in equation (14.3). In the competitive case, therefore, the condition for profit maximization is that the price of the input be equal to marginal expenditure:

$$\text{ME} = w \tag{14.6}$$

In our example, the price of the fabric ($10 per yard) is determined in the competitive fabric market shown in Figure 14.7(a) at the intersection of the demand and supply curves. Figure 14.7(b) shows the amount of fabric purchased by a firm at the intersection of the marginal expenditure and marginal revenue product curves. When 50 yards of fabric are purchased, the marginal expenditure of $10 is equal to the marginal revenue from the sale of clothing made possible by the increased use of fabric in the production process. If less than 50 yards of fabric were purchased, the firm would be forgoing an opportunity to make additional profit from clothing sales. If more than 50 yards were purchased, the cost of the fabric would be greater than the additional revenue from the sale of the extra clothing.

The Market Supply of Inputs

In §8.6, we explain that the short-run market supply curve shows the amount of output that will be produced by firms in the market for every possible price.

The market supply curve for a factor input is usually upward sloping. We saw in Chapter 8 that the market supply for a good sold in a competitive market is usually upward sloping because the marginal cost of producing the good is typically increasing. This is also the case for fabric and other raw material inputs.

When the input is labor, however, people rather than firms are making supply decisions. In this case, utility maximization by workers rather than profit maximization by firms determines supply. In the discussion that follows, we use the analysis of income and substitution effects from Chapter 4 to show that although the market supply curve for labor can be upward sloping, it may also, as in Figure 14.8, be *backward bending*. In other words, a higher wage rate can lead to less labor being supplied.

In §4.2, we explain that an increase in the price of a good has two effects: The real purchasing power of each consumer decreases (the income effect) and the good becomes relatively expensive (the substitution effect).

To see why a labor supply curve may be backward bending, divide the day into hours of work and hours of leisure. *Leisure* is a term that describes enjoyable non-work activities, including sleeping and eating. *Work* benefits the worker only through the income that it generates. We also assume that a worker has the flexibility to choose how many hours per day to work.

The wage rate measures the price that the worker places on leisure time, because his or her wage measures the amount of money that the worker gives up to enjoy leisure. As the wage rate increases, therefore, the price of leisure also increases. This price change brings about both a substitution effect (a change in relative price with utility held constant) and an income effect (a change in utility with relative prices unchanged). There is a substitution effect because the higher price of leisure encourages workers to substitute work for leisure. An income effect occurs because the higher wage rate increases the worker's purchasing

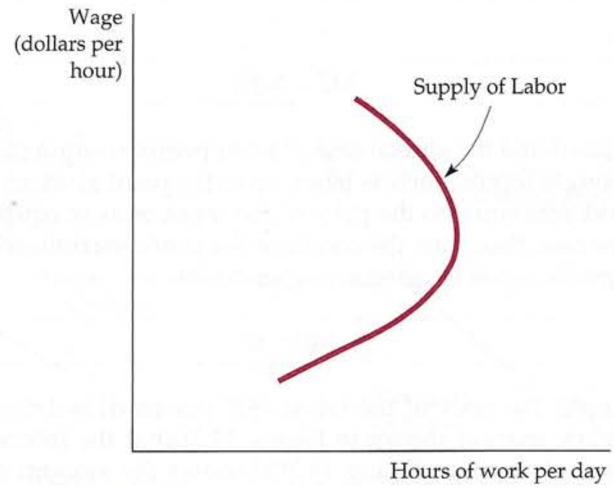

FIGURE 14.8 Backward-Bending Supply of Labor

When the wage rate increases, the hours of work supplied increase initially but can eventually decrease as individuals choose to enjoy more leisure and to work less. The backward-bending portion of the labor supply curve arises when the income effect of the higher wage (which encourages more leisure) is greater than the substitution effect (which encourages more work).

power. With higher income, the worker can buy more of many goods, one of which is leisure. If more leisure was chosen, it is because the income effect has encouraged the worker to work fewer hours. Income effects can be large because wages are the primary component of most people's income. When the income effect outweighs the substitution effect, the result is a backward-bending supply curve.

Figure 14.9 illustrates how a backward-bending supply curve for labor can result from the work-leisure decision for a typical weekday. The horizontal axis shows hours of leisure per day, the vertical axis income generated by work. (We assume there are no other sources of income.) Initially the wage rate is \$10 per hour, and the budget line is given by PQ. Point P, for example, shows that if an individual worked a 24-hour day he would earn an income of \$240.

The worker maximizes utility by choosing point A, thus enjoying 16 hours of leisure per day (with 8 hours of work) and earning \$80. When the wage rate increases to \$30 per hour, the budget line rotates about the horizontal intercept to line RQ. (Only 24 hours of leisure are possible.) Now the worker maximizes utility at B by choosing 19 hours of leisure per day (with 5 hours of work), while earning \$150. If only the substitution effect came into play, the higher wage rate would encourage the worker to work 12 hours (at C) instead of 8. However, the income effect works in the opposite direction. It overcomes the substitution effect and lowers the work day from 8 hours to 5.

In real life, a backward-bending labor supply curve might apply to a college student working during the summer to earn living expenses for the school year. As soon as a target level of earnings is reached, the student stops working and allocates more time to leisure. An increase in the wage rate will then lead to fewer hours worked because it enables the student to reach the target level of earnings more quickly.

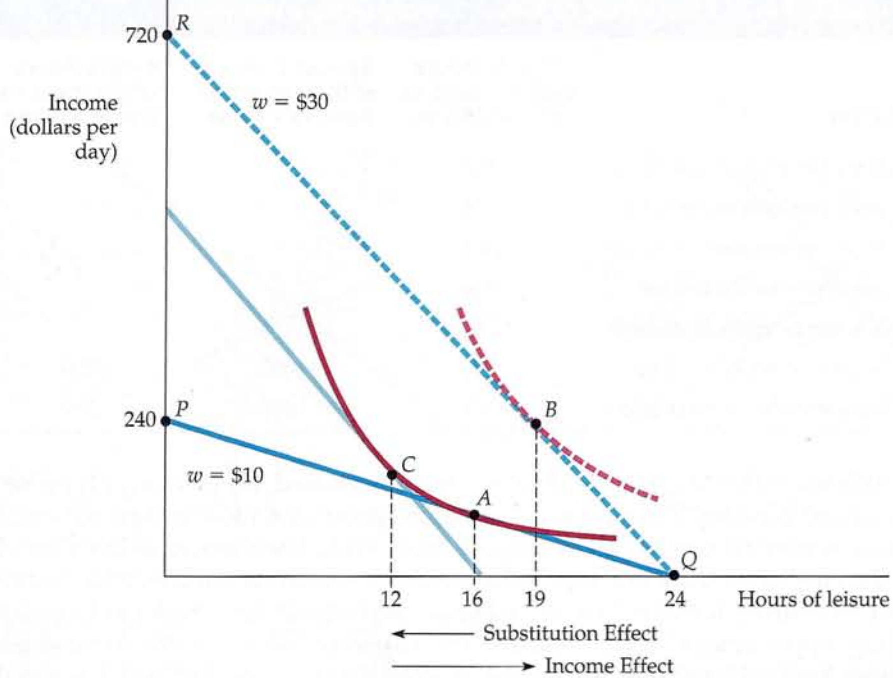

FIGURE 14.9 Substitution and Income Effects of a Wage Increase

When the wage rate increases from $10 to $30 per hour, the worker's budget line shifts from *PQ* to *RQ*. In response, the worker moves from *A* to *B* while decreasing work hours from 8 to 5. The reduction in hours worked arises because the income effect outweighs the substitution effect. In this case, the supply of labor curve is backward bending.

