CHAPTER 16

Employment and Pricing of Inputs

What determines the employment level and prices of inputs used to produce final products?

Chapter Outline

16.1 The Input Demand Curve of a Competitive Firm
The Firm's Demand Curve: One Variable Input *The Firm's Demand Curve: All Inputs Variable* *The Firm's Demand Curve: An Alternative Approach*
16.2 Industry and Market Demand Curves for an Input
A Competitive Industry's Demand Curve for an Input *The Elasticity of an Industry's Demand Curve for an Input*
Application 16.1 Explaining Sky-High Pilot Salaries Under Airline Regulation
The Market Demand Curve for an Input
16.3 The Supply of Inputs
16.4 Industry Determination of Price and Employment of Inputs
Process of Input Price Equalization Across Industries
Application 16.2 The Net Benefits of High-Tech Immigrants
16.5 Input Price Determination in a Multi-Industry Market
16.6 Input Demand and Employment by an Output Market Monopoly
16.7 Monopsony in Input Markets
Application 16.3 Major League Monopsony

Learning Objectives

- Explore the factors influencing the demand for an input by an individual competitive firm.
- Derive the market demand curve for an input by aggregating the demand curves of the various firms interested in hiring the input.
- Investigate the general shape of an input supply curve.
- Show how an input's price and employment level is determined in a multi-industry market.
- Examine input demand and employment by an output market monopoly.
- Define what is meant by monopsony in input markets.

In previous chapters we emphasized factors determining the output and price of final products. In this and the next two chapters we turn our attention to factors determining the employment level and prices of inputs used to produce final products. There are many similarities in the processes of analyzing product markets and input markets, since both involve the interaction of buyers and sellers. The leading actors' roles, however, are reversed. Firms are suppliers in product markets but demanders in input markets. Households and individuals are demanders in product markets and suppliers in input markets.

An examination of input markets helps us answer many interesting and important questions, such as "What determines a country's income level and distribution?", "Why do

women earn less than men and African Americans less than whites?", "Why have the salaries of professional baseball players risen so rapidly during the past three decades?", "Why do entrepreneurs earn more than their non-entrepreneurial counterparts?", "Does immigration make the United States better off?", and "What are the effects of unions, and why have unions in the United States been declining in the private sector but growing in the public sector?"

This chapter discusses the basic principles common to all input market analysis, whether the input is labor, capital, land, or raw materials. We first examine the demand for inputs by competitive firms, then turn to the supply of inputs, and finally bring the two together to complete the general model. The chapter's final two sections analyze input markets under noncompetitive conditions.

16.1 THE INPUT DEMAND CURVE OF A COMPETITIVE FIRM[1]

The market demand curve for an input shows the total quantity of the input that will be purchased at various prices by all demanders as a group. To determine the market demand, we begin with the factors influencing the employment decision of the individual demander, usually a business firm. Then, we aggregate the demands of individual firms to obtain the total or market demand for the input. This treatment is similar to the derivation of the market demand curve for a consumer good. In that case, we first derived the individual consumer's demand curve; then, we combined those curves to obtain the market demand curve.

In this section we learn how a change in an input's price will affect its use by a competitive firm. At the outset, note that there are no new assumptions in the analysis. We use the same competitive model developed in Chapters 8 and 9, but now our focus is on input rather than output markets. For example, we still assume the firm's goal is profit maximization, but now we want to see why profit maximization, along with the competitive model's other assumptions, implies that a firm will employ more of an input when its price is lower.

How does a firm determine how many workers to hire when its goal is to maximize profit? For the moment we assume workers are homogeneous—that is, interchangeable as far as the firm is concerned—so the only question is how many the firm should employ. Each additional worker hired adds to the firm's cost, since the firm must pay the going wage rate. At the same time, each additional worker also adds to the firm's revenue, since a larger work force produces more output. Thus, benefits (greater revenue) and costs (wages) are associated with the firm's employment decision. The firm will increase profit by hiring additional workers as long as the additional revenue generated by the output expansion exceeds the wages paid. A comparison of the marginal benefit of hiring workers, in the form of added revenue, to the marginal cost of hiring workers, in the form of added wage costs, guides the firm's decision of how many workers to employ.

The Firm's Demand Curve: One Variable Input

Imagine a short-run setting where the quantities of nonlabor inputs (such as raw materials and machines) are fixed and only the number of workers can be varied. In this setting the law of diminishing marginal returns applies to labor: beyond some point, each additional worker results in a smaller addition to output. The contribution to output made by increasing the number of workers is an important determinant of the firm's demand for labor, and the marginal product curve (as described in Chapter 7) contains the relevant information.

[1]A mathematical treatment of some of the material in this section is given in the appendix at the back of the book (pages 575–576).

In Figure 16.1, the downward-sloping portion of the marginal product curve for labor is MP_L. The marginal product curve indicates that if the firm employs 20 workers per day, the output produced by an additional worker (MP_L) is three units; if employment increases to 25 workers, the marginal product of labor is lower (in this case two units) due to the law of diminishing marginal returns.

Starting from any given employment level, let's consider how hiring an additional worker affects the firm's total revenue. If 20 workers are employed, one more worker increases final output by 3 units. If the final output sells for a price, P, of \$100 per unit, the additional 3 units of output generated by hiring another worker add \$300 to revenue. Multiplying the marginal product by the price per unit of output ($MP_L \times P$) gives the marginal value product of labor (MVP_L). In general, the **marginal value product** measures *the extra revenue a competitive firm receives by selling the additional output generated when employment of an input is increased by one unit.*

MARGINAL VALUE PRODUCT (MVP) the extra revenue a competitive firm receives by selling the additional output generated when employment of an input is increased by one unit

A downward-sloping marginal value product curve is derived by multiplying the constant price of output (recall that we are dealing with a competitive firm; the price is unchanged when more output is sold) by the declining marginal product of labor. The marginal value product curve coincides with the marginal product curve. The only difference is that now we are measuring the marginal product of labor in terms of what it sells for on the vertical axis. For example, when 20 workers are employed and the price of output is \$100, MP_L is 3 units and MVP_L is \$300; when 25 workers are employed, MP_L falls to 2 units and MVP_L to \$200.

The marginal value product curve is the firm's demand curve for a given input when all other inputs are fixed. To see this relationship, suppose that the daily wage rate is \$300 per worker. The firm can hire as many workers as it wants at this wage rate, so each additional employee adds \$300 to the firm's total cost. Every extra worker, however, also adds an amount equal to the marginal value product to the firm's total revenue. Comparing the cost and revenue effects tells the firm how many workers to hire. For instance, if the firm is currently employing 15 workers, the marginal value product of an additional worker is \$400. Hiring another worker thus adds more to revenue (\$400) than to cost (\$300), so profit increases by employing more workers. The firm should expand employment up to the point where the marginal value product has fallen to \$300, the wage rate. In Figure 16.1, the most profitable employment level is 20 workers when the wage rate is \$300. If the firm hires more than 20 workers, total cost goes up by more than total revenue (the wage rate, \$300, is greater than MVP_L beyond 20 workers), so profit declines.

FIGURE 16.1

A Competitive Firm's Demand for Labor: One Variable Input

With labor as the only variable input, we can convert labor's marginal product curve, MP_L, into the marginal value product curve, MVP_L, by multiplying the marginal product of labor by the price of the commodity produced. The MVP_L curve is the competitive firm's demand curve for labor if other inputs cannot be varied.

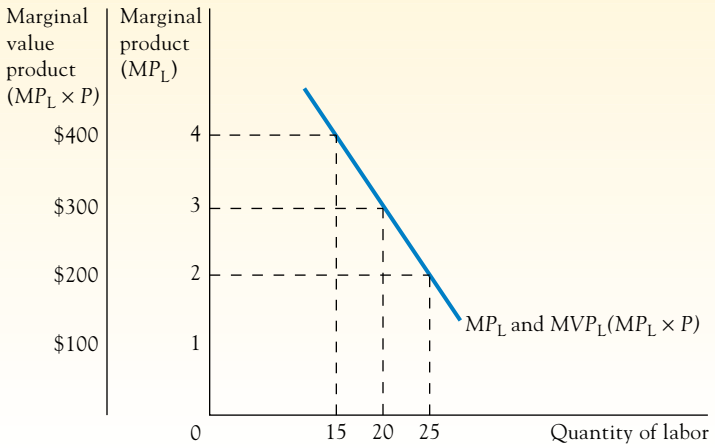

Note that at a lower wage rate, hiring more workers is profitable. For instance, if the wage rate drops to $200 per day, the firm maximizes profit by expanding employment to 25 workers. At the initial employment level of 20 workers and the lower wage rate, the marginal value product of hiring another worker ($300) is now greater than the wage cost ($200), so the firm adds more to revenue than to cost by employing 5 more workers.

Two important conclusions emerge from this analysis. First, the marginal value product curve identifies the most profitable employment level for the input at each alternative cost. The firm will hire up to the point where the input's marginal value product equals its cost. Second, the marginal value product curve—the firm's demand curve when other inputs are not varied—slopes downward. This follows directly from the law of diminishing marginal returns: If an input's marginal product declines as more is employed, so must the marginal value product.

The preceding analysis assumes that the firm is a profit maximizer in a competitive market. It may be helpful to relate this analysis to the discussion in Chapter 9 that emphasized the most profitable output for the competitive firm. We have just seen that the firm maximizes profit by employing an input—in this case, labor—up to the point where its *MVP* equals the cost of the input—in this case, the wage rate w. When profit is at a maximum, therefore, the following condition holds:

$$w = MVP_L. \tag{1}$$

Because MVP_L equals $MP_L \times P$, if we divide both sides of equation (1) by MP_L we obtain:

$$w/MP_L = P. \tag{2}$$

Recall from Chapter 8 that the ratio w/MP_L is equal to the marginal cost (MC) of producing one more unit of output by using additional amounts of labor. Therefore, equation (2) is equivalent to the price-equals-marginal-cost condition for profit maximization in a competitive output market. *When the competitive firm is hiring workers so that* $w = MVP_L$, *then* $MC = P$, *and vice versa*. We have been looking at the same process of profit maximization that we examined in earlier chapters, but now from the perspective of its implications for the employment decisions of the firm.

