CHAPTER 4

Labor Demand Elasticities

In 1995, a heated debate broke out among economists and policymakers about the employment effects of minimum wage laws. Clearly, the standard theory developed in chapter 3 predicts that if wages are raised above their market level by a minimum wage law, employment opportunities will be reduced as firms move up (and to the left) along their labor demand curves. Two prominent labor economists, however, after reviewing previous work on the subject and doing new studies of their own, published a 1995 book in which they concluded that the predicted job losses associated with increases in the minimum wage simply could not be observed to occur, at least with any regularity.[1]

The book triggered a highly charged discussion of a long-standing question: just how responsive is employment demand to given changes in wages?[2] Hardly anyone doubts that jobs would be lost if mandated wage increases were huge, but how many are lost with modest increases?

[1]David Card and Alan B. Krueger, *Myth and Measurement*: *The New Economics of the Minimum Wage* (Princeton, N.J.: Princeton University Press, 1995).

[2]Six reviews of Card and Krueger, *Myth and Measurement*, appear in the book review section of the July 1995 issue of *Industrial and Labor Relations Review* 48, no. 4. More recent reviews of findings can be found in Richard V. Burkhauser, Kenneth A. Couch, and David C. Wittenburg, "A Reassessment of the New Economics of the Minimum Wage Literature with Monthly Data from the Current Population Survey," *Journal of Labor Economics* 18 (October 2000): 653–680; and David Neumark and William Wascher, "Minimum Wages and Employment: A Review of Evidence from the New Minimum Wage Research," working paper no. 12663, National Bureau of Economic Research (Cambridge, Mass., January 2007).

The focus of this chapter is on the degree to which employment responds to changes in wages. The responsiveness of labor demand to a change in wage rates is normally measured as an *elasticity*, which in the case of labor demand is the percentage change in employment brought about by a 1 percent change in wages. We begin our analysis by defining, analyzing, and measuring *own-wage* and *cross-wage* elasticities. We then apply these concepts to analyses of minimum wage laws and the employment effects of technological innovations.

The Own-Wage Elasticity of Demand

The *own-wage elasticity of demand* for a category of labor is defined as the percentage change in its employment (E) induced by a 1 percent increase in its wage rate (W):

$$\eta_{ii} = \frac{\%\Delta E_i}{\%\Delta W_i} \tag{4.1}$$

In equation (4.1), we have used the subscript i to denote category of labor i, the Greek letter η (eta) to represent elasticity, and the notation %Δ to represent "percentage change in." Since the previous chapter showed that labor demand curves slope downward, an increase in the wage rate will cause employment to decrease; the own-wage elasticity of demand is therefore a negative number. What is at issue is its magnitude. The larger its *absolute value* (its magnitude, ignoring its sign), the larger the percentage decline in employment associated with any given percentage increase in wages.

Labor economists often focus on whether the absolute value of the elasticity of demand for labor is greater than or less than 1. If it is greater than 1, a 1 percent increase in wages will lead to an employment decline of greater than 1 percent; this situation is referred to as an *elastic* demand curve. In contrast, if the absolute value is less than 1, the demand curve is said to be *inelastic*: a 1 percent increase in wages will lead to a proportionately smaller decline in employment. If demand is elastic, aggregate earnings (defined here as the wage rate times the employment level) of individuals in the category will decline when the wage rate increases, because employment falls at a faster rate than wages rise. Conversely, if demand is inelastic, aggregate earnings will increase when the wage rate is increased. If the elasticity just equals −1, the demand curve is said to be *unitary elastic*, and aggregate earnings will remain unchanged if wages increase.

Figure 4.1 shows that the flatter of the two demand curves graphed (D_1) has greater elasticity than the steeper (D_2). Beginning with any wage (W, for example), a given wage change (to W', say) will yield greater responses in employment with demand curve D_1 than with D_2. To judge the different elasticities of response brought about by the same percentage wage increase, compare $(E_1 - E'_1)/E_1$ with $(E_2 - E'_2)/E_2$. Clearly, the more elastic response occurs along D_1.