EXAMPLE 14.2 Labor Supply for One- and Two-Earner Households

One of the most dramatic changes in the labor market in the twentieth century has been the increase in women's participation in the labor force. Whereas only 34 percent of women had entered the labor force in 1950, the number had risen to well over 60 percent by 2007. Married women account for a substantial portion of this increase. The increased role of women in the labor market has also had a major impact on housing markets: Where to live and work has increasingly become a joint husband-and-wife decision.

The complex nature of the work choice was analyzed in a study that compared the work decisions of 94 unmarried females with the work decisions of heads of households and spouses in 397 families.[3] One way to describe the work decisions of the various family groups is to calculate labor supply elasticities. Each elasticity relates the numbers of hours worked not only to the wage earned by the head of the household, but also to the wage of the other member of two-earner households. Table 14.2 summarizes the results.

[3]See Janet E. Kohlhase, "Labor Supply and Housing Demand for One- and Two-Earner Households," *Review of Economics and Statistics* 68 (1986): 48–56; and Ray C. Fair and Diane J. Macunovich, "Explaining the Labor Force Participation of Women 20–24" (unpublished, February 1997).

TABLE 14.2 Elasticities of Labor Supply (Hours Worked)

Group	Head's Hours with Respect to Head's Wage	Spouse's Hours with Respect to Spouse's Wage	Head's Hours with Respect to Spouse's Wage
Unmarried males, no children	.026		
Unmarried females, children	.106		
Unmarried females, no children	.011		
One-earner family, children	–.078		
One-earner family, no children	.007		
Two-earner family, children	–.002	–.086	–.004
Two-earner family, no children	–.107	–.028	–.059

When a higher wage rate leads to fewer hours worked, the labor supply curve is backward bending: The income effect, which encourages more leisure, outweighs the substitution effect, which encourages more work. The elasticity of labor supply is then negative. Table 14.2 shows that heads of one-earner families with children and two-earner families (with or without children) all have backward-bending labor supply curves, with elasticities ranging from –.002 to –.078. Most single-earner heads of households are on the upward-sloping portion of the labor supply curve, with the largest elasticity of .106 associated with single women with children. Married women (listed as spouses of heads of households) are also on the backward-bending portion of the labor supply curve, with elasticities of –.028 and –.086.

14.2 EQUILIBRIUM IN A COMPETITIVE FACTOR MARKET

A competitive factor market is in equilibrium when the price of the input equates the quantity demanded to the quantity supplied. Figure 14.10(a) shows such an equilibrium for a labor market. At point *A*, the equilibrium wage rate is w_C and the equilibrium quantity supplied is L_C. Because they are well informed, all workers receive the identical wage and generate the identical marginal revenue product of labor wherever they are employed. If any worker had a wage lower than her marginal product, a firm would find it profitable to offer that worker a higher wage.

If the output market is also perfectly competitive, the demand curve for an input measures the benefit that consumers of the product place on the additional use of the input in the production process. The wage rate also reflects the cost to the firm and to society of using an additional unit of the input. Thus, at *A* in Figure 14.10(a), the marginal benefit of an hour of labor (its marginal revenue product MRP_L) is equal to its marginal cost (the wage rate w).

In §9.2, we explain that in a perfectly competitive market, efficiency is achieved because the sum of aggregate consumer and producer surplus is maximized.

When output and input markets are both perfectly competitive, resources are used efficiently because the difference between total benefits and total costs is maximized. Efficiency requires that the additional revenue generated by employing an additional unit of labor (the marginal revenue product of labor, MRP_L) equal the benefit to consumers of the additional output,

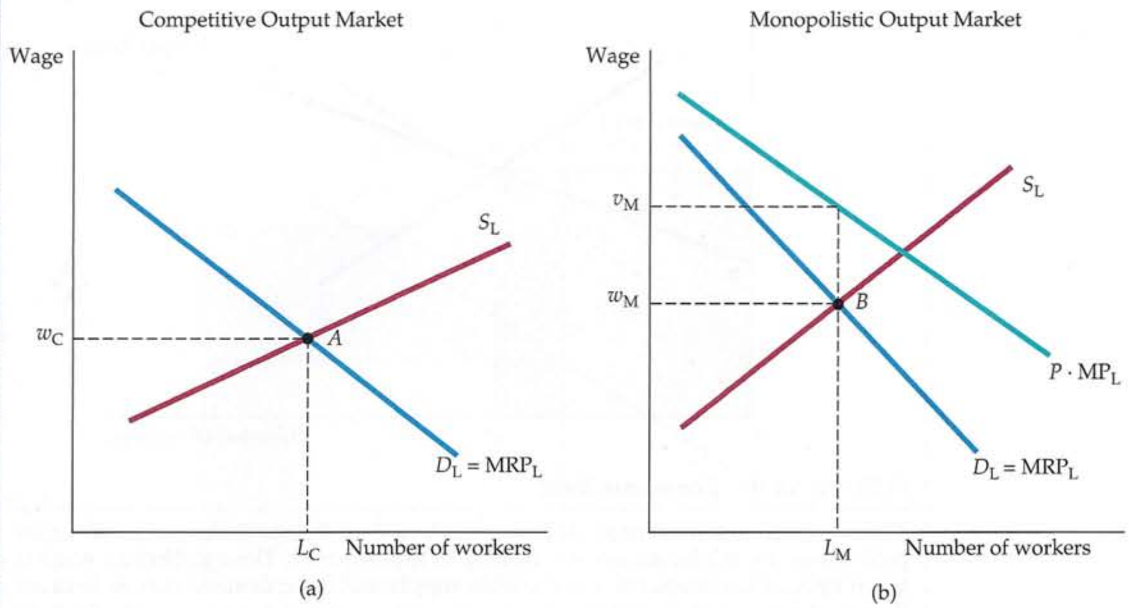

FIGURE 14.10 Labor Market Equilibrium

In a competitive labor market in which the output market is competitive, the equilibrium wage w_c is given by the intersection of the demand for labor (marginal revenue product) curve and the supply of labor curve. This is point A in part (a) of the figure. Part (b) shows that when the producer has monopoly power, the marginal value of a worker v_M is greater than the wage w_M. Thus too few workers are employed. (Point B determines the quantity of labor that the firm hires and the wage rate paid.)

which is given by the price of the product times the marginal product of labor, $(P)(\text{MP}_L)$.

When the output market is not perfectly competitive, the condition $\text{MRP}_L = (P)(\text{MP}_L)$ no longer holds. Note in Figure 14.10(b) that the curve representing the product price multiplied by the marginal product of labor $[(P)(\text{MP}_L)]$ lies above the marginal revenue product curve $[(\text{MR})(\text{MP}_L)]$. Point B is the equilibrium wage w_M and the equilibrium labor supply L_M. But because the price of the product is a measure of the value to consumers of each additional unit of output that they buy, $(\text{P})(\text{MP}_L)$ is the value that consumers place on additional units of labor. Therefore, when L_M laborers are employed, the marginal cost to the firm w_M is less than the marginal benefit to consumers v_M. Although the firm is maximizing its profit, its output is below the efficient level and it uses less than the efficient level of the input. Economic efficiency would be increased if more laborers were hired and, consequently, more output produced. (The gains to consumers would outweigh the firm's lost profit.)