The Firm's Demand Curve: All Inputs Variable

In identifying a firm's *MVP* curve as its demand curve for an input, we assumed that the quantities of other inputs are fixed. In general, however, a change in an input's price leads a firm to alter its employment of not only that input but also other inputs. For example, a reduction in the cost of computers may lead to the employment of more computer programmers as well as more computers. Consequently, an input demand curve should allow a firm to adjust its use of other inputs as well as the one whose price has changed.

We can easily extend the analysis to allow for variation in the quantities of all inputs. Suppose a competitive firm is initially in equilibrium, employing the appropriate quantities of all inputs. In Figure 16.2, the firm is operating at point A on MVP_L, employing 20 workers when the daily wage rate is $300. Note that the quantity of capital, assumed to be the only other input, is constant at 10 units at all points along MVP_L. Now suppose the wage rate falls to $200. If the quantity of capital is kept constant at 10 units, the firm increases its employment of workers to 25 units, at point B on MVP_L. This increased employment does not represent a complete adjustment to the lower wage rate, since it will normally be in the interest of the firm to expand its employment of capital, too.

An increase in the quantity of capital, though, shifts the MVP_L curve upward. If the quantity of capital increases to 12 units, the MVP_L curve shifts to MVP'_L. With 12 units of capital, each worker has more "tools" to work with than before, so the marginal productivity

FIGURE 16.2

A Competitive Firm's Demand for Labor: All Inputs Variable

When all inputs are variable, an input's *MVP* curve shifts with changes in the employment of other inputs. The firm's labor demand curve is then the *d* curve, which takes into account the way changes in the amount of capital employed affect the *MVP* of labor.

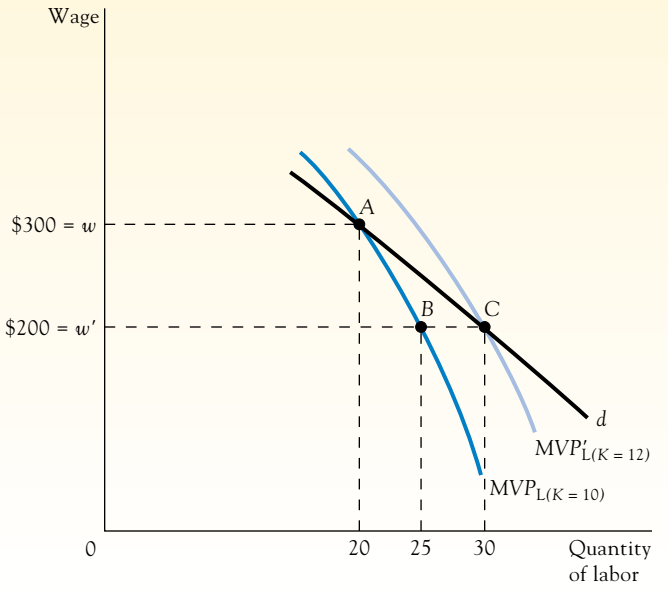

of workers is greater. The greater marginal productivity, coupled with an unchanged product price, implies that the new marginal value product curve has a greater height at each possible employment level than before. The adjustment leads to a further increase in the employment of workers, to point C on the marginal value product curve for a constant 12 units of capital. Thus, the firm's full response to the lower wage rate is an employment increase from 20 to 30 workers.[2]

A and C are two points on the firm's labor demand curve, *d*, when the employment of all inputs can be varied in response to a change in the wage rate. The firm still employs workers up to the point where their marginal value product equals the wage rate, but we have now allowed for the effect of variations in the employment of other inputs on the marginal productivity of labor. We can think of the demand curve, *d*, therefore, as a generalized marginal value product curve. Since this curve allows the firm to vary all inputs, it is the competitive firm's long-run demand curve for an input. Note that in deriving this long-run demand curve, we assume that other inputs' prices are unchanged (only the quantities employed are variable) and that the final product's price is also constant.

The Firm's Demand Curve: An Alternative Approach

We can gain further insight into a firm's adjustment to an input price change with an alternative approach. Our first method has the advantage of linking the demand for an input to its marginal productivity but the disadvantage of obscuring what happens in the output mar-

[2]Our analysis deals with the situation, thought to be typical, where labor and capital are complements. (Two inputs are complements if an increase in the quantity of one leads to an increase in the marginal product of the other.) When the firm uses more labor and the same amount of capital at point *B*, the marginal product curve of capital increases so the firm also expands its use of capital. The demand curve for labor, however, still slopes downward if the two inputs are substitutes. (They are substitutes if an increase in the quantity of one decreases the marginal product of the other.) In that case, when the firm increases labor (from point *A* to point *B*), the marginal product of capital declines and the firm employs less capital. A reduction in capital increases the marginal product of labor when the inputs are substitutes, so the MVP_L curve shifts upward in this case as well, just as it does in Figure 16.2.

ket and the market for other inputs. The approach we develop next also shows that the demand curve for an input slopes downward, but pays more explicit attention to the output market and the demand for other inputs.

Figure 16.3b shows the firm's position, q_1, in its output market where its marginal cost curve MC crosses the horizontal demand curve at point E. Figure 16.3a shows the same initial situation from the perspective of the firm's employment of inputs. The least-cost method of producing output q_1 occurs at the tangency between the IQ_1 isoquant and the isocost line AZ at point E. The firm employs 20 workers and 10 units of capital at point E.

Now let's work through the effects of a reduction in the wage rate. First, how will a lower wage rate affect the firm if we *tentatively* assume it continues to produce the same output? Recall from Chapter 8 that the isocost line's slope equals the ratio of the wage rate to the price of capital. A lower wage rate and an unchanged price of capital imply that isocost lines will be flatter because labor becomes cheaper relative to capital when the wage rate falls. Isocost line $A'Z'$ in Figure 16.3a reflects the lower wage rate, and the least costly way of producing q_1 units of output occurs at E_1, where $A'Z'$ is tangent to isoquant IQ_1. *To produce an unchanged output, the firm uses more labor and less capital when the relative cost of labor falls; that is, the firm substitutes labor for capital.*

In Figure 16.3b, point E_1 on MC′ shows the same adjustment. A wage rate reduction lowers the entire marginal cost curve to MC′ (see Chapter 8 for a more detailed discussion), so when output is q_1, the product's price is greater than its marginal cost. The firm, therefore, has an incentive to expand output as a result of the lower wage rate. Now consider the subsequent effects as the firm expands output from q_1 to the new profit-maximizing level q_2. Figure 16.3a shows this effect as the movement along the new expansion path (based on the lower wage rate and unchanged price of capital) from point E_1 to point E_2. As the firm pro-

FIGURE 16.3

A Competitive Firm's Demand for Labor: All Inputs Variable

(a) A lower wage rate causes the firm to substitute toward labor and away from capital—the move from E to E_1. (b) At a lower wage rate, output expands from q_1 to q_2, as the lower wage rate shifts the marginal cost curve downward. This output effect further increases labor employment—the move from E_1 to E_2 in part (a).

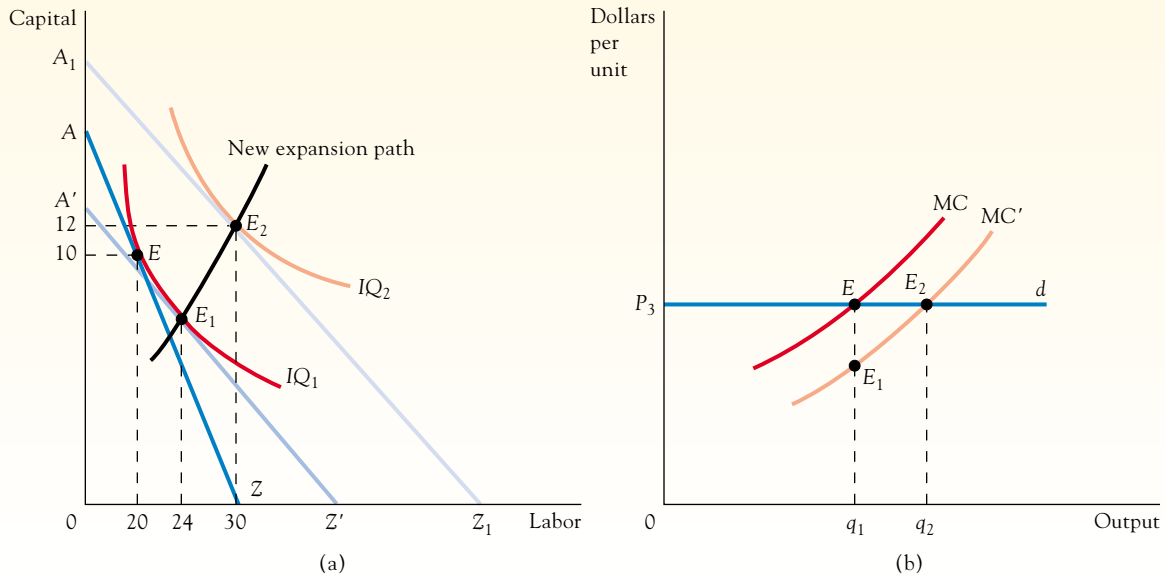

SUBSTITUTION EFFECT OF AN INPUT PRICE CHANGE
the change in input employment when output is held constant and one input is substituted for another in response to an input price change

OUTPUT EFFECT OF AN INPUT PRICE CHANGE
the change in input employment when output is altered in response to a change in the price of an input

duces more output, it moves to a higher isocost line, A_1Z_1, and it employs more of both inputs than it did at point E_1. At E_2 the firm employs 30 workers and 12 units of capital. The wage rate reduction thus increases labor employment from 20 to 30 workers.