To speak of a demand curve as having "an" elasticity, however, is technically incorrect. Given demand curves will generally have elastic and inelastic ranges, and while we are usually interested only in the elasticity of demand in the range

Figure 4.1

Relative Demand Elasticities

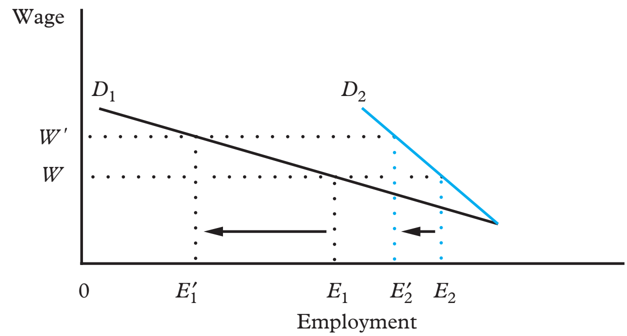

around the current wage rate in any market, we cannot fully understand elasticity without comprehending that it can vary along a given demand curve.

To illustrate, suppose we examine the typical straight-line demand curve that we have used so often in chapters 2 and 3 (see Figure 4.2). One feature of a straight-line demand curve is that at *each* point along the curve, a unit change in wages induces the *same* response in terms of units of employment. For example, at any point along the demand curve shown in Figure 4.2, a $2 decrease in wages will increase employment by 10 workers.

However, the same responses in terms of *unit* changes along the demand curve do *not* imply equal *percentage* changes. To see this point, look first at the upper end of the demand curve in Figure 4.2 (the end where wages are high

Figure 4.2

Different Elasticities along a Demand Curve

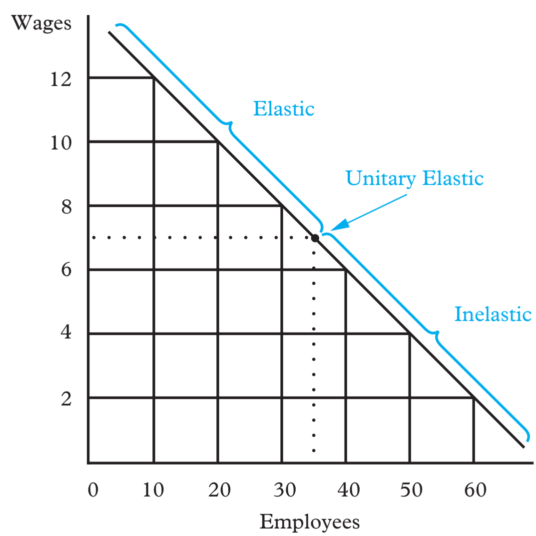

and employment is low). A $2 decrease in wages when the base is $12 represents a 17 percent reduction in wages, while an addition of 10 workers when the starting point is also 10 represents a 100 percent increase in demand. Demand at this point is clearly *elastic*. However, if we look at the same unit changes in the lower region of the demand curve (low wages, high employment), demand there is inelastic. A $2 reduction in wages from a $4 base is a 50 percent reduction, while an increase of 10 workers from a base of 50 is only a 20 percent increase. Since the percentage increase in employment is smaller than the percentage decrease in wages, demand is seen to be inelastic at this end of the curve.

Thus, the upper end of a straight-line demand curve will exhibit greater elasticity than the lower end. Moreover, a straight-line demand curve will actually be elastic in some ranges and inelastic in others (as shown in Figure 4.2).

The Hicks–Marshall Laws of Derived Demand

The factors that influence own-wage elasticity can be summarized by the Hicks–Marshall laws of derived demand—four laws named after two distinguished British economists, John Hicks and Alfred Marshall, who are closely associated with their development.[3] These laws assert that, other things equal, the own-wage elasticity of demand for a category of labor is high under the following conditions:

1. When the price elasticity of demand for the product being produced is high.
2. When other factors of production can be easily substituted for the category of labor.
3. When the supply of other factors of production is highly elastic (that is, usage of other factors of production can be increased without substantially increasing their prices).
4. When the cost of employing the category of labor is a large share of the total costs of production.