Economic Rent

The concept of economic rent helps explain how factor markets work. When discussing output markets in the long run in Chapter 8, we defined economic rent as the payments received by a firm over and above the minimum cost of producing its output. For a factor market, *economic rent is the difference between the payments made to a factor of production and the minimum amount that must be*

In §8.7, we explain that economic rent is the amount that firms are willing to pay for an input less the minimum amount necessary to buy it.

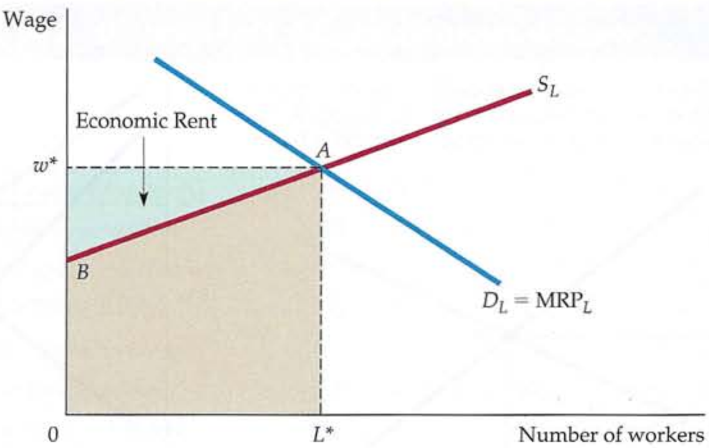

FIGURE 14.11 Economic Rent

The economic rent associated with the employment of labor is the excess of wages paid above the minimum amount needed to hire workers. The equilibrium wage is given by A, at the intersection of the labor supply and labor demand curves. Because the supply curve is upward sloping, some workers would have accepted jobs for a wage less than w^*. The green-shaded area ABw^* is the economic rent received by all workers.

spent to obtain the use of that factor. Figure 14.11 illustrates the concept of economic rent as applied to a competitive labor market. The equilibrium price of labor is w^*, and the quantity of labor supplied is L^*. The supply of labor curve is the upward-sloping curve, and the demand for labor is the downward-sloping marginal revenue product curve. Because the supply curve tells us how much labor will be supplied at each wage rate, the minimum expenditure needed to employ L^* units of labor is given by the tan-shaded area AL^*0B, below the supply curve to the left of the equilibrium labor supply L^*.

In perfectly competitive markets, all workers are paid the wage w^*. This wage is required to get the last "marginal" worker to supply his or her labor, but all other workers earn rents because their wage is greater than the wage that would be needed to get them to work. Because total wage payments are equal to the rectangle $0w^*AL^*$, the economic rent earned by labor is given by the area ABw^*.

Note that if the supply curve were perfectly elastic, economic rent would be zero. There are rents only when supply is somewhat inelastic. And when supply is perfectly inelastic, all payments to a factor of production are economic rents because the factor will be supplied no matter what price is paid.

As Figure 14.12 shows, one example of an inelastically supplied factor is land. The supply curve is perfectly inelastic because land for housing (or for agriculture) is fixed, at least in the short run. With land inelastically supplied, its price is determined entirely by demand. The demand for land is given by D_1, and its price per unit is s_1. Total land rent is given by the green-shaded rectangle. But when the demand for land increases to D_2, the rental value per unit of land increases to s_2; in this case, total land rent includes the blue-shaded area as well. Thus, an increase in the demand for land (a shift to the right in the demand curve) leads both to a higher price per acre and to a higher economic rent.

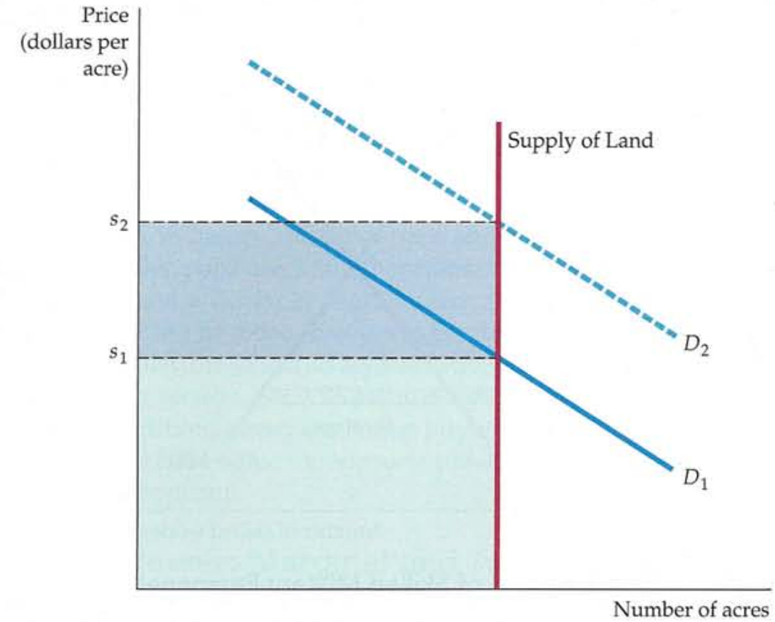

FIGURE 14.12 Land Rent

When the supply of land is perfectly inelastic, the market price of land is determined at the point of intersection with the demand curve. The entire value of the land is then an economic rent. When demand is given by D_1, the economic rent per acre is given by s_1, and when demand increases to D_2, rent per acre increases to s_2.

EXAMPLE 14.3 Pay in the Military

The U.S. Army has had a personnel problem for many years. During the Civil War, roughly 90 percent of the armed forces were unskilled workers involved in ground combat. Since then, however, the nature of warfare has evolved. Ground combat forces now make up only 16 percent of the armed forces. Meanwhile, changes in technology have led to a severe shortage in skilled technicians, trained pilots, computer analysts, mechanics, and others needed to operate sophisticated military equipment. Why has such a shortage developed? Why has the military been unable to keep skilled personnel? An economic study provides some answers.[4]

The rank structure of the army has remained essentially unchanged over the years. Among the officer ranks, pay increases are determined primarily by the number of years of service. Consequently, officers with differing skill levels and abilities are usually paid similar salaries. Moreover, some skilled workers are underpaid relative to salaries that they could receive in the private sector. As a

[4]Walter Y. Oi, "Paying Soldiers: On a Wage Structure for a Large Internal Labor Market" (unpublished, undated paper).

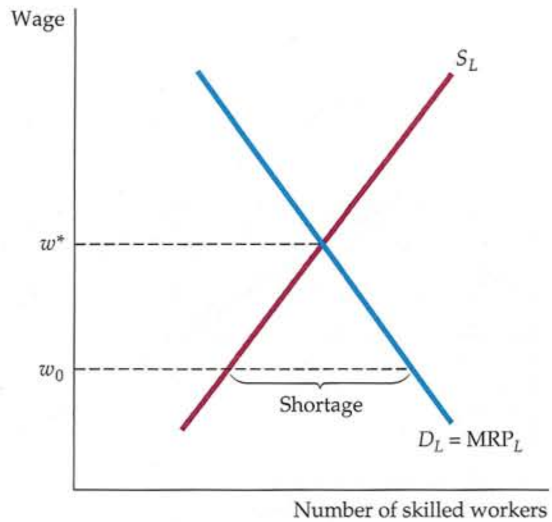

FIGURE 14.13 The Shortage of Skilled Military Personnel

When the wage w^* is paid to military personnel, the labor market is in equilibrium. When the wage is kept below w^*, at w_0, there is a shortage of personnel because the quantity of labor demanded is greater than the quantity supplied.

result, skilled workers who join the army because of attractive salaries find that their marginal revenue products are eventually higher than their wages. Thus, a married Air Force pilot with eight years of training would have earned an annual salary of $45,000 in 1989. This salary would increase gradually to $61,000 in 2009, but is substantially less than a commercial airline pilot would make with only 10 years of service.[5] Some officers remain in the army, but many leave.