This approach to an input's demand curve involves separating the total effect of a price change into two components, similar to what we did with a consumer's demand curve. The change in input employment when output is held constant and one input is substituted for another is called the **substitution effect of an input price change**. The movement along the IQ_1 isoquant from point E to point E_1 (from 20 to 24 workers) shows the substitution effect in our case when the wage rate falls. The change in input employment when output is altered is called the **output effect of an input price change** and is shown by the movement along the new expansion path from point E_1 to point E_2 (from 24 to 30 workers).[3] When summed, these two effects identify a firm's full input employment response to an input price change.

Since the substitution and output effects both imply greater employment at a lower input price, the firm's input demand curve slopes downward. The firm employs more workers at a lower wage rate because it uses more labor per unit of output (the substitution effect) and because it is profitable to increase output when production cost falls (the output effect).

16.2 Industry and Market Demand Curves for an Input

To derive the market demand curve for an input, we aggregate the demand curves of the various firms interested in hiring the input. Since an input such as labor is likely to be demanded by firms in many different industries, we proceed in two steps. First, we need to determine each industry's demand curve for labor. Second, we (horizontally) aggregate each industry's demand curve for labor to obtain the total or market demand curve for it.

A Competitive Industry's Demand Curve for an Input

In determining a particular industry's demand curve for an input such as labor, we must recognize one problem. When deriving the competitive firm's input demand curve, we assumed that the product's price was fixed. Recall, though, that when the wage rate fell, the firm expanded output and sold the larger output at an unchanged price. The assumption of a fixed product price is appropriate when we are dealing with just one firm. But now we are interested in the response of all firms in an industry to a lower wage rate. *When all firms simultaneously increase output, they can sell more total industry output only at a lower price.*

Figure 16.4 illustrates how this factor affects the derivation of an industry's demand curve for an input. In Figure 16.4a, d is the labor demand curve for a single firm, assuming that the price, P, of the final product is \$100 at all points along d. If the initial wage rate is \$300, the firm in Figure 16.4a hires out to point A on its demand curve d and employs 20 workers. With 100 identical firms in the industry interested in hiring labor, total employment is 2,000 workers at the \$300 wage rate, or point A' in Figure 16.4b. Point A' lies on the Σd curve, which is the horizontal summation of the d curves of the firms in the industry.

By assumption, A' in Figure 16.4b is a point on the industry demand curve for labor. The simple summation of the firms' demand curves, Σd, does not, however, show the amount of labor demanded by the industry at other wage rates. To see this, suppose that the wage rate falls to \$200. In Figure 16.4a, individual firms begin expanding employment from point A to point B, and in the process output rises. As industry output rises, the product price falls since consumers will buy the larger output only at a lower price. But a lower product price shifts each firm's labor demand curve downward, since labor's marginal value product curve is

[3]In contrast to a consumer's demand curve, the firm's labor demand curve is not derived by pivoting the isocost line at point A in Figure 16.3a. That approach would be valid only if the firm continues operating at the same total cost, which is generally untrue. The most profitable total cost for the firm to incur depends on the demand and cost conditions shown in Figure 16.3b.

FIGURE 16.4

The Competitive Industry's Demand for Labor
(a) The competitive firm's demand curve for labor, *d*, assumes a given product price. (b) In deriving the industry demand curve for labor, we must take into account that as industry output changes, so will the product price. The industry demand curve for labor is *D*.

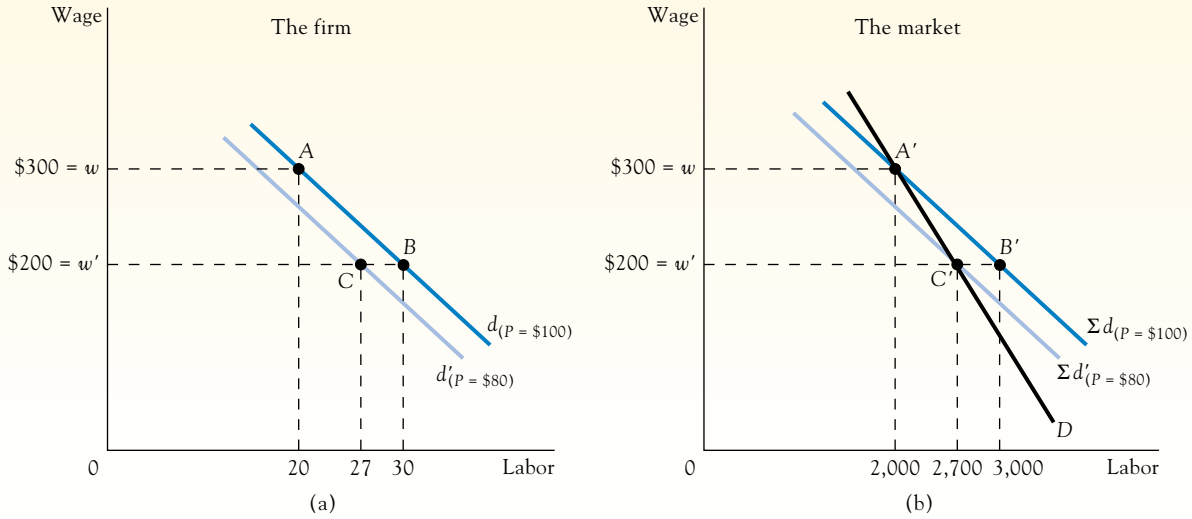

lower when the value (price) of the product is reduced. (Recall that the marginal value product of labor is equal to the marginal product of labor multiplied by the price of the output. Since the output price has fallen, the marginal value product of labor is lower for each level of input use.) If the product price falls to $80 per unit, for example, the firm's demand curve becomes d' in Figure 16.4a, and the firm employs 27 workers. In Figure 16.4b, all the firms together begin by increasing employment from point A' to point B', but the increase is cut short by the falling product price, and they end up at point C'. Point C' is a second point on the industry demand curve for labor, D.

DERIVED DEMAND another name for an industry's input demand curve, reflecting the fact that the industry's demand for an input ultimately derives from consumers' demand for the final product produced by that input

This derivation of an industry's labor demand curve accounts for the effect of increased employment, and hence output, on the product price. The industry demand curve is less elastic than the Σd curve, which is based on a fixed product price but still slopes downward. This relationship becomes easier to see when we recognize that there are still substitution and output effects, implying greater employment at the lower wage rate even when the declining product price is accounted for. The product price falls only because more output is produced, and greater production involves the use of more labor (the output effect). In addition, firms use more labor per unit of output when the wage rate is lower because they substitute labor for other inputs (the substitution effect) at each given output level.

Note that we assume the demand curve for the final product is fixed when deriving an industry's demand curve for an input. In fact, economists often refer to an industry's input demand curve as a **derived demand**: the textile industry's demand for workers, for example, is derived from the demand of consumers for textiles. Firms will pay workers to produce textiles only because consumers are willing to pay for textiles. If the demand curve for textiles shifts, the textile industry's demand curve for workers also shifts. The consumer demand curve for a product is thus an important determinant of the industry demand curve for an input used in the production of the product.

The Elasticity of an Industry's Demand Curve for an Input

The price elasticity of an industry's demand for an input is defined and measured in the same way as a consumer demand curve. The magnitude of the price elasticity of demand for an input can be critical. For example, in evaluating the minimum wage law, whether a 10 percent increase in the legal minimum wage reduces employment of low-wage workers in the fast-food industry by 20 percent (an elasticity of 2.0) or 5 percent (an elasticity of 0.5) makes a big difference.

There are four major determinants of the elasticity of an industry's demand for an input. First, the greater the elasticity of demand for the product produced by the industry, the more elastic the input demand. Recall that an industry's input demand curve is a derived demand curve. If consumers will purchase a great deal more of the good at a slightly lower price (highly elastic product demand), firms in the industry will produce much more when an input price falls, and employment will increase sharply. An elastic product demand gives rise to a large output effect, which, in turn, contributes to the elasticity of the industry's demand for inputs. For instance, consider Figure 16.4b, and suppose the wage rate falls from $300 to $200. If the consumers' demand curve was perfectly elastic, the greater industry output could be sold at an unchanged $100 price. The firms would expand employment to point B', and, in fact, Σd would be the industry's demand curve for labor in this case.

Second, an industry's input demand is more elastic when it is easier to substitute one input for another in production. This condition refers to the technology of production reflected in the curvature of the production isoquants. When it is technically easy to substitute among inputs, the substitution effect of an input price change is large, implying a large (elastic) employment change. For example, if machines can adequately do the work performed by workers and at only a slightly higher cost, a wage increase can lead firms in an industry to switch entirely to machines and reduce employment of workers to zero—implying a highly elastic industry demand for workers.

Third, an industry's demand for an input is more elastic when the supply of other inputs is more elastic. If machine prices rise sharply when firms switch from workers to machines (implying an inelastic supply), only a limited amount of profitable switching can occur, resulting in a small substitution effect and a low elasticity of demand for workers. The output effect reinforces this impact. If machines rise in price as more are used when output increases, the additional output firms can profitably produce as wage rates fall is limited.

Fourth, the longer the time allowed for adjustment, the more elastic an industry's demand for an input becomes. This is true because substitution possibilities among inputs become greater as firms in an industry have more time to alter their usage of various inputs. For example, a rise in the wage rate may mean that replacing workers with machines is profitable. It takes time, however, for machines to be built and installed, so in the short run few workers will be discharged.

Application 16.1 Explaining Sky-High Pilot Salaries Under Airline Regulation

As mentioned in Chapter 10, employee wages in the domestic airline industry were significantly higher when the industry was subject to regulation. For example, veteran union pilots at the major carriers made $150,000 per year in 1978, just prior to airline deregulation. By contrast, firms entering the industry following deregulation paid their top pilots only $45,000 per year.