Not only are these laws generally valid as an empirical proposition, but the first three can be shown to always hold. There are conditions, however, under which the final law does not hold.

In seeking to explain why these laws hold, it is useful to act as if we could divide the process by which an increase in the wage rate affects the demand for labor into two steps. First, an increase in the wage rate increases the relative cost of the category of labor in question and induces employers to use less of it and more of other inputs (the *substitution effect*). Second, when the wage increase causes the marginal costs of production to rise, there are pressures to

[3]John R. Hicks, *The Theory of Wages*, 2nd ed. (New York: St. Martin's Press, 1966): 241–247; and Alfred Marshall, *Principles of Economics*, 8th ed. (London: Macmillan, 1923): 518–538.

increase product prices and reduce output, causing a fall in employment (the *scale effect*). The four laws of derived demand each deal with substitution or scale effects.

Demand for the Final Product We noted above that wage increases cause production costs to rise and tend to result in product price increases. The greater the price elasticity of demand for the final product, the larger the percentage decline in output associated with a given percentage increase in price—and the greater the percentage decrease in output, the greater the percentage loss in employment (other things equal). Thus, *the greater the elasticity of demand for the product*, *the greater the elasticity of demand for labor*.

One implication of this first law is that, other things equal, the demand for labor at the *firm* level will be more elastic than the demand for labor at the *industry*, or market, level. For example, the product demand curves facing *individual* carpet-manufacturing companies are highly elastic because the carpet of company X is a very close substitute for the carpet of company Y. Compared with price increases at the *firm* level, however, price increases at the *industry* level will not have as large an effect on demand because the closest substitutes for carpeting are hardwood, ceramic, or some kind of vinyl floor covering—none a very close substitute for carpeting. (For the same reasons, the labor demand curve for a monopolist is less elastic than for an individual *firm* in a competitive industry. Monopolists, after all, face *market* demand curves for their product because they are the only seller in the particular market.)

Another implication of this first law is that *wage elasticities will be higher in the long run than in the short run*. The reason for this is that price elasticities of demand in product markets are higher in the long run. In the short run, there may be no good substitutes for a product or consumers may be locked into their current stock of consumer durables. After a period of time, however, new products that are substitutes may be introduced, and consumers will begin to replace durables that have worn out.

Substitutability of Other Factors As the wage rate of a category of labor increases, firms have an incentive to try to substitute other, now relatively cheaper, inputs for the category. Suppose, however, that there were no substitution possibilities; a given number of units of the type of labor *must* be used to produce one unit of output. In this case, there is no reduction in employment due to the substitution effect. In contrast, when substitution possibilities do present themselves, a reduction in employment owing to the substitution effect will accompany whatever reductions are caused by the scale effect. Hence, other things equal, *the easier it is to substitute other factors of production*, *the greater the wage elasticity of labor demand*.

Limitations on substitution possibilities need not be solely technical ones. For example, as we shall see in chapter 13, unions often try to limit substitution possibilities by including specific work rules in their contracts (e.g., minimum crew size for railroad locomotives). Alternatively, the government may legislate limitations by specifying minimum employment levels for safety reasons (e.g., each public swimming pool in New York State must always have a lifeguard present). Such restrictions make the demand for labor less elastic, but substitution possibilities that are not feasible in the short run may well become feasible over longer periods of time. For example, if the wages of railroad workers went up, companies could buy more powerful locomotives and operate with larger trains and fewer locomotives. Likewise, if the wages of lifeguards rose, cities might build larger, but fewer, swimming pools. Both adjustments would occur only in the long run, which is another reason the demand for labor is more elastic in the long run than in the short run.