This study of army pay applies to all of the armed forces. Figure 14.13 shows the inefficiency that can result from the military pay policy. The equilibrium wage rate w^* is the wage that equates the demand for labor to the supply. Because of inflexibility in its pay structure, however, the military pays the wage w_0, which is below the equilibrium wage. At w_0, demand is greater than supply, and there is a shortage of skilled labor. By contrast, competitive labor markets pay more productive workers higher wages than their less productive counterparts.

So how can the military attract and keep a skilled labor force?

The military's choice of wage structure affects the nation's ability to maintain an effective fighting force, and a 7.2-percent military pay raise went into effect in 2007. Even so, military pay remains low: A first-class private earns $18,400, a sergeant $22,250, a captain $39,500, and a major $44,950. In response to its personnel problems, the military has begun to change its salary structure by expanding the number and size of its reenlistment bonuses. Selective reenlistment bonuses targeted at skilled jobs for which there are shortages can be an effective recruiting device. The immediate bonuses create more of an incentive than the uncertain promise of higher wages in the future. As the demand for skilled military jobs increases, we can expect the armed forces to make greater use of reenlistment bonuses and other market-based incentives.

[5]Department of Defense, Department of Defense Aviator Retention Study—1988, Table 2-4. (Washington: GPO, November 28, 1988).

14.3 FACTOR MARKETS WITH MONOPSONY POWER

In some factor markets, individual buyers have *buyer power* that allows them to affect the prices they pay. Often this happens either when one firm is a monopsony buyer or there are only a few buyers, in which case each firm has some monopsony power. For example, we saw in Chapter 10 that automobile companies have monopsony power as buyers of parts and components. GM and Toyota, for example, buy large quantities of brakes, radiators, and other parts and can negotiate lower prices than those charged smaller purchasers. In other cases, there might be only two or three sellers of a factor and a dozen or more buyers, but each buyer nonetheless has *bargaining power*—it can negotiate low prices because it makes large and infrequent purchases and can play the sellers off against each other when bargaining over price.

In §10.5, we explain buyer has monopsor when his purchasing can affect the price c product.

Throughout this section, we will assume that the output market is perfectly competitive. In addition, because a single buyer is easier to visualize than several buyers who all have some monopsony power, we will restrict our attention at first to pure monopsony.

Monopsony Power: Marginal and Average Expenditure

When you are deciding how much of a good to purchase, you keep increasing the number of units purchased until the additional value from the last unit purchased—the *marginal value*—is just equal to the cost of that unit—the *marginal expenditure*. In perfect competition, the price that you pay for the good—the *average expenditure*—is equal to the marginal expenditure. However, when you have monopsony power, the marginal expenditure is greater than the average expenditure, as Figure 14.14 shows.

In §10.5, we explain marginal expenditure cost of one more uni average expenditure average price paid p

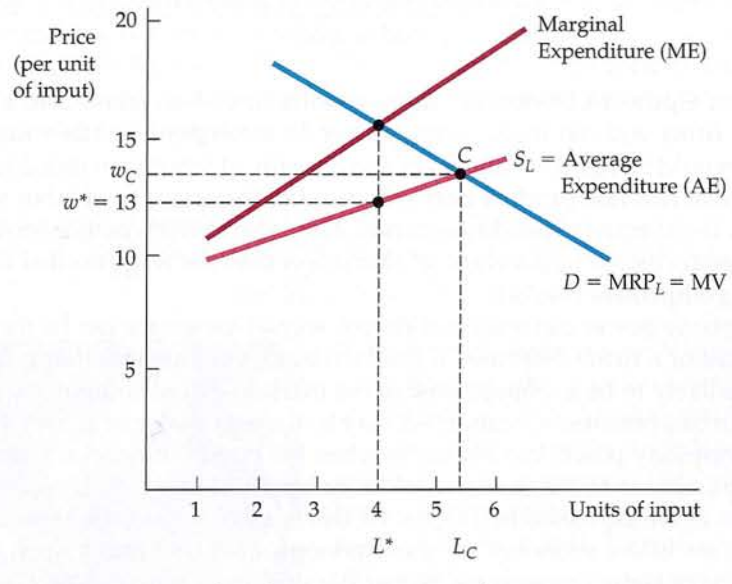

FIGURE 14.14 Marginal and Average Expenditure

When the buyer of an input has monopsony power, the marginal expenditure curve lies above the average expenditure curve because the decision to buy an extra unit raises the price that must be paid for all units, not just for the last one. The number of units of input purchased is given by L^*, at the intersection of the marginal revenue product and marginal expenditure curves. The corresponding wage rate w^* is lower than the competitive wage w_c.

The factor supply curve facing the monopsonist is the market supply curve, which shows how much of the factor suppliers are willing to sell as its price increases. Because the monopsonist pays the same price for each unit, the supply curve is its *average expenditure curve*. The average expenditure curve is upward sloping because the decision to buy an extra unit raises the price that must be paid for all units, not just the last one. For a profit-maximizing firm, however, the *marginal expenditure curve* is relevant in deciding how much to buy. The marginal expenditure curve lies above the average expenditure curve: When the firm increases the price of the factor to hire more units, it must pay *all* units that higher price, not just the last unit hired.

Purchasing Decisions with Monopsony Power

How much of the input should the firm buy? As we saw earlier, it should buy up to the point where marginal expenditure equals marginal revenue product. Here the benefit from the last unit bought (MRP) is just equal to the cost (ME). Figure 14.14 illustrates this principle for a labor market. Note that the monopsonist hires L^* units of labor; at that point, $\text{ME} = \text{MRP}_L$. The wage rate w^* that workers are paid is given by finding the point on the average expenditure or supply curve with L^* units of labor.

As we showed in Chapter 10, a buyer with monopsony power maximizes net benefit (utility less expenditure) from a purchase by buying up to the point where marginal value (MV) is equal to marginal expenditure:

$$\text{MV} = \text{ME}$$

For a firm buying a factor input, MV is just the marginal revenue product of the factor MRP. Thus, we have (as in the case of a competitive factor market)

$$\text{ME} = \text{MRP} \quad \textbf{(14.7)}$$

Note from Figure 14.14 that the monopsonist hires less labor than a firm or group of firms with no monopsony power. In a competitive labor market, L_C workers would be hired: At that level, the quantity of labor demanded (given by the marginal revenue product curve) is equal to the quantity of labor supplied (given by the average expenditure curve). Note also that the monopsonistic firm will be paying its workers a wage w^* that is less than the wage w_C that would be paid in a competitive market.

Monopsony power can arise in different ways. One source can be the specialized nature of a firm's business. If the firm buys a component that no one else buys, it is likely to be a monopsonist in the market for that component. Another source can be a business's location—it may be the only major employer within an area. Monopsony power can also arise when the buyers of a factor form a cartel to limit purchases of the factor, in order to buy it at less than the competitive price. (But as we explained in Chapter 10, this is a violation of the antitrust laws.)