The higher employee wages during regulation stemmed from the actions of the Civil Aeronautics Board

(CAB). Specifically, the CAB limited competition between airlines on any given route during regulation. Only a few carriers were initially issued licenses to operate on each route and the CAB refused to grant additional licenses to other firms wanting to enter a market.

With limited competition, airlines with operating licenses for particular routes faced more inelastic demand curves for their output. More inelastic demands for final output implied that the derived demands for inputs, such as pilots, also were less price sensitive. In other words, because the cost of pilot salaries could more easily be passed on to an airline's customers in the form of higher air fares, domestic airlines were less price sensitive in their demand for pilots during regulation. Consequently, employee wages in the domestic airline industry were higher under regulation than after deregulation.

The Market Demand Curve for an Input

Once we have derived an industry's demand curve for an input, determining the market demand curve for the input is straightforward. What we need to do is recognize that firms in more than one industry may be vying for the services of a particular input. For example, the automobile industry is not the only buyer of steel or the only employer of engineers. The aerospace and construction industries also may compete for the services of the same inputs.

The market demand curve for an input is determined by (horizontally) aggregating the various industry demand curves for the input. The aggregation is analogous to the manner in which the demand curves of individual consumers are aggregated to obtain the market demand curve for a product. As we will see in a later section, it is the market demand for an input (the total demand of all industries using the input) that interacts with the total supply to all industries to determine input prices.

16.3 THE SUPPLY OF INPUTS

The supply side of input markets deals with the quantities of inputs available at alternative prices. This subject is somewhat complicated because the shape of the supply curve is likely to differ according to the type of input. In this section we make some general observations that are applicable to all inputs; we defer a discussion of specific inputs until the next chapter.

A broad definition of inputs might classify them as either labor, land, or capital. A narrow definition might distinguish between skilled and unskilled workers, land in New York City and Iowa, and buildings and trucks. The appropriate definition depends on the problem. The broad classification serves to make the general points here, but in the next two chapters we will see examples of cases for which it is fruitful to be more specific.

People own the inputs used by firms to produce goods. Our problem is to understand the conditions under which the owners of inputs will offer them for sale or rent. At the outset we should distinguish between the amount of inputs in existence at any given time—the stock of resources—and the amount offered for sale or rent. At any time, a fixed number of people are capable of working. There is a fixed area of land, and a fixed number of buildings, machinery, and other capital equipment. The amount in existence can differ significantly from the amount owners offer for use. Since the amount that owners offer depends on the price they are paid, we must be concerned with the supply curves of inputs, not just the stock of inputs in existence.

The general shape of an input supply curve depends critically on the market for which the supply curve is drawn. Consider the supply curve of labor to all industries in the economy. To simplify, suppose that all workers are identical, so there is only one wage rate. If the wage rate goes up, will the total amount of labor offered increase? We will give a fuller analysis in the next chapter, but here we simply note that the total amount of labor can increase only if workers decide to work longer hours or if more people enter the labor force. Such responses to a higher wage rate may be so small, though, that the supply curve of labor to all industries

together will be approximately vertical, as in Figure 16.5a. A vertical supply curve indicates that an increase in the wage rate from w_1 to w_2 leaves the number of workers unchanged at 100 million. We are not asserting that the supply curve will necessarily be vertical, but it could be; so for the moment let's suppose it is.

Although the supply of labor to all industries taken together may be vertical, this does not imply that the supply curve of labor confronting *any particular industry* is vertical. While the total number of workers employed in the economy may not change, the number employed by a particular industry is subject to great variation. If the wage rate paid to software programmers increases, workers in other industries will leave their jobs to go to work as software programmers. This adjustment doesn't change the total number of workers employed, but it does change employment in the software programming industry. The supply curve of workers to the software programming industry thus slopes upward, as illustrated in Figure 16.5b. An increase in the software programming wage from w_1 to w_2 induces 10,000 workers to move from other jobs into the software programming industry.

Both labor supply curves in Figure 16.5 are correct in the sense that they can both exist simultaneously. Figure 16.5a shows the supply curve of labor *to the entire economy*; Figure 16.5b shows the supply curve of labor *to the software programming industry*. Because the software programming industry is only a small part of the entire (economy-wide) labor market, its labor supply curve is more price sensitive than the supply curve of labor for the economy. Indeed, if its share were as small as the numbers used in Figure 16.5, the software programming industry would likely face a virtually horizontal labor supply curve.

This discussion explains why referring to *the* supply curve of an input is ambiguous. We must always specify that it is the input supply curve to a particular set of demanders. Otherwise, we can fall into the trap of thinking that the supply curve of engineers to the defense contracting industry, for example, is vertical, because in the short run there are only a limited number of trained engineers. In fact, the supply curves of *most* inputs to *most* industries

FIGURE 16.5

The Supply Curve of Labor to the Economy and to a Particular Industry
Distinguishing (a) the supply curve of an input to all industries together from (b) the supply curve of the input to one industry is important. The supply curve to one industry will always be more elastic than the supply curve to the economy as a whole.

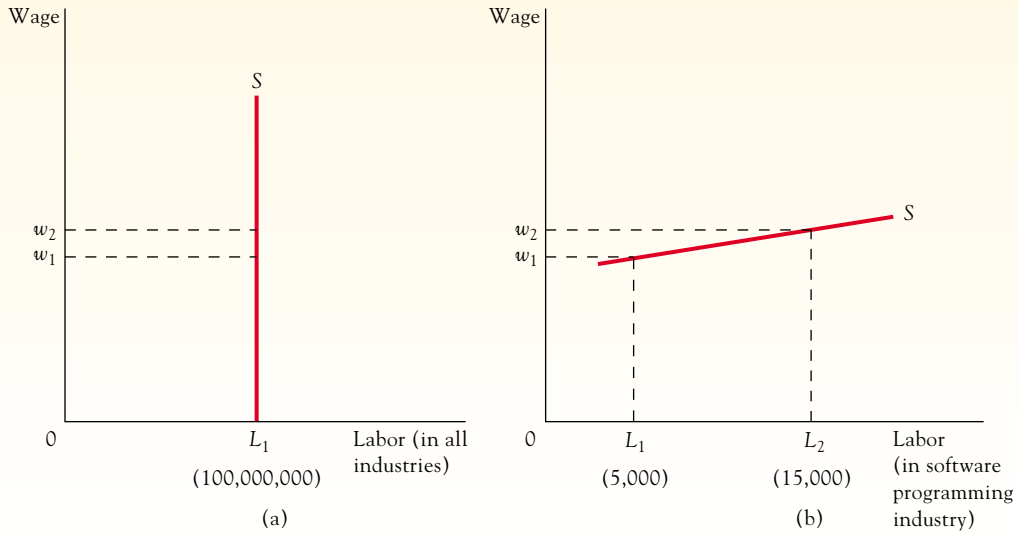

are likely to be upward-sloping, as in Figure 16.5b, regardless of the shape of the supply curve of the input to the economy as a whole, because most industries employ only a small portion of the total amount of any input.

This concept applies to other inputs besides labor. Although the total supply of land to the economy may plausibly be fixed (a vertical supply curve, as in Figure 16.5a), the supply available to any given industry is not. The corn industry can bid land away from other uses if it expands, just as homeowners can bid land away from farmers to build homes on.

Consequently, for individual industries, input supply curves will generally be quite elastic. However, supply curves of inputs to more broadly defined markets will be less elastic, and there are some types of situations where we should use this kind of supply curve. The proper pairing of supply and demand concepts is important to the analysis of input markets, as we will show in detailed examples ahead.

16.4 INDUSTRY DETERMINATION OF PRICE AND EMPLOYMENT OF INPUTS

Firms in an industry compete with one another to acquire inputs, and the industry demand curve for an input summarizes the way an input's price influences the firms' hiring decisions. Input owners provide resources to firms in the industry, and the supply curve of an input to the industry reflects the way the input owners' decisions depend on the price they receive. As with other competitive markets, the interaction of supply and demand determines the equilibrium price and quantity. Figure 16.6 illustrates this process, once again using labor as an example. As before, we assume that all workers are identical.

Figure 16.6b shows a particular industry's demand and supply curves for labor. Perfect competition results in an equilibrium where the industry's firms employ 10,000 workers at a

FIGURE 16.6

The Equilibrium Wage and Employment Level for a Competitive Industry
(a) The position of the firm in equilibrium is shown. Each firm faces a horizontal supply curve at the industry-determined wage rate of $300. (b) Supply and demand in an industry determine the equilibrium employment level and wage rate: 10,000 workers at $300 per day.

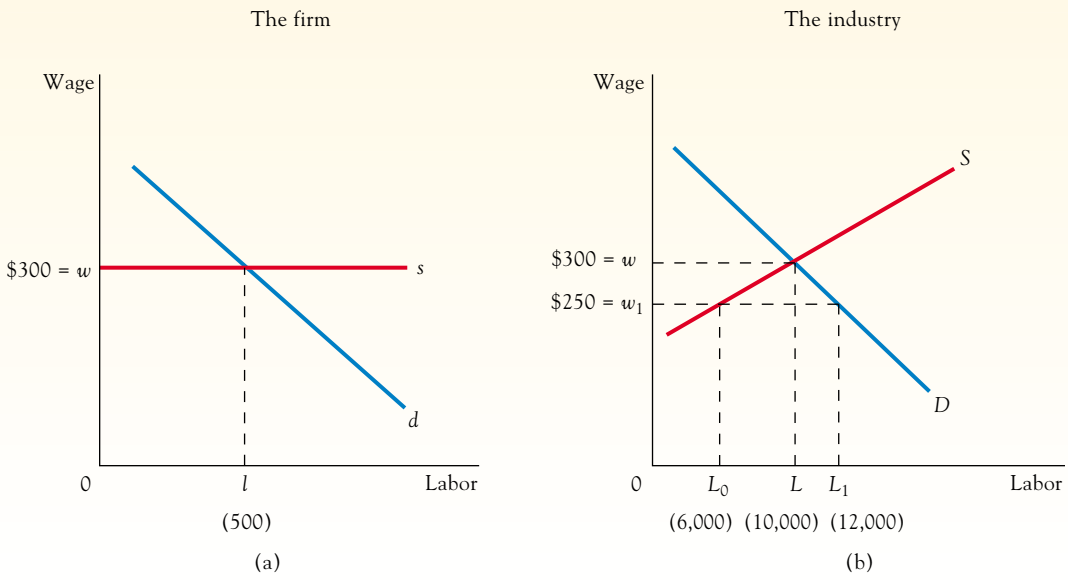

(daily) wage rate of \$300 (w). To understand why the behavior of firms and workers results in this outcome, consider what would happen if some other rate prevailed. Suppose, for example, that the wage rate is \$250 (w_1) instead. At that wage, only 6,000 (L_0) workers agree to work. With such low labor costs, firms as a group find it profitable to employ 12,000 (L_1) workers, but only 6,000 are available. A shortage of labor will exist at the \$250 wage rate, resulting in a tight labor market. Firms will advertise for workers but have too few applicants at the current wage. Workers currently employed by one firm will receive job offers from others at higher wages. Firms unable to recruit workers from within the industry will try to hire workers from other industries, but to do so, they will have to offer higher wages.