The Supply of Other Factors Suppose that, as the wage rate increased and employers attempted to substitute other factors of production for labor, the prices of these other factors were bid up. This situation might occur, for example, if we were trying to substitute capital equipment for labor. If producers of capital equipment were already operating their plants near capacity, so that taking on new orders would cause them substantial increases in costs because they would have to work their employees overtime and pay them a wage premium, they would accept new orders only if they could charge a higher price for their equipment. Such a price increase would dampen firms' "appetites" for capital and thus limit the substitution of capital for labor.

For another example, suppose an increase in the wages of unskilled workers caused employers to attempt to substitute skilled employees for unskilled employees. If there were only a fixed number of skilled workers in an area, their wages would be bid up by employers. As in the prior example, the incentive to substitute alternative factors would be reduced, and the reduction in unskilled employment due to the substitution effect would be smaller. In contrast, if the prices of other inputs did not increase when employers attempted to increase their use, other things equal, the substitution effect—and thus the wage elasticity of labor demand—would be larger.

Note again that prices of other inputs are less likely to be bid up in the long run than in the short run. In the long run, existing producers of capital equipment can expand their capacity and new producers can enter the market. Similarly, in the long run, more skilled workers can be trained. This observation is an additional reason the demand for labor will be more elastic in the long run.

The Share of Labor in Total Costs Finally, the share of the category of labor in total costs is crucial to the size of the elasticity of labor demand. If the category's initial share were 20 percent, a 10 percent increase in the wage rate, other things equal, would raise total costs by 2 percent. In contrast, if its initial share were 80 percent, a 10 percent increase in the wage rate would increase total costs by 8 percent. Since employers would have to increase their product prices by more in the latter case, output and employment would fall more in that case. *Thus, the greater the category's share in total costs, the greater the wage elasticity of demand*.[4]

Estimates of Own-Wage Labor Demand Elasticities

We now turn to the results of studies that estimate own-wage demand elasticities for labor as a generic input (that is, labor undifferentiated by skill level). The estimates we discuss are based on studies that utilize wage, output, and employment data from firms or narrowly defined industries. Thus, the employment responses being estimated approximate those that would be expected to occur in a firm that had to raise wages to remain competitive in the labor market. These estimates are suggestive of what might be a "typical" response but, of course, are not indicative of what would happen with any particular firm.

As our analysis has indicated, employers' labor demand responses to a wage change can be broken down into two components: a scale effect and a substitution effect. These two effects can themselves be expressed as elasticities, and their sum is the own-wage labor demand elasticity. In Table 4.1, we display the results of estimates of (a) the short-run scale effect, (b) the substitution effect, and (c) the overall elasticity of demand for labor in the long run.

The scale effect (expressed as an elasticity) is defined as the percentage change in employment associated with a given percentage change in the wage, *holding production technology constant*; that is, it is the employment response that occurs without a substitution effect. By definition, the *short-run* labor demand elasticity includes *only* the scale effect, although we noted earlier that the scale effect is likely to be greater in the long run than it is in the short run (owing to greater possibilities for *product market* substitutions in the long run). Therefore, estimates of short-run labor demand elasticities will be synonymous with the short-run scale effect, which may approximate the long-run scale effect if product market substitutions are relatively swift. A study using data from British manufacturing plants estimated the short-run, own-wage labor demand elasticity to be −0.53 (see

[4]An exception to the law occurs when it is easier for employers to substitute other factors of production for the category of labor than it is for customers to substitute other products for the product being produced; in this case, the law is reversed. Suppose, for example, that the elasticity of product demand among a firm's customers were zero; in this case, a rising wage rate would create only a substitution effect. With a larger labor share, and thus a higher ratio of labor to capital, the *percentage* fall in labor usage as wages rise will tend to be smaller, thus causing the elasticity of demand for labor to be smaller. For more on the effects of labor's share on the elasticity of demand, see Saul D. Hoffman, "Revisiting Marshall's Third Law: Why Does Labor's Share Interact with the Elasticity of Substitution to Decrease the Elasticity of Labor Demand?" *Journal of Economic Education* 40 (Fall 2009): 437–445.