Few firms in our economy are pure monopsonists. But many firms (or individuals) have some monopsony power because their purchases account for a large portion of the market. The government is a monopsonist when it hires volunteer soldiers or buys missiles, aircraft, and other specialized military equipment. A mining firm or other company that is the only major employer in a community also has monopsony power in the local labor market. Even in these cases, however, monopsony power may be limited because the government competes to some extent with other firms that offer similar jobs. Likewise, the mining firm competes to some extent with companies in nearby communities.

Bargaining Power

In some factor markets, there are a small number of sellers and a small number of buyers. In such cases, an individual buyer and an individual seller will negotiate with each other to determine a price. The resulting price might be high or low, depending on which side has more bargaining power.

The amount of bargaining power that a buyer or seller has is determined in part by the number of competing buyers and competing sellers. But it is also determined by the nature of the purchase itself. If each buyer makes large and infrequent purchases, it can sometimes play the sellers off against each other when negotiating a price and thereby amass considerable bargaining power.

An example of this kind of bargaining power occurs in the market for commercial aircraft. Airplanes are clearly key factor inputs for airlines, and airlines want to buy planes at the lowest possible prices. There are dozens of airlines, however, and only two major producers of commercial aircraft—Boeing and Airbus. One might think that as a result, Boeing and Airbus would have a considerable advantage when negotiating prices. The opposite is true. It is important to understand why.

Airlines do not buy planes every day, and they do not usually buy one plane at a time. A company like American Airlines will typically order new planes only every three or four years, and each order might be for 20 or 30 planes, at a cost of several billion dollars. As big as Boeing and Airbus are, this is no small purchase, and each seller will do all it can to win the order. American Airlines knows this and can use it to its advantage. If, for example, American is choosing between 20 new Boeing 777s or 20 new Airbus A340s (which are similar airplanes), it can play the two companies off against each other when negotiating a price. Thus if Boeing offers a price of, say, \$150 million per plane, American might go to Airbus and ask it to do better. Whatever Airbus offers, American will then go back to Boeing and demand a bigger discount, claiming (truthfully or otherwise) that Airbus is offering large discounts. Then back to Airbus, back to Boeing, and so on, until American has succeeded in obtaining a large discount from one of the two companies.

EXAMPLE 14.4 Monopsony Power in the Market for Baseball Players

In the United States, major league baseball is exempt from the antitrust laws, the result of a Supreme Court decision and the policy of Congress not to apply those laws to labor markets.[6] This exemption allowed baseball team owners (before 1975) to operate a monopsonistic cartel. Like all cartels, this one depended on an agreement among owners. The agreement involved an annual draft of players and a *reserve clause* that effectively tied each player to one team for life, thereby eliminating most interteam competition for players. Once a player was drafted by a team, he could not play for another team unless rights

[6]This example builds on an analysis of the structure of baseball players' salaries by Roger Noll, who has kindly supplied us with the relevant data.

were sold to that team. As a result, baseball owners had monopsony power in negotiating new contracts with their players: The only alternative to signing an agreement was to give up the game or play it outside the United States.

During the 1960s and early 1970s, baseball players' salaries were far below the market value of their marginal products (determined in part by the incremental attention that better hitting or pitching might achieve). For example, if the players' market had been perfectly competitive, those players receiving a salary of about $42,000 in 1969 would have instead received a salary of $300,000 in 1969 dollars (which is $1.7 million in year 2007 dollars).

Fortunately for the players, and unfortunately for the owners, there was a strike in 1972 followed by a lawsuit by one player (Curt Flood of the St. Louis Cardinals) and an arbitrated labor-management agreement. This process eventually led in 1975 to an agreement by which players could become free agents after playing for a team for six years. The reserve clause was no longer in effect, and a highly monopsonistic labor market became much more competitive.

The result was an interesting experiment in labor market economics. Between 1975 and 1980, the market for baseball players adjusted to a new post-reserve clause equilibrium. Before 1975, expenditures on players' contracts made up approximately 25 percent of all team expenditures. By 1980, those expenditures had increased to 40 percent. Moreover, the average player's salary doubled in real terms. By 1992, the average baseball player was earning $1,014,942—a very large increase from the monopsonistic wages of the 1960s. In 1969, for example, the average baseball salary was approximately $42,000 adjusted for inflation, about $236,000 in year 2007 dollars.

Salaries for baseball players continued to grow. Whereas the average salary was just less than $600,000 in 1990, it had risen to $1,998,000 by 2000 and $2,950,000 by 2007, and many players earned much more. The New York Yankees as a team averaged over $8,010,000 in 2005.

EXAMPLE 14.5 Teenage Labor Markets and the Minimum Wage

Increases in the national minimum wage rate (which was $4.50 in early 1996 and $5.15 in 1999) were controversial, raising the question of whether the cost of any unemployment that might be generated would be outweighed by the benefit of higher incomes to those whose wages have been increased.[7] A study of the effects of the minimum wage on employment in fast-food restaurants in New Jersey added to that controversy.[8]

[7]See Example 1.4 (page 14) for an initial discussion of the minimum wage, and Section 9.3 (page 319) for an analysis of its effects on employment.

[8]David Card and Alan Krueger, "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," *American Economic Review* 84 (September 1994). See also David Card and Alan B. Krueger, "A Reanalysis of the Effect of the New Jersey Minimum Wage on the Fast-Food Industry with Representative Payroll Data," Working Paper No. 6386, Cambridge, MA: National Bureau of Economic Research, 1998; and Madeline Zavodny, "Why Minimum Wage Hikes May Not Reduce Employment," Federal Reserve Bank of Atlanta, Economic Review, Second Quarter, 1998.

In §9.3, we explain that setting a minimum wage in a perfectly competitive market can create unemployment and a deadweight loss.

Some states have minimum wages above the Federal level. In April 1992 the New Jersey minimum wage was increased from \$4.25 to \$5.05 per hour. Using a survey of 410 fast-food restaurants, David Card and Alan Krueger found that employment had actually *increased* by 13 percent after the minimum wage went up. What is the explanation for this surprising result? One possibility is that restaurants responded to the higher minimum wage by reducing fringe benefits, which usually take the form of free and reduced-price meals for employees. A related explanation is that employers responded by providing less on-the-job training and by lowering the wages for those with experience who had previously been paid more than the new minimum wage.

An alternative explanation for the increased New Jersey employment holds that the labor market for teenage (and other) unskilled workers is not highly competitive. If so, the analysis of Chapter 9 does not apply. If the unskilled fast-food labor market were monopsonistic, for example, we would expect a different effect from the increased minimum wage. Suppose that the wage of \$4.25 was the wage that fast-food employers with monopsony power in the labor market would offer their workers even if there were no minimum wage. Suppose also that \$5.10 would be the wage enjoyed by workers if the labor market were fully competitive. As Figure 14.14 shows, the increase in the minimum wage would not only raise the wage, but would also increase the employment level (from L^* to L_C).

Does the fast-food study show that employers have monopsony power in this labor market? The evidence suggests no. If firms do have monopsony power but the fast-food market is competitive, then the increase in the minimum wage should have no effect on the price of fast food. Because the market for fast food is so competitive, firms paying the higher minimum wage would be forced to absorb the higher wage cost themselves. The study suggests, however, that prices did increase after the introduction of the higher minimum wage.