As a result, the wage rate will not stay at \$250; it will rise. As firms bid the wage up, workers will quit their jobs in other industries to seize the better opportunity in this market, resulting in an increase in the quantity supplied—a movement up the supply curve. As the wage increases, firms will find that it is no longer profitable to try to fill 12,000 jobs, and the quantity demanded will decrease—a movement up the demand curve. The process will continue until the wage rate reaches the point where the number of workers willing and able to work for firms in the industry equals the number of workers firms are willing to employ. Graphically, the equilibrium is shown by the intersection of the industry demand and supply curves.

When the labor market for a particular industry is in competitive equilibrium, the situation from an individual firm's perspective is depicted in Figure 16.6a. The equilibrium wage is \$300, and each firm faces a horizontal supply curve at that wage. The individual firm is a small part of the total market, so it has no option but to pay the going wage determined in the broader market (illustrated in Figure 16.6b). And at the equilibrium wage it can hire more or fewer workers without appreciably affecting the wage rate. Faced with the \$300 wage, the firm in Figure 16.6a maximizes its profit by hiring 500 workers.

One implication of this analysis is that the market-determined input price equals the marginal value product of the input. Each firm is in the position shown in Figure 16.6a, employing an input quantity for which the marginal value product equals the input price. Recall that an input's marginal value product is the value consumers place on the addition to output made by the input. Thus, in a competitive market, input owners are compensated according to how much value the inputs they supply add to output. This happens because input demands are derived demands. Consumers, in their product purchases, are indirectly expressing how valuable the services of the inputs are to them.

Process of Input Price Equalization Across Industries

Several different industries often employ the same inputs. Most industries, for example, use land, electricity, unskilled labor, and buildings. To understand how the prices of these widely used inputs are determined, we must look beyond the boundaries of a single industry and recognize that there is competition among many industries for these inputs. We emphasize one characteristic of this situation here: the tendency for identical inputs to receive the same price, regardless of the industry employing them.

Suppose that the aerospace and telecommunications industries both employ computer programmers. Suppose also that for some reason wages are higher in the aerospace industry. The wage difference won't persist for long because programmers can move from one industry to another. Programmers in the low-paid telecommunications industry will leave their jobs and seek work in the aerospace industry, where pay is higher. This movement will simultaneously reduce the supply of programmers to the telecommunications industry and increase the supply of programmers to the aerospace industry. With more programmers seeking employment in the aerospace industry, the wage rate there will decline, and the wage rate will rise in the telecommunications industry, where supply has decreased. This process will continue until wages in the two industries are equal, since only then do programmers have no further incentive to change jobs.

Figure 16.7 illustrates the process of input price equalization. Let's assume that the aerospace and telecommunications industries together employ 1,500 programmers: aerospace employs 500 while telecommunications hires 1,000. At their respective levels of employment, the daily wage rate is $400 in the aerospace industry (Figure 16.7a) and $300 in the telecommunications industry (Figure 16.7b). To show why these markets are not in equilibrium and how they will ultimately adjust to an equilibrium, we use a specially defined supply curve. Consider the momentary, or very-short-run, supply curve of programmers that identifies the number of programmers in each industry at a specific time. These supply curves will be vertical; that is, at any given time each industry employs a certain number of programmers. These curves are not the supply curves we use in most applications. We use them here for the specific purpose of illustrating explicitly how the movement of programmers from one industry to another affects both markets.

Initially, the momentary supply curve of programmers in the aerospace industry is SS_A, and in the telecommunications industry it is SS_T. Given the initial allocation of programmers between the industries, the wage rate of the 500 aerospace programmers is $400, and the wage rate of the 1,000 telecommunications workers is $300. Clearly, this result is not an equilibrium since telecommunications programmers have an incentive to quit their jobs and seek employment in the aerospace industry. Shifts in the SS curves show the movement of programmers between industries. The momentary supply curve in the telecommunications industry shifts to the left as programmers leave, and when these programmers seek jobs in the aerospace industry, the momentary supply curve there shifts to the right. The movement of workers decreases the difference in wages between the two industries: wages fall in aerospace and rise in telecommunications. Moreover, the movement of programmers will persist as long as the wage rate is

FIGURE 16.7

Input Price Equalization Across Industries

When several industries employ the same input, the input tends to be allocated so that its price is equalized across industries. If this were not true—if, say, programmers were receiving $400 ($w_A$) in the aerospace industry and $300 ($w_T$) in the telecommunications industry—input owners would have an incentive to shift inputs to industries where pay is higher. This process tends to equalize input prices.

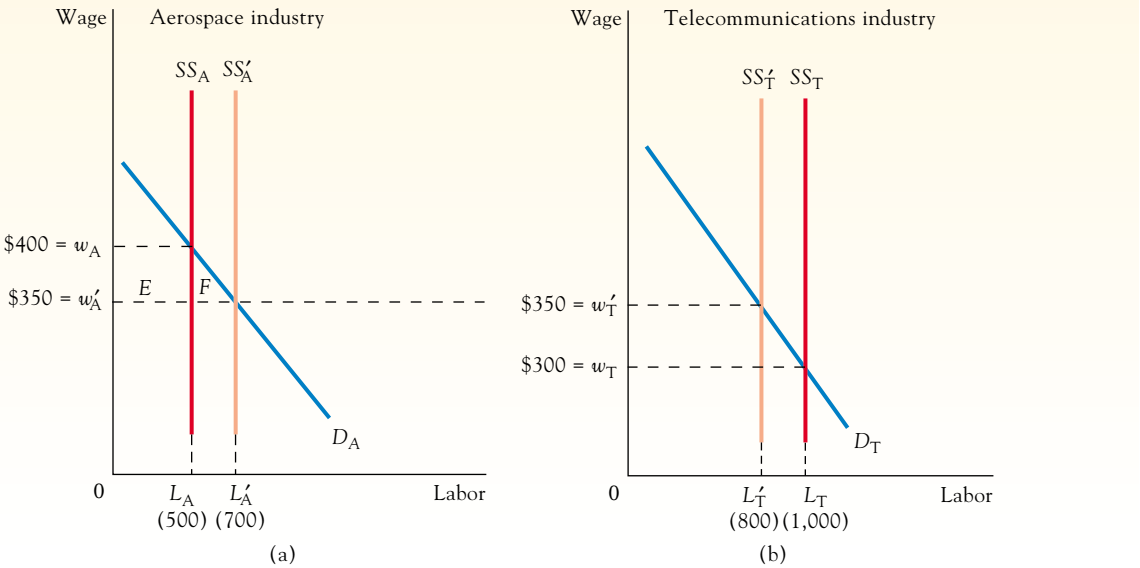

higher in the aerospace industry, meaning it will continue until wages in the two industries are equal. In the diagram, equilibrium occurs when 200 programmers have moved from telecommunications to aerospace, and the common wage rate of $350 is established.

Note that this process works without requiring all programmers to change jobs in search of higher salaries. Relatively few programmers need to relocate to bring wages in the two industries into equilibrium. In our example, only 200 of the 1,500 programmers have to change jobs to produce a uniform wage rate. In fact, in some cases it is unnecessary for the workers to move, geographically speaking, at all. Suppose that the aerospace industry is in California and the telecommunications industry is in Colorado. Suppose also that programmers in Colorado are paid less but are unwilling to relocate to California. Labor immobility of this sort does not forestall the adjustment process. The wage differential creates an incentive for aerospace firms to relocate in Colorado and take advantage of the lower wage rate there, so the impact on wages will be the same as if workers had moved from Colorado to California—a uniform wage will be established.[4]

As this discussion suggests, competitive markets establish uniform input prices across firms, industries, and regions when identically productive inputs are compared. Competitive markets thus promote "equal pay for equal work." This conclusion is true if *equal work* is interpreted to mean equally productive work from the viewpoint of consumers—that is, equal in terms of marginal productivity. The possibility that discrimination by employers leads to differences in wage rates among equally productive workers, an exception to this conclusion, is discussed in Chapter 18.

APPLICATION 16.2 THE NET BENEFITS OF HIGH-TECH IMMIGRANTS

In the latter half of the 1990s, the information-technology sector added more than 1 million net new jobs to the U.S. economy, at wages 75 percent higher than in the rest of the economy.[5]

To fill the need for workers, high-tech companies lobbied for lifting the government restrictions on the number of immigrant skilled workers who could obtain H-1B visas in the United States. In 1999, the number of such visas was nearly doubled to 115,000 per year. In 2000, Congress again raised the immigration limit—this time to 200,000 skilled foreign workers per year.