Table 4.1

Components of the Own-Wage Elasticity of Demand for Labor: Empirical Estimates Using Plant-Level Data

	Estimated Elasticity
Short-Run Scale Effect	
British manufacturing firms, 1974–1982	−0.53
Substitution Effect	
32 studies using plant or narrowly defined industry data	Average: −0.45 (typical range: −0.15 to −0.75)
Overall Labor Demand Elasticity	
British plants, 1984	−0.93
British coal mines, 1950–1980	−1.0 to −1.4

Source: Daniel S. Hamermesh, *Labor Demand* (Princeton, N.J.: Princeton University Press, 1993): 94–104.

Table 4.1). The short-run labor demand curve for a typical firm or narrowly defined sector, therefore, would appear to be inelastic.

The substitution effect, when expressed as an elasticity, is the percentage change in employment associated with a given percentage change in the wage rate, *holding output constant*. That is, it is a measure of how employers change their production techniques in response to wage changes, even if output does not change (that is, even if the scale effect is absent). It happens that substitution effects are easier to credibly estimate, so there are many more studies of these effects. One careful summary of 32 studies estimating substitution-effect elasticities placed the average estimated elasticity at −0.45 (which is what is displayed in Table 4.1), with most estimates falling into the range of −0.15 to −0.75.[5]

With the short-run scale elasticity and the substitution elasticity each very close to −0.5, it is not surprising that estimates of the long-run overall elasticity of demand for labor are close to unitary in magnitude. Table 4.1 indicates that a study of plants across several British industries estimated an own-wage elasticity of −0.93, whereas another of British coal mines placed the elasticity of demand for labor in the range of −1.0 to −1.4.[6] Thus, these estimates suggest that if the wages a firm must pay rise by 10 percent, the firm's employment will shrink by close to 10 percent in the long run, other things being equal (that is, unless something else occurs that also affects the demand for labor).

[5]Daniel S. Hamermesh, *Labor Demand* (Princeton, N.J.: Princeton University Press, 1993): 103.

[6]A more recent analysis of the wages and employment of American women in the period following World War II estimates that the overall elasticity of demand for their labor was very similar—in the range of −1.0 to −1.5. See Daron Acemoglu, David H. Autor, and David Lyle, "Women, War and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury," *Journal of Political Economy* 112 (June 2004): 497–551. Estimates of the own-wage elasticity of demand for skilled and unskilled manufacturing labor in Germany are somewhat lower (–0.6 to –1.3); see John T. Addison, Lutz Bellmann, Thorsten Schank, and Paulino Teixeira, "The Demand for Labor: An Analysis Using Matched Employer–Employee Data from the German LIAB. Will the High Unskilled Worker Own-Wage Elasticity Please Stand Up?" *Journal of Labor Research*, 29 (June 2008): 114–137.

Applying the Laws of Derived Demand: Inferential Analysis

Because empirical estimates of demand elasticities that may be required for making particular decisions are often lacking, it is frequently necessary to guess what these elasticities are likely to be. In making these guesses, we can apply the laws of derived demand to predict at least relative magnitudes for various types of labor. Consider first the demand for unionized New York City garment workers. As we shall discuss in chapter 13, because unions are complex organizations, it is not always possible to specify what their goals are. Nevertheless, it is clear that most unions value both wage *and* employment opportunities for their members. This observation leads to the simple prediction that, other things equal, the more elastic the demand for labor, the smaller the wage gain that a union will succeed in winning for its members. The reason for this prediction is that the more elastic the demand curve, the greater the percentage employment decline associated with any given percentage increase in wages. As a result, we can expect the following:

1. Unions would win larger wage gains for their members in markets with inelastic labor demand curves.
2. Unions would strive to take actions that reduce the wage elasticity of demand for their members' services.
3. Unions might first seek to organize workers in markets in which labor demand curves are inelastic (because the potential gains to unionization are higher in these markets).