The Card-Krueger analysis of the minimum wage remains hotly debated. A number of critics argued that the New Jersey study was atypical. Others questioned the reliability of the data, arguing that a higher minimum wage reduces employment (see our discussion in Chapter 9).[9] In response, Card and Krueger repeated their study, using a more comprehensive and accurate data set. They obtained the same results. Where does this leave us? Perhaps a better characterization of low-wage labor markets requires a more sophisticated theory (e.g., the efficiency wage theory discussed in Chapter 17). In any case, new empirical analyses should shed more light on the effects of the minimum wage.

14.4 FACTOR MARKETS WITH MONOPOLY POWER

Just as buyers of inputs can have monopsony power, sellers of inputs can have monopoly power. In the extreme, the seller of an input may be a monopolist, as when a firm has a patent to produce a computer chip that no other firm can duplicate. Because the most important example of monopoly power in factor

[9]For example, see Donald Deere, Kevin M. Murply, and Finis Welch, "Employment and the 1990–1991 Minimum Wage Hike," *American Economic Review, Papers and Proceedings* 85 (May 1995): 232–37; and David Neumark and William Wascher, "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Comment," *American Economic Review* 90 (2000): 1362–96.

In §10.2, we explain that a seller of a product has some monopoly power if it can profitably charge a price greater than marginal cost.

markets involves labor unions, we will concentrate most of our attention there. In the subsections that follow, we show how a labor union, which is a monopolist in the sale of labor services, might increase the well-being of its members and substantially affect nonunionized workers.

Monopoly Power over the Wage Rate

Figure 14.15 shows a demand for labor curve in a market with no monopsony power: It aggregates the marginal revenue products of firms that compete to buy labor. The labor supply curve describes how union members would supply labor *if* the union exerted no monopoly power. In that case, the labor market would be competitive, and L^* workers would be hired at a wage of w^*, where demand D_L equals supply S_L.

Because of its monopoly power, however, the union can choose any wage rate and the corresponding quantity of labor supplied, just as a monopolist seller of output chooses price and the corresponding quantity of output. If the union wanted to maximize the number of workers hired, it would choose the competitive outcome at A. However, if the union wished to obtain a higher-than-competitive wage, it could restrict its membership to L_1 workers. As a result, the firm would pay a wage rate of w_1. Although union members who work would be better off, those who cannot find jobs would be worse off.

Is a policy of restrictive union membership worthwhile? If the union wishes to maximize the economic rent that its workers receive, the answer is yes. By

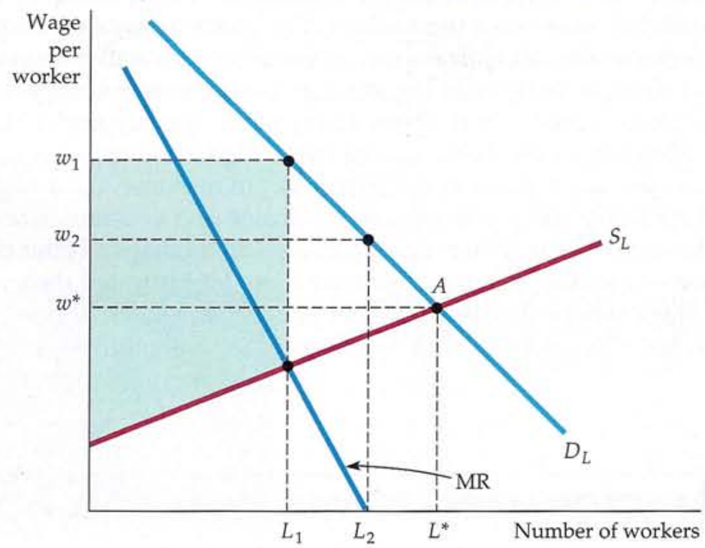

FIGURE 14.15 Monopoly Power of Sellers of Labor

When a labor union is a monopolist, it chooses among points on the buyer's demand for labor curve D_L. The seller can maximize the number of workers hired, at L^*, by agreeing that workers will work at wage w^*. The quantity of labor L_1 that maximizes the rent earned by employees is determined by the intersection of the marginal revenue and supply of labor curves; union members will receive a wage rate of w_1. Finally, if the union wishes to maximize total wages paid to workers, it should allow L_2 union members to be employed at a wage rate of w_2. At that point, the marginal revenue to the union will be zero.

restricting membership, the union would be acting like a monopolist, which restricts output in order to maximize profit. To a firm, profit is the revenue that it receives less its opportunity costs. To a union, rent represents the wages that its members earn as a group in excess of their opportunity cost. To maximize rent, the union must choose the number of workers hired so that the marginal revenue to the union (the additional wages earned) is equal to the extra cost of inducing workers to work. This cost is a *marginal* opportunity cost because it is a measure of what an employer has to offer an additional worker to get him or her to work for the firm. However, the wage that is necessary to encourage additional workers to take jobs is given by the supply of labor curve S_L.

In §7.1, we explain that opportunity cost is the cost associated with opportunities that are foregone by not putting a firm's resources to their best alternative use.

The rent-maximizing combination of wage rate and number of workers is given by the intersection of the MR and S_L curves. We have chosen the wage-employment combination of w_1 and L_1 with the rent-maximization premise in mind. The shaded area below the demand for labor curve, above the supply of labor curve and to the left of L_1, represents the economic rent that all workers receive.

A rent-maximizing policy might benefit nonunion workers if they can find nonunion jobs. However, if these jobs are not available, rent maximization could create too sharp a distinction between winners and losers. An alternative objective is to maximize the aggregate wages that all union members receive. Look again at the example in Figure 14.15. To achieve this goal, the number of workers hired is increased from L_1 until the marginal revenue to the union is equal to zero. Because any further employment decreases total wage payments, aggregate wages are maximized when the wage is equal to w_2 and the number of workers is equal to L_2.

Unionized and Nonunionized Workers

When the union uses its monopoly power to increase members' wages, fewer unionized workers are hired. Because these workers either move to the nonunion sector or choose initially not to join the union, it is important to understand what happens in the nonunionized part of the economy.

Assume that the total supply of unionized and nonunionized workers is fixed. In Figure 14.16, the market supply of labor in both sectors is given by S_L. The demand for labor by firms in the unionized sector is given by D_U, the demand in the nonunionized sector by D_{NU}. Total market demand is the horizontal sum of the demands in the two sectors and is given by D_L.

Suppose the union chooses to increase the wage rate of its workers above the competitive wage w^*, to w_U. At that wage rate, the number of workers hired in the unionized sector falls by an amount ΔL_U, as shown on the horizontal axis. As these workers find employment in the nonunionized sector, the wage rate in that sector adjusts until the labor market is in equilibrium. At the new wage rate in the nonunionized sector, w_{NU}, the additional number of workers hired in that sector, ΔL_{NU}, is equal to the number of workers who left the unionized sector.

Figure 14.16 shows an adverse consequence of a union strategy directed toward raising union wages: Nonunionized wages fall. Unionization can improve working conditions and provide useful information to workers and management. But when the demand for labor is not perfectly inelastic, union workers are helped at the expense of nonunion workers.