The net benefit associated with less-restricted immigration of high-tech workers to the United States is illustrated by Figure 16.7a. "Consumers" of labor in the United States are better off by the sum of areas E and F—the difference between the old and the new wage rate (w_A versus w'_A) out to the labor demand curve in the United States.[6] Recall that the height of the demand curve for any product—in this case, labor—reflects the maximum consumers are willing to pay for that product. Due to immigration and the fall in the wage rate, or price of labor, the L_A workers hired in the United States prior to immigration are now cheaper to employ by the amount of the wage rate decline. Consumers of labor in the United States thus benefit from the lower wage rate by rectangular area E on the first L_A workers hired. Moreover, U.S. labor employment rises from L_A to L'_A. On these newly employed workers, consumers of labor realize a net benefit equal to the difference between the height

[5]Laura D'Andrea Tyson, "Open the Gates Wide to High-Skill Immigrants," *Business Week,* July 5, 1999, pp. 16.

[6]The *consumers* of labor who benefit from the lower wage rate need not be the firms hiring the labor. For example, to the extent that product markets are perfectly competitive, the gains from the lower wage rate will be passed on to consumers of the final products produced by the firms hiring the labor.

[4]In some cases, wages can differ between locations because of workers' geographic preferences, but we defer that topic to Chapter 17.

of the demand curve (the marginal value product of the workers) and the cost of hiring the new employees (the wage rate of w'_A). Consequently, a net benefit equal to triangular area F accrues to consumers of U.S. labor on the additional workers hired from L_A to L'_A.

What about the native suppliers of labor represented by the pre-immigration labor supply curve SS_A? They are harmed by immigration since it leads to a decline in the going wage rate and thereby reduces their earnings. The decline in the total earnings accruing to native laborers is represented by rectangular area E—the difference between the old and new wage rate out to the initial labor supply curve SS_A. As might be expected, union organizations representing domestic high-skill workers have been vehemently opposed to any loosening of immigration restrictions by Congress.

On net, however, immigration makes initial U.S. residents better off by triangular area F if Figure 16.7a appropriately reflects the U.S. labor market. The gain to labor consumers from immigration (area E plus area F) outweighs the loss imposed by immigration on native laborers (area E).

16.5 Input Price Determination in a Multi-Industry Market

When several industries compete for the available supply of a particular input, the impact of any one industry on the input's price is likely to be slight, since it usually composes only a small part of the total demand for the input. The broader multi-industry conditions of demand and supply determine the input's price. Now we wish to see how a single industry fits into this broader input market and, in particular, to identify the factors that determine the shape and position of the input supply curve confronting each industry.

Let's consider the hypothetical market for engineers, workers we assume to be identical. Several industries employ engineers. Industry B's demand curve for engineers is D_B in Figure 16.8b, and D_A in Figure 16.8a reflects the demand for engineers by all other industries, *excluding* industry B. Think of industry A as a *group* of industries; each has a demand curve for engineers, and their demands are aggregated as D_A. Therefore, A and B together constitute the total market demand for engineers. The total market demand curve for engineers, the sum of D_A and D_B, is D_T in Figure 16.8c. The market supply curve of engineers to all industries together is S_T, and we have drawn it as upward-sloping on the assumption that higher wage rates will be needed to encourage more people to enter the engineering profession. (Note that we are looking at a time period long enough for people to complete their training. In the short run the market supply curve for engineers will be more inelastic.) *The interaction between the number of people willing and able to work as engineers and the total demand for engineers by all firms and industries determines the wage rate for engineers*. This interaction is shown in Figure 16.8c, with an equilibrium involving employment of 6,000 engineers at a daily wage rate of w, or $350. Each individual industry will then employ the number of engineers it wants at that wage. Industry B will hire 1,000 engineers and industry A will hire 5,000, for a total of 6,000.

Our primary purpose in this section is to explain what determines the shape of the supply curve of engineers to a particular industry—in this example, industry B. We have already identified one point on this supply curve, point F in Figure 16.8b. At F, 1,000 engineers are willing to work in industry B at a wage of $350. Recall that we derive the supply curve of *output* for an industry by assuming a shift in the demand curve and tracing the consequences. We can use the same approach to derive the supply curve of an *input* to an industry—that is, the supply curve of engineers to industry B.

Let's assume that industry B's demand for engineers increases to D'_B, perhaps because of an increase in consumer demand for industry B's output. As a result, the total market demand for engineers increases, but the effect on the market demand is proportionately less because demand has increased in only one market segment. With the market demand rising to D'_T

FIGURE 16.8

Input Price Determination in a Multi-Industry Setting
Total labor demand D_T is the sum of the demands of industries A and B, and it intersects with total supply S_T in part (c) to determine the uniform wage rate. In part (b) the supply curve confronting industry B alone is derived by assuming that D_B increases. Supply curve S_B is highly elastic because industry B is a small part of the total labor market.

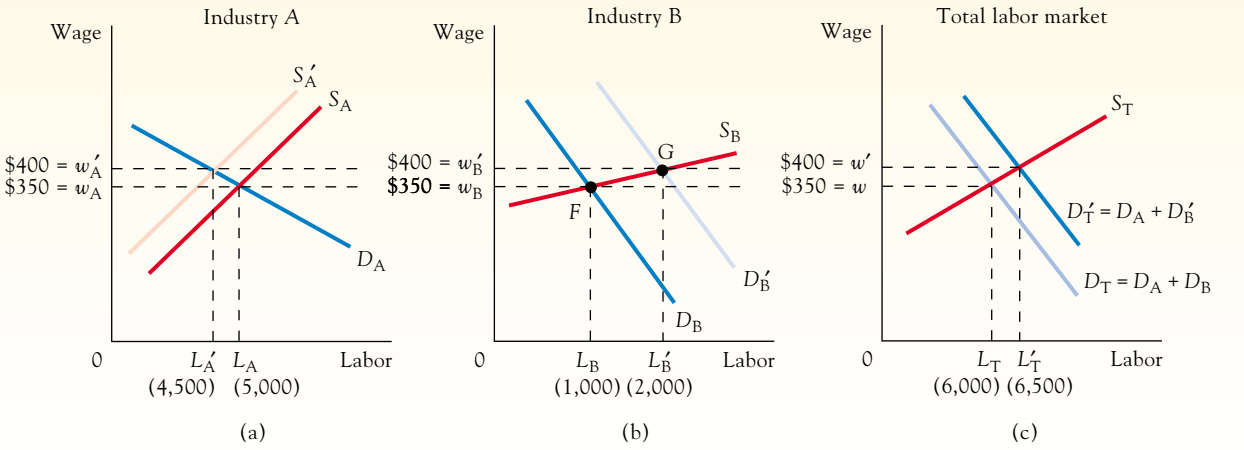

(equal to D_A plus the new demand by B, D'_B), the wage of engineers is bid up to w', and total employment increases to 6,500. At the new wage of $400, industry B hires 2,000 engineers, at point G. Point G is a second point on the supply curve of engineers to industry B; that is, at a wage of $400, 2,000 engineers are willing to work in industry B. In industry A, where the demand curve has not shifted, employment falls to 4,500 when the wage rises to $400.

Note that the additional 1,000 engineers employed in industry B come partly from industry A, where 500 fewer are employed, and partly from an expansion in the total number of engineers from 6,000 to 6,500. In effect, as industry B bids for more engineers by offering higher wages, it attracts some from industry A and also induces some new entrants into the engineering profession. Having derived the supply curve of engineers to industry B, we can understand more easily why the supply curve of an input to a particular industry will normally be highly elastic. In the example given here, when the wage rate in industry B rose from $350 to $400, the number of engineers willing to work there increased from 1,000 to 2,000, implying an elasticity of supply of about 5 (using the formula for arc elasticity). The elasticity of the supply curve of engineers to the total market, however, is only about 0.6. The reason for this difference is straightforward: industry B is only a part of the market for engineers, and it can bid some away from other industries—perhaps only a few from each of dozens of industries—without greatly affecting the general wage level for them all. *The smaller the share of the total market accounted for by an industry, the more elastic its input supply curve*. In our example, industry B initially employed one-sixth of the total number of engineers, but this proportion rose to nearly one-third after the demand increase. In many real-world cases, a single industry will compose a much smaller part of the total market, so its input supply curve can easily be perfectly elastic (horizontal). Recall the significance of high input supply elasticities for the elasticity of the output supply curve, discussed in Chapter 9.

We should also note that the increase in demand by industry B causes the supply curve of engineers to industry A to shift. An input supply curve to a given industry is based on given demand conditions in other industries (in drawing S_B, we assumed D_A was fixed). When other industries compete more aggressively for inputs, industry A will find its workers being

bid away, causing a reduction in input supply to industry A. The result is a higher wage in industry A as well. Remember that input prices will be equalized across industries, so industry A will be unable to retain engineers if it pays less than industry B.

If we were concerned solely with industry B, we would simply need to consider Figure 16.8b. We should, however, integrate the supply of an input to a particular industry into the broader market for the input, a market that usually contains several industries. Indeed, the concept of an industry—a group of firms producing the same product—was designed primarily to study how output markets work. For that purpose, grouping the firms producing the same product makes sense. Relying on the same classification scheme when analyzing input markets is much less helpful since many different industries compete for the same supply of inputs. The notion of a multi-industry input market (Figure 16.8c) is more appropriate.

16.6 INPUT DEMAND AND EMPLOYMENT BY AN OUTPUT MARKET MONOPOLY[7]

A monopoly is defined as a firm that is the sole seller of some product, but a firm that has monopoly power in its output market does not necessarily have market power in its input markets. A firm can be the sole seller of a product and still compete with a large number of firms in hiring inputs. In that case, the firm is a monopoly in its output market and a competitor in its input markets; this situation is the subject of this section's discussion.

Like a competitive firm, a monopoly bases its decisions about input use on the way profit is affected. It expands input employment as long as hiring one more unit adds more to revenue than to cost. The price that must be paid for an input measures the added cost of employing it—just as it did for a competitive firm. The difference in the two market settings rests on the way hiring one more input unit affects the firm's revenue.