Because of foreign competition, the price elasticity of demand for the clothing produced by New York City garment workers is extremely high. Furthermore, employers can easily find other inputs to substitute for these workers—namely, lower-paid nonunion garment workers in the South or in other countries. These facts lead one to predict that the wage elasticity of demand for New York City unionized garment workers is very high. Consequently, union wage demands have historically been moderate. The union has also sought to reduce the elasticity of product demand by supporting policies that reduce foreign competition, and it has pushed for higher federal minimum wages to reduce employers' incentives to move their plants to the South. (For another illustration of how an elastic *product* demand inhibits union wage increases, see Example 4.1.)

Next, consider the wage elasticity of demand for unionized airplane pilots in the United States. Only a small share of the costs of operating large airplanes goes to pay pilots' salaries; such salaries are dwarfed by fuel and capital costs. Furthermore, substitution possibilities are limited; there is little room to substitute unskilled labor for skilled labor (although airlines can substitute capital for labor by reducing the number of flights they offer while increasing the size of airplanes). In addition, before the deregulation of the airline industry in 1978, many airlines faced no competition on many of their routes or were prohibited from reducing their prices to compete with other airlines that flew the same routes. These factors all suggest that the wage elasticity of demand for airline pilots was quite low (inelastic). As one might expect, pilots' wages were also quite high because their

EXAMPLE 4.1

Why Are Union Wages So Different in Two Parts of the Trucking Industry?

The trucking industry's "general freight" sector, made up of motor carriers that handle nonspecialized freight requiring no special handling or equipment, is split into two distinct segments. One type of general freight carrier exclusively handles full truckloads (TLs), taking them directly from a shipper to a destination. The other type of carrier handles less-than-truckload (LTL) shipments, which involve multiple shipments on each truck and an intricate coordination of pickups and deliveries. These two segments of the general freight industry have vastly different *elasticities of product demand*; thus, the union that represents truck drivers has a very different ability to raise wages (without suffering unacceptable losses of employment) in each segment.

The TL part of the industry has a product market that is very competitive, because it is relatively easy for firms or individuals to enter the market; one needs only a truck, the proper driver's license, and access to a telephone (to call a freight broker, who matches available drivers with shipments needing delivery). Because this part of the industry has many competing firms, with the threat of even more if prices rise, each firm faces a relatively elastic product demand curve.

Firms specializing in LTL shipments must have a complex system of coordinated routes running between and within cities, and they must therefore be large enough to support their own terminals for storing and transferring shipments from one route to another. The LTL segment of the industry is not easily entered and thus is partially monopolized. From 1980 to 1995—a time period over which the number of TL carriers tripled—virtually the only new entrants into the LTL market were regional subsidiaries of pre-existing national carriers! To contrast competition in the two product markets somewhat differently, in 1987, the four largest LTL carriers accounted for 37 percent of total LTL revenues, while the four largest TL carriers accounted for only 11 percent of TL revenues.

The greater extent of competition in the TL part of the industry implies that at the firm level, *product* demand is more elastic there than in the LTL sector; other things being equal, then, we would expect the *labor* demand curve to also be more elastic in the TL sector. Because unions worry about potential job losses when negotiating with carriers about wages, we would expect to find that union wages are lower in the TL than in the LTL part of the industry. In fact, a 1991 survey revealed that the union mileage rates (drivers are typically compensated on a cents-per-mile basis) were dramatically different in the two sectors:

TL sector

Average union rate: 28.4 cents per mile
Ratio, union to nonunion rate: 1.23

LTL sector

Average union rate: 35.8 cents per mile
Ratio, union to nonunion rate: 1.34

The above data support the theoretical implication that a union's power to raise wages is greater when product (and therefore labor) demand is relatively inelastic. In the less-competitive LTL segment of the trucking industry, union drivers' wages are higher, both absolutely and relative to nonunion wages, than they are in the more competitive TL sector.