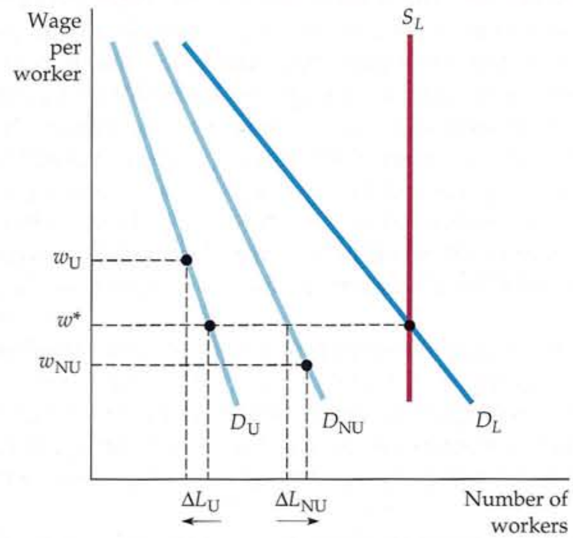

FIGURE 14.16 Wage Discrimination in Unionized and Nonunionized Sectors

When a monopolistic union raises the wage in the unionized sector of the economy from w^* to w_U, employment in that sector falls, as shown by the movement along the demand curve D_U. For the total supply of labor, given by S_L, to remain unchanged, the wage in the nonunionized sector must fall from w^* to w_{NU}, as shown by the movement along the demand curve D_{NU}.

EXAMPLE 14.6 The Decline of Private-Sector Unionism

For several decades, both the membership and bargaining power of labor unions have been declining.[10] A decline in union monopoly power can lead to different responses by union negotiators and can also affect the wage rate and level of employment. During the 1970s, most of the impact was on union wages: Although levels of employment did not change much, the differential between union and nonunion wages decreased substantially. We would have expected the same pattern to occur in the 1980s because of heavily publicized wage freezes and the rapid growth of two-tier wage provisions in which newer union members are paid less than more experienced counterparts.

Surprisingly, however, the union-management bargaining process changed during this period. From 1980 to 1984, the level of unionized employment fell from 23 percent to 19 percent. Yet the union-nonunion wage differential remained stable—and in fact grew wider in some industries. For example, the union wage rate in mining, forestry, and fisheries declined only slightly, from 25 percent higher than the nonunion wage in 1980 to 24 percent higher in 1984. On the other hand, the union wage rate in manufacturing increased slightly—from approximately 14 percent higher than the nonunion wage in 1980 to 16 percent in 1984. This same pattern has continued over the years. As Figure 14.17 shows, by 2006, unionized employment had fallen to below 12 percent of total employment. The union-nonunion wage differential remained essentially unchanged. In recent years, nonunion wages have increased faster than union wages.

[10]This example is based on Richard Edwards and Paul Swaim, "Union-Nonunion Earnings Differentials and the Decline of Private-Sector Unionism," *American Economic Review* 76 (May 1986): 97–102.

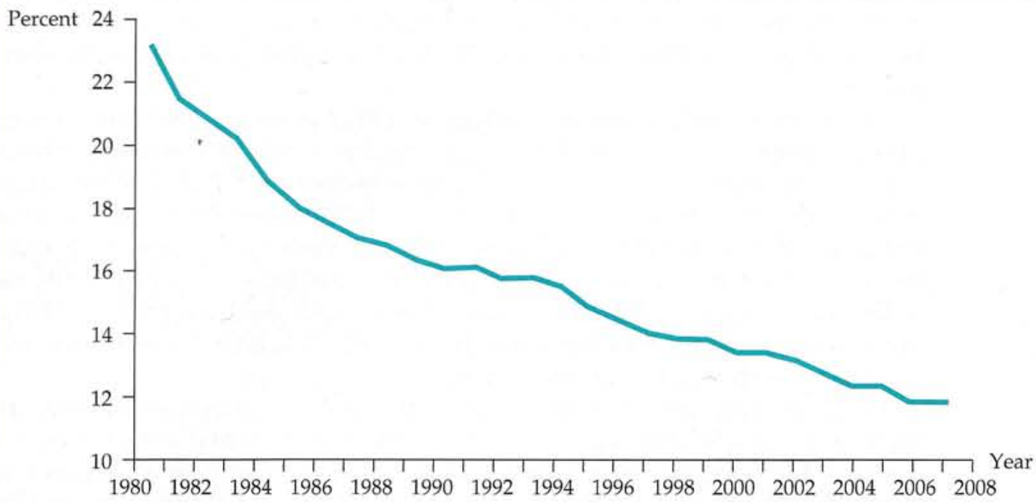

FIGURE 14.17 Union Workers as a Percentage of Total

The percentage of workers that are unionized has been declining steadily over the past 25 years.

Source: Bureau of Labor Statistics.

Nevertheless, the union-nonunion wage differential remains significant. For example, data from the Employer Costs for Employee Compensation Survey show that in 2006, wages and salaries for private-industry union workers averaged \$21.65 per hour, compared with \$17.59 for nonunion workers.

One explanation for this pattern of wage-employment responses is a change in union strategy—a move to maximize the wage rate for its members rather than the total wages paid to all union members. However, the demand for unionized employees has probably become increasingly elastic over time as firms find it easier to substitute capital for skilled labor in the production process. Faced with an elastic demand for its services, the union would have little choice but to maintain the wage rate of its members and allow employment levels to fall. Of course, the substitution of nonunion for union workers may cause further losses in the bargaining power of labor unions. How this will affect the differential between union and nonunion wages remains to be seen.

EXAMPLE 14.7 Wage Inequality—Have Computers Changed the Labor Market?

In Example 2.2 (page 29), we explained how the rapid growth in the demand for skilled relative to unskilled labor has been partly responsible for the growing inequality in the distribution of income in the United States. What is the underlying source of that change in relative demand? Is it the decline in private-sector unionism and the failure of the minimum wage to keep up with inflation? Or is it the increasing role that computers now play in the labor market? A recent study, which focuses on the wages of college relative to high school graduates, provides some answers.[11]

From 1950 to 1980, the relative wages of college graduates (the ratio of their average wages to those of high school graduates) hardly changed. In contrast, college graduates' (relative) wages grew rapidly from 1980 to 1995. This pattern is not consistent with what one would expect if the decline of unionism (and/or changes in the minimum wage) was the primary reason for the growing inequality. A clue to what happened is given by the dramatic increase in the use of computers by workers. In 1984, 25.1 percent of all workers used computers; that figure increased to 45.8 percent by 1993 and 56.1 percent in 2003. For managers and professionals, the figure was over 80 percent.

While computer use increased for all workers, the largest increases were registered by workers with college degrees—from 42 to 82 percent. For those without a high school degree, the increase was only 11 percentage points (from 5 to 16 percent); for those with high school degrees, the increase was 21 percentage points (from 19 to 40 percent).

Further analysis of data on jobs and wages confirms the importance of computers. Education and computer use have gone hand-in-hand to increase the demand for skilled workers. The wages of college graduates who use computers (relative to high school graduates) grew by about 11 percent from 1983 to 1993; for noncomputer users, wages grew by less than 4 percent. A statistical analysis shows that, overall, the spread of computer technology is responsible for nearly half the increase in relative wages during this period. Furthermore, the growth in the demand for skilled workers has occurred primarily within industries where computers have become increasingly useful.

Is this increase in the relative wages of skilled workers necessarily a bad thing? At least one economist suggests that the answer is no.[12] It is true that the growing inequality can disadvantage low-wage workers, whose limited opportunities might lead them to drop out of the labor force; in the extreme, they might even turn to crime. However, it can also motivate workers, whose opportunities for upward mobility through high-wage jobs have never been better.

Consider the circumstances facing men and women who are deciding whether to complete high school or college. We'll take the median wage of someone who completed high school as the norm. In 2005, college graduates age 25 and over earned on average $500 more per week than workers who stopped with a high school diploma. This figure translates into a real-wage increase for college graduates and a real-wage decrease for high school dropouts in comparison to 1979. Furthermore, the unemployment rate among college graduates is only one-third that among high school dropouts. The college wage premium has more than doubled over the past 30 years and provides strong motivation for college students to finish their studies.