For a competitive firm, employing one more input unit adds to revenue an amount equal to the marginal value product. The marginal value product is the additional output produced multiplied by the price at which it can be sold. For a monopoly, one more input unit also adds to revenue by expanding output, but revenue does not increase by the price at which the additional output is sold. Recall that to sell more, a monopoly must reduce the price for all its output; that is, the price received for the hundredth unit of output, for example, is greater than its contribution to revenue. Marginal revenue, which is always lower than price, measures the effect on revenue of selling one more unit of output. Consequently, for a monopoly, the contribution to revenue from employing one more input unit is the additional output (the input's marginal product) multiplied by the marginal revenue associated with the additional output. The product of marginal product and marginal revenue is called the input's **marginal revenue product**.

MARGINAL REVENUE PRODUCT
the product of an input's marginal product and the marginal revenue that can be derived from selling that marginal product

Consider a situation in which all inputs but labor are fixed in quantity for the firm. In Figure 16.9, the marginal value product curve, MVP_L, would be the demand curve for the input under competitive conditions, as we explained in Section 16.1. If the firm is a monopoly, the marginal revenue product curve, MRP_L, is the demand curve for the input. For a monopoly, marginal revenue is below price at each level of output and at each level of labor employment, so the MRP_L curve lies below the MVP_L curve. The monopoly's demand curve for labor, the MRP_L curve, slopes downward for two reasons. First, the marginal product of labor declines as more labor is employed (this relationship also holds true in the competitive case). Second, the marginal revenue associated with selling more output also declines as more labor is employed, since the additional output can be sold only at a lower price.

[7]A mathematical treatment of some of the material in this section is given in the appendix at the back of the book (page 576).

In Figure 16.9, at a wage rate of \$400, the monopoly employs L_1 workers. This amount is the profit-maximizing level of employment. At any lower level of employment the revenue generated by hiring another worker exceeds the cost ($MRP_L > w$), so profit increases if employment increases to L_1, where $MRP_L = w$. If the monopoly employed more workers, they would add more to cost (w) than to revenue (MRP_L), and profit would be lower.

Employment of labor, or any other input, is lower under monopoly than competition. Under competition, L_2 workers are employed at the point where $w = MVP_L$; only L_1 workers are employed under monopoly. This result should come as no surprise, since it was already implicit in our conclusion in Chapter 11 that a monopoly produces less output than a competitive industry. To produce less, it uses fewer inputs. Figure 16.9 therefore depicts the monopolistic reduction in output from the perspective of the input market.

Figure 16.9 also shows an output market monopoly's deadweight loss. The marginal value product of labor measures how much one more worker's output is worth to consumers. At the monopoly outcome, L_1, MVP_L equals \$700, indicating that consumers are willing to pay more for another worker than it costs the monopolist to hire the worker (a wage of \$400). The benefit to consumers of more output is greater than the cost of producing the output, but the monopoly does not hire more workers to expand output. This discussion describes the same deadweight loss due to monopoly explained in Chapter 11, but focuses on the input side of the picture.

We can extend the analysis of the monopoly demand for an input to a case where all inputs are variable in the same manner as we did for the competitive firm. No significant new conclusions emerge, and two important points remain. First, the input demand curves of an output monopoly slope downward, both because the input's marginal productivity declines and because marginal revenue from selling output declines as more of any input is consumed. Second, the input demand curves of a monopoly are lower than they would be if the output market were competitive.

Finally, it is important to note that our focus in this section has been on the input demand curve of a firm that is a monopoly in its output market. *An output market monopoly is not the same as an input market monopoly*. The latter involves a single seller of an input confronting the entire market demand curve for the input (where the market demand curve for an input is derived by aggregating the demands for the input across various firms operating in either competitive or monopoly output markets). We address the topic of input market monopoly in the next chapter in the context of labor unions.

FIGURE 16.9

An Output Monopolist's Demand for an Input
When only one input is variable, we derive a monopoly's input demand curve by multiplying the input's marginal product by the marginal revenue, *MR*, from selling the commodity produced: $MRP_L = MP_L \times MR$. The marginal revenue product curve MRP_L is the monopoly's demand curve, and it lies below the competitive demand curve, MVP_L.

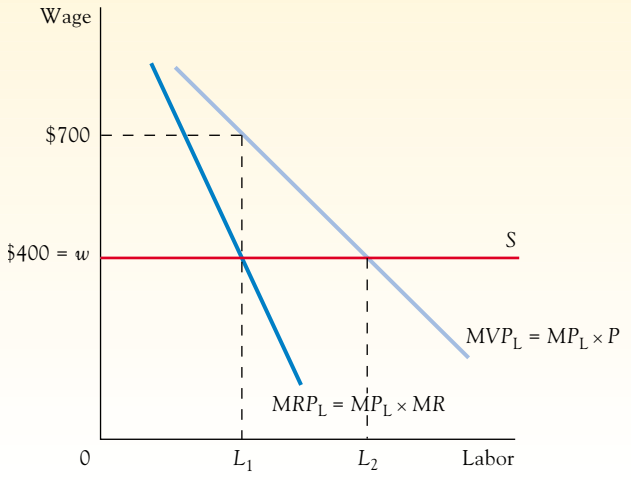

16.7 MONOPSONY IN INPUT MARKETS

MONOPSONY
an input market in which a firm is the sole purchaser of an input

Monopsony means "single buyer." Pure monopsony in input markets occurs when a firm is the sole purchaser of an input. An example of pure monopsony is purchases by General Motors of an automobile part that has been uniquely tailored by suppliers for GM cars. As a sole buyer, a monopsony faces the market supply curve of the input, a curve that is often upward-sloping. An upward-sloping supply curve means that the monopsonist has market power in the input market and can reduce the price paid without losing all the input.

An input market monopsony is analogous to an output market monopoly. An output market monopoly has some discretion over its product's price (as determined by the downward-sloping demand curve), while an input market monopsony has some discretion over the input's price (as determined by the input's upward-sloping supply curve).

Graphically, an upward-sloping supply curve for an input confronting the firm indicates the presence of monopsony. We'll use labor once more as an example. An upward-sloping supply curve means that the firm must pay a higher wage rate to increase the number of workers it employs. (Up until now, we have assumed that every firm faces a horizontal supply curve for each input.) When the firm faces an upward-sloping input supply curve, **marginal input cost** is not the same as **average input cost**. For example, suppose that a firm employs 10 workers at a wage of $300, but to employ 11 workers, the firm must pay a wage rate of $310 to them all. The marginal cost of hiring the eleventh worker is therefore $410, because total labor cost rises from $3,000 to $3,410 when employment increases by one worker. Put differently, we can think of the wage rate as the average cost of labor (AC_L). When the average cost rises as more workers are employed (as is the case with an upward-sloping labor supply curve), the marginal cost of labor (MC_L) must be greater than average cost. In this case the $310 wage rate is equal to the average cost of labor, or the total wage bill divided by the number of workers, and the marginal cost is $410, or the additional cost associated with hiring an extra worker.

MARGINAL INPUT COST
the cost of using an additional unit of an input

AVERAGE INPUT COST
the total cost of an input divided by the units of that input used by a firm

We can identify the profit-maximizing employment level of a monopsony by comparing the impact on total revenue from a change in employment with the effect on total cost, just as we did before. Now, however, the effect on cost is determined by the marginal cost of labor, which is not equal to the wage rate. Figure 16.10 illustrates how we determine the profit-

FIGURE 16.10

An Input Market Monopsony
An input market monopsony faces an upward-sloping input supply curve, so the marginal cost of employing the input is greater than the input price (the average cost of the input): MC_L lies above $S = AC_L$. The intersection of MC_L and D determines employment, but the wage rate is determined by the height of the supply curve.

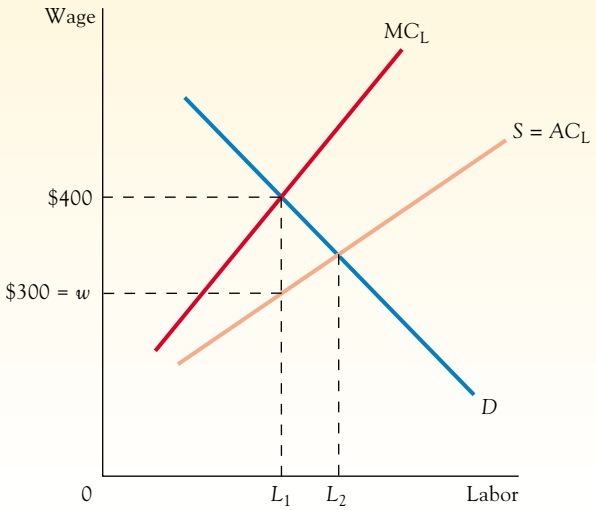

maximizing level of employment. The demand curve of the firm indicates the amount that hiring an additional worker adds to revenue. (The demand curve will be the MVP_L curve if the firm is a competitor in its output market; it will be the MRP_L curve if the firm is a monopoly in its output market. Conceivably, a firm could be competitive in its output market and be a monopsony in its labor market, although this situation is unlikely.) The supply curve to the firm slopes upward—the graphical characteristic of monopsony—so the marginal cost of labor curve, MC_L, lies above the supply curve. Note that the average cost of labor curve, AC_L, is identical to the supply curve if the firm can select any employment level along the labor supply curve. But once the level is selected, each worker hired is paid the same wage. For example, if each worker is paid \$300 per day, the average daily wage per worker is also \$300.

The firm maximizes profit by employing L_1 workers. At that point the marginal cost of labor equals the addition to total revenue from hiring one more worker. The firm, however, does not pay an amount equal to the worker's marginal contribution to revenue (\$400). Instead, the firm pays a wage of \$300. The intersection of MC_L and the demand curve determines the most profitable employment level, while the wage rate is determined by the height of the supply curve at the corresponding level of employment.