Data from: Michael H. Belzer, "Collective Bargaining after Deregulation: Do the Teamsters Still Count?" *Industrial and Labor Relations Review* 48 (July 1995): 636–655; and Michael H. Belzer, *Paying the Toll: Economic Deregulation of the Trucking Industry* (Washington, D.C.: Economic Policy Institute, 1994).

union could push for large wage increases without fear that these increases would substantially reduce pilots' employment levels. However, after airline deregulation, competition among airline carriers increased substantially, leading to a more elastic labor demand for pilots. As a result, many airlines "requested," and won, reduced wages from their pilots.

The Cross-Wage Elasticity of Demand

Because firms may employ several categories of labor and capital, the demand for any one category can be affected by price changes in the others. For example, if the wages of carpenters rose, more people might build brick homes and the demand for *masons* might increase. An increase in carpenters' wages might decrease the overall level of home building in the economy, however, which would decrease the demand for *plumbers*. Finally, changes in the price of *capital* could increase or decrease the demand for workers in all three trades.

The direction and magnitude of the above effects can be summarized by examining the elasticities of demand for inputs with respect to the prices of *other* inputs. The *elasticity of demand for input j with respect to the price of input k* is the percentage change in the demand for input j induced by a 1 percent change in the price of input k. If the two inputs are both categories of labor, these *cross-wage elasticities of demand* are given by

$$\eta_{jk} = \frac{\%\Delta E_j}{\%\Delta W_k} \tag{4.2}$$

and

$$\eta_{kj} = \frac{\%\Delta E_k}{\%\Delta W_j}$$

where, again, the Greek letter η is used to represent the elasticity. If the cross-elasticities are positive (with an increase in the price of one "category" increasing the demand for the other), the two are said to be *gross substitutes*. If these cross-elasticities are negative (and an increase in the price of one "category" reduces the demand for the other), the two are said to be *gross complements* (refer back to Figure 3.3).

It is worth reiterating that whether two inputs are gross substitutes or gross complements depends on the relative sizes of the scale and substitution effects. To see this, suppose we assume that adults and teenagers are substitutes in production. A decrease in the teenage wage will thus have opposing effects on adult employment. On the one hand, there is a substitution effect: for a given level of output, employers will now have an incentive to substitute teens for adults in the production process and reduce adult employment. On the other hand, there is a scale effect: a lower teenage wage reduces costs and provides employers with an incentive to increase employment of all inputs, including adults.

If the scale effect proves to be smaller than the substitution effect, adult employment will move in the same direction as teenage wages, and the two groups will be gross substitutes. In contrast, if the scale effect is larger than the substitution effect, adult employment and teenage wages will move in opposite directions, and the two groups will be gross complements. Knowing that two groups are substitutes in production, then, is not sufficient to tell us whether they are gross substitutes or gross complements.[7]

Because economic theory cannot indicate in advance whether two given inputs will be gross substitutes or gross complements, the major policy questions about cross-wage elasticities of demand relate to the issue of their *sign*; that is, we often want most to know whether a particular cross-elasticity is positive or negative. Before turning to a review of actual findings, we analyze underlying forces that determine the signs of cross-elasticities.

Can the Laws of Derived Demand Be Applied to Cross-Elasticities?

The Hicks–Marshall laws of derived demand are based on four technological or market conditions that determine the size of *own-wage* elasticities. Each of the four conditions influences the substitution or the scale effect, and as noted above, the relative strengths of these two effects are also what determine the sign of *cross-elasticities*. The laws that apply to own-wage elasticities cannot be applied directly to cross-elasticities, because with cross-elasticities, the substitution effect (if there is one) and the scale effect work in opposite directions. The same underlying considerations, however, are basic to an analysis of cross-elasticities.

As we discuss these four considerations in the context of cross-elasticities, it will be helpful to have an example in mind. Let us return, then, to the question of what might happen to the demand for adult workers if the wages of teenage workers were to fall. As noted above, the answer depends on the relative strengths of the scale and substitution effects. What determines the strength of each?