[11]David H. Autor, Lawrence Katz, and Alan B. Krueger, "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (November 1998): 1169–1213.

[12]Finis Welch, "In Defense of Inequality," *American Economic Association, Papers and Proceedings* 89 (May 1999): 1–17.

SUMMARY

1. In a competitive input market, the demand for an input is given by the marginal revenue product, the product of the firm's marginal revenue, and the marginal product of the input.
2. A firm in a competitive labor market will hire workers to the point at which the marginal revenue product of labor is equal to the wage rate. This principle is analogous to the profit-maximizing output condition that production be increased to the point at which marginal revenue is equal to marginal cost.
3. The market demand for an input is the horizontal sum of industry demands for the input. But industry demand is not the horizontal sum of the demands of all the firms in the industry. To determine industry demand, one must remember that the market price of the product will change in response to changes in the price of an input.
4. When factor markets are competitive, the buyer of an input assumes that its purchases will have no effect on its price. As a result, the firm's marginal expenditure and average expenditure curves are both perfectly elastic.
5. The market supply of a factor such as labor need not be upward sloping. A backward-bending labor supply curve can result if the income effect associated with a higher wage rate (more leisure is demanded because it is a normal good) is greater than the substitution effect (less leisure is demanded because its price has gone up).
6. Economic rent is the difference between the payments to factors of production and the minimum payment that would be needed to employ them. In a labor market, rent is measured by the area below the wage level and above the marginal expenditure curve.
7. When a buyer of an input has monopsony power, the marginal expenditure curve lies above the average expenditure curve, which reflects the fact that the monopsonist must pay a higher price to attract more of the input into employment.
8. When the input seller is a monopolist, such as a labor union, the seller chooses the point on the marginal revenue product curve that best suits its objective. Maximization of employment, economic rent, and wages are three plausible objectives for labor unions.

QUESTIONS FOR REVIEW

1. Why is a firm's demand for labor curve more inelastic when the firm has monopoly power in the output market than when the firm is producing competitively?
2. Why might a labor supply curve be backward bending?
3. How is a computer company's demand for computer programmers a derived demand?
4. Compare the hiring choices of a monopsonistic and a competitive employer of workers. Which will hire more workers, and which will pay the higher wage? Explain.
5. Rock musicians sometimes earn several million dollars per year. Can you explain such large incomes in terms of economic rent?
6. What happens to the demand for one input when the use of a complementary input increases?
7. For a monopsonist, what is the relationship between the supply of an input and the marginal expenditure on it?
8. Currently the National Football League has a system for drafting college players by which each player is picked by only one team. The player must sign with that team or not play in the league. What would happen to the wages of both newly drafted and more experienced football players if the draft system were repealed and all teams could compete for college players?
9. The government wants to encourage individuals on welfare to become employed. It is considering two possible incentive programs:
   a. Give firms $2 per hour for every individual on welfare who is hired.
   b. Give each firm that hires one or more welfare workers a payment of $1000 per year, irrespective of the number of hires.

   To what extent is each of these programs likely to be effective at increasing the employment opportunities for welfare workers?
10. A small specialty cookie company, whose only variable input is labor, finds that the average worker can produce 50 cookies per day, the cost of the average worker is $64 per day, and the price of a cookie is $1. Is the company maximizing its profit? Explain.
11. A firm uses both labor and machines in production. Explain why an increase in the average wage rate causes both a movement along the labor demand curve and a shift of the curve.

EXERCISES

1. Suppose that the wage rate is \$16 per hour and the price of the product is \$2. Values for output and labor are in units per hour.

q	L
0	0
20	1
35	2
47	3
57	4
65	5
70	6

   a. Find the profit-maximizing quantity of labor.
   b. Suppose that the price of the product remains at \$2 but that the wage rate increases to \$21. Find the new profit-maximizing level of L.
   c. Suppose that the price of the product increases to \$3 and the wage remains at \$16 per hour. Find the new profit-maximizing L.
   d. Suppose that the price of the product remains at \$2 and the wage at \$16, but that there is a technological breakthrough that increases output by 25 percent for any given level of labor. Find the new profit-maximizing L.
2. Assume that workers whose incomes are less than \$10,000 currently pay no federal income taxes. Suppose a new government program guarantees each worker \$5000, whether or not he or she earns any income. For all earned income up to \$10,000, the worker must pay a 50-percent tax. Draw the budget line facing the worker under this new program. How is the program likely to affect the labor supply curve of workers?
3. Using your knowledge of marginal revenue product, explain the following:
   a. A famous tennis star is paid \$200,000 for appearing in a 30-second television commercial. The actor who plays his doubles partner is paid \$500.
   b. The president of an ailing savings and loan is paid *not* to stay in his job for the last two years of his contract.
   c. A jumbo jet carrying 400 passengers is priced higher than a 250-passenger model even though both aircraft cost the same to manufacture.
4. The demands for the factors of production listed below have increased. What can you conclude about changes in the demands for the related consumer goods? If demands for the consumer goods remain unchanged, what other explanation is there for an increase in derived demands for these items?
   a. Computer memory chips
   b. Jet fuel for passenger planes
   c. Paper used for newsprint
   d. Aluminum used for beverage cans
5. Suppose there are two groups of workers, unionized and nonunionized. Congress passes a law that requires all workers to join the union. What do you expect to happen to the wage rates of formerly nonunionized workers? Of those workers who were originally unionized? What have you assumed about the union's behavior?
6. Suppose that a firm's production function is given by $Q = 12L - L^2$, for $L = 0$ to 6, where L is labor input per day and Q is output per day. Derive and draw the firm's demand for labor curve if the firm's output sells for \$10 in a competitive market. How many workers will the firm hire when the wage rate is \$30 per day? \$60 per day? (*Hint:* The marginal product of labor is $12 - 2L$.)
7. The only legal employer of military soldiers in the United States is the federal government. If the government uses its knowledge of its monopsonistic position, what criteria will it employ when determining how many soldiers to recruit? What happens if a mandatory draft is implemented?
8. The demand for labor by an industry is given by the curve $L = 1200 - 10w$, where L is the labor demanded per day and w is the wage rate. The supply curve is given by $L = 20w$. What is the equilibrium wage rate and quantity of labor hired? What is the economic rent earned by workers?
9. Using the same information as in Exercise 8, suppose now that the only labor available is controlled by a monopolistic labor union that wishes to maximize the rent earned by union members. What will be the quantity of labor employed and the wage rate? How does your answer compare with your answer to Exercise 8? Discuss. (*Hint:* The union's marginal revenue curve is given by $MR = 120 - 0.2L$.)

*10. A firm uses a single input, labor, to produce output q according to the production function $q = 8\sqrt{L}$. The commodity sells for \$150 per unit and the wage rate is \$75 per hour.
   a. Find the profit-maximizing quantity of L.
   b. Find the profit-maximizing quantity of q.
   c. What is the maximum profit?
   d. Suppose now that the firm is taxed \$30 per unit of output and that the wage rate is subsidized at a rate of \$15 per hour. Assume that the firm is a price taker, so the price of the product remains at \$150. Find the new profit-maximizing levels of L, q, and profit.
   e. Now suppose that the firm is required to pay a 20 percent tax on its profits. Find the new profit-maximizing levels of L, q, and profit.