In comparison with competitive input market conditions, employment is lower under monopsony and so is the wage rate. If there were competition in the input market in Figure 16.10, employment would be L_2, and the wage would be higher, since the supply curve slopes upward. A similarity between monopsony and monopoly becomes apparent. An output market monopoly restricts output to obtain a higher price; an input market monopsony restricts employment to pay a lower wage. An output market monopoly is able to charge a higher price because it faces a downward-sloping demand curve for its product; an input market monopsony is able to pay a lower wage because it faces an upward-sloping input supply curve.

APPLICATION 16.3 MAJOR LEAGUE MONOPSONY

Until arbitration led to its dismantling in 1975, a provision in major league baseball known as the *reserve clause* limited players to negotiating salary with the first team to sign them.[8] That is, unless the first team to offer a player a contract terminated or traded him, he was at the mercy of the team's owners. The reserve clause thus effectively tied each player to one team and eliminated competition between teams for individual players.

As the monopsony model predicts, when the reserve clause was in effect, players' salaries were lower than their contributions to total revenues (as measured by gate receipts, broadcast fees, concession sales, and so on). In the late 1960s, for example, superstar hitters received an average salary of \$68,000 per season but contributed \$384,000 to their teams' total revenues (net of any training and transportation costs). Superstar pitchers averaged annual salaries of \$86,000, but added \$480,000 to their teams' total revenues.

The dismantling of the reserve clause in 1975 and resulting competition between teams over individual players has caused a fly-up in players' salaries. As of 1976, players who have served a team for five years are eligible to become free agents and sell their skills to the highest bidder. Average annual salaries shot up by 175 percent between 1975 and 1980, from \$53,000 to \$146,000. By 2002, the average salary of a Major League baseball player had risen to \$2,380,000.

[8]Susan Lee, "The Baseball Strike Is a Monopolists' Slugfest," *Wall Street Journal*, June 30, 1981, p. 15; and Gerald W. Scully, "Pay and Performance in Major League Baseball," *American Economic Review*, 64 No. 6 (December 1974), pp. 915–930.

SUMMARY

- A competitive firm's demand for any input slopes downward. This can be shown either by focusing on the marginal value product curve or by considering the substitution and output effects of a change in the input's price.
- A competitive industry's demand curve for an input is derived by aggregating firms' demand curves for the input, but this summation is not a simple one, since a product's price changes as total industry employment and thus output vary.
- In terms of input supply, the demander or set of demanders needs to be specified. Firms that are perfect competitors in input markets face horizontal input supply curves.
- An input's supply curve to an industry may be either horizontal or upward-sloping.
- The supply curve to the economy as a whole may be very inelastic.
- In competitive markets the interaction of supply and demand determines input prices and employment, although the relevant market is frequently broader than a single industry.
- When many industries employ the same input, the input price tends to be equalized across them, and the supply curve of the input to any single industry will be very elastic; no single industry will have much effect on the input price.
- Firms that are monopolies in output markets also have downward-sloping input demand curves, termed marginal revenue product curves.
- Compared with competitive firms, output market monopolies have lower input demand curves. This outcome does not necessarily imply, however, that an output market monopoly will pay lower prices for inputs, since it may face a horizontal input supply curve and thus have to pay the market price for an input.
- An input market monopsony is the sole employer of the input, and so it faces the market supply curve. If that supply curve slopes upward, the monopsony's employment decision affects the price of the input: The monopsonist has market power in the input market.
- Compared to the competitive outcome, the price and employment of an input will be lower with input market monopsony.

REVIEW QUESTIONS AND PROBLEMS

Questions and problems marked with an asterisk have solutions given in Answers to Selected Problems at the back of the book (page 583).

16.1. The data in Table 7.1 relate total output to the amount of labor employed when the amounts of other inputs such as capital are held constant. Use it to answer the following: If the price of final output is $10 and the wage rate is $120, how many workers will be hired? (Assume that the output market is competitive.) Illustrate your answer with a graph similar to Figure 16.1.

16.2. What factors (other than the wage rate) affect the amount of labor a firm that operates in perfectly competitive output markets will hire? How will a change in each of these factors affect the firm's demand curve for labor?

16.3. "The law of demand does not apply to professional baseball players. Since each team already has the maximum number of players allowed on its squad, a reduction in the wage rate that must be paid for baseball players would not lead to any more being hired." Evaluate this statement.

***16.4.** If the demand for personal computers rises, the productivity of workers engaged in making personal computers does not necessarily increase. Why, then, does the demand for such workers increase?

16.5. Distinguish between the short-run and the long-run supply curves of geologists to the domestic economy. Which curve will be more inelastic? Why?

16.6. Rank the following labor supply curves in terms of their elasticities. How does your answer depend on whether you consider short-run or long-run supply curves? Explain your answer.

a. The supply of economists to the federal government.
b. The supply of taxi drivers to Chicago.
c. The supply of college professors to Ohio State University.

16.7. Discuss the determination of equilibrium input price and employment by a competitive industry. Concerning the equilibrium, firms would prefer to pay less for an input; why don't they? Input owners would prefer to receive a higher price; why don't they refuse to supply the input unless the price is higher?

***16.8.** "Employers set wage rates equal to marginal value products." True or false? Explain.

***16.9.** "If the supply of labor increases and depresses wages in a competitive industry, this outcome will benefit firms at the expense of workers." True or false? Explain.

***16.10.** Writing about the nineteenth century, C. Vann Woodward observes: "There was nothing but the urging of con-

science and the weak protest of labor to keep employers from cutting costs at the expense of their workers." Analyze this statement. Was it conscience that kept wages from being zero?

16.11. "College teachers are no more productive today than they were 50 years ago, yet they are paid three times as much today. They are obviously not being paid according to their marginal productivity." Discuss.

***16.12.** If the demand for automobiles rises sharply, how will the price of refrigerators be affected? (Steel is an important input in the production of both products.)

16.13. "Recently, the demand for DVD players has increased rapidly, while the demand for radios has hardly budged. Therefore, the fact that workers are better paid in the DVD player industry is not surprising." Would you be surprised if the assertion about wages turned out to be correct? Support your position with a graphical analysis.

16.14. "If among many good-hearted employers there is one determined to exploit workers, the actions of the single employer may suffice to neutralize the good-heartedness of the rest. This is because if a single employer succeeds in paying lower wages, his fellow-employers may have no alternative but to follow suit, or to see themselves undersold in the product market." Why is this statement wrong?

16.15. "If Mexicans are allowed to immigrate to the United States, they will take jobs away from U.S. citizens." Evaluate this statement.

16.16. Explain why an output market monopoly will employ more of an input when its price is lower in terms of the substitution and output effects of the lower input price. (*Hint:* How would Figures 16.3a and 16.3b be different for a monopoly?)

16.17. Do output market monopolies cause unemployment?

16.18. What is a monopsony? Graphically, what distinguishes a monopsony from a competitive employer of inputs? What does this difference imply for the relative levels of employment and input prices under monopsony versus competition?

16.19. "Along a downward-sloping competitive industry's demand curve for labor, such as the one depicted in Figure 16.4, firm profits will be greater the lower the wage rate is." Explain why this statement is true, false, or uncertain.

16.20. Workers belonging to the Malevolent Association of Microeconomics Teaching Assistants Union at your school are disgruntled over their salaries and would like to request a $100 increase in their weekly wage. Is their request more or less likely to be honored if alternative means of assisting faculty in the instruction of economics, such as study guides, videotapes, and robots, become more plentiful? If the demand for the undergraduate education provided by your school becomes more price elastic due to greater competition from other schools?

16.21. Bad Breath, Inc., sells its output at $1 per unit into competitive markets. Bad Breath's factory is the only employer of labor in Gilroy, California (garlic capital of the world). It faces a supply from competitive workers of $Q_L = w$ where Q_L is the number of workers hired per year and w is the annual wage. Each additional worker hired adds one less unit of output than was added by the previous worker. The 30,000th worker adds nothing to total output. Bad Breath must pay all workers the same wage and, because it has to raise wages to get more labor, each additional worker costs the company $2Q_L$ dollars per year. To maximize profit, how much labor should Bad Breath hire and what wage should it pay? Does efficiency prevail in the Gilroy labor market? If not, what is the size of the deadweight loss?

16.22. If nominal wages in the South are less than in the North, does this imply that the economic theory of labor mobility is invalid? Explain why or why not.

16.23. Per capita income is 600 percent higher in the United States than in Mexico. No other two countries sharing a border have a wider disparity in income levels. Explain why this leads to immigration of workers from Mexico to the United States. Is such immigration beneficial, on net, to the United States?

16.24. Due to OPEC, the price of jet fuel increased by over 700 percent between 1970 and 1980. Explain why airlines responded hardly at all to the increase in the price of jet fuel in the short run (an estimated demand elasticity of between 0 and 0.15). Also explain why the long run response was more substantial and involved substituting Boeing 737s and 767s for 727s and 747s and investing in superior tracking systems whereby a take-off can be delayed until a landing slot is assured at the destination city.

16.25. Studies find that controlling for other factors, university professors' earnings tend to decline with experience. In other words, the more seniority a faculty member has with a particular institution, the lower his or her salary after accounting for other factors such as productivity, degree, and field of study. Explain why this phenomenon may reflect monopsony power being exercised by universities because of the greater moving costs faced by more senior faculty members.

16.26. Suppose that firms in an industry use two inputs, labor and capital. If the price of labor increases, then the firms will demand less labor and more capital. True, false, or uncertain? Explain your answer.

16.27. Economists Ross Eckert and Richard Leftwich have noted that in the early 1950s over 60 percent of MBA graduates from leading business schools took their first jobs in manufacturing and 10 percent in investment banking and consulting. Nowadays, no more than 20 percent of MBA graduates from leading business schools take their first jobs in manufacturing and over 60 percent in investment banking and consulting. Between the 1950s and today, furthermore, MBA starting salaries in investment banking and consulting have risen dramatically relative to starting salaries in manufacturing. Using a graphical analysis, explain why this phenomenon is related to a decline in demand for manufactured products and an increase in demand for services over the last half century.