The Scale Effect The most immediate effect of a fall in the wages of teenagers would be reduced production costs for those firms that employ them. Competition in the product market would ensure that lower costs are followed by price reductions, which should stimulate increases in both product demand and the level of output. Increased levels of output will tend to cause increases in employment of all kinds of workers, including adults. This chain of events obviously describes

[7]As noted in chapter 3, if two groups are complements in production, a decrease in the price of one should lead to increased employment of the other. Complements in production are always gross complements.

behavior underlying the scale effect, and we now investigate what conditions are likely to make for a strong (or weak) scale effect.

The initial cost (and price) reductions would be greater among those employers for whom teenage wages constituted a higher proportion of total costs. Other things equal, greater price reductions would result in greater increases in both product demand and overall employment. Thus, *the share of total costs devoted to the productive factor whose price is changing* will influence the size of the scale effect. The larger this share, other things equal, the greater the scale effect (and the more likely it is that gross complementarity will exist). This tendency is analogous to the fourth Hicks–Marshall law discussed earlier; the difference is that with cross-elasticities, the factor whose *price* is changing is not the same as the one for which *employment* changes are being analyzed.

The other condition that greatly influences the size of the scale effect is product demand elasticity. In the earlier case of teenage wage reductions, the greater the increase in product demand when firms reduce their prices, the greater the tendency for employment of all workers, including adults, to increase. More generally, *the greater the price elasticity of product demand*, *other things equal*, *the greater the scale effect* (*and thus the greater the likelihood of gross complementarity*). The effects of product demand elasticity are thus similar for both own-wage and cross-wage elasticities.

The Substitution Effect After teenage wages fall, firms will also have incentives to alter their production techniques so that teenagers are more heavily used. Whether the greater use of teenagers causes an increase or some loss of adult jobs partially depends on a technological question: are teenagers and adults substitutes or complements in production? If they are complements in production, the effect on adults of changing productive techniques will reinforce the scale effect and serve to unambiguously increase adult employment (meaning, of course, that adults and teenagers would be gross complements). If they are substitutes in production, however, then changing productive techniques involves using a higher ratio of teenagers to adults, and the question then becomes whether this substitution effect is large or small relative to the scale effect.

A technological condition affecting the size of the substitution effect is a direct carryover from the second Hicks–Marshall law discussed previously: *the substitution effect will be greater when the category of labor whose price has changed is easily substituted for other factors of production*. When analyzing the effects on adult employment of a decline in the teenage wage, it is evident that when teenagers are more easily substituted for adults, the substitution effect (and therefore the chances of gross substitutability between the two categories of labor) will be greater.

Another condition influencing the size of the substitution effect associated with a reduction in the teenage wage relates to the labor supply curve of adults. If the adult labor supply curve were upward-sloping and rather steep, then adult wages would tend to fall as teenagers were substituted for adults and the demand curve for adults shifted left. This fall would blunt the substitution effect, because adults would also become cheaper to hire. Conversely, if the adult labor supply curve were relatively flat, adult wages would be less affected by reduced demand and the substitution effect would be less blunted. As in the case of own-wage elasticities, *more-elastic supply curves of substitute inputs also lead to a greater substitution effect*, *other things equal*, *in the case of cross-wage elasticities*.[8]

Estimates Relating to Cross-Elasticities

Estimating at least the *sign* of cross-wage labor demand elasticities is useful for answering many public-policy questions. For example, if we were to reduce the *teenage* minimum wage, how would this affect the demand for *adult* labor? If *capital* were to be subsidized, how would this affect the demand for *labor*? Or, to take a hotly debated issue in recent years (and one we will return to in chapter 10), when *immigrant* labor becomes cheaper and more available, what are the likely effects on the demand for various grades of *native* labor? These questions, of course, are really asking whether the pairs of inputs italicized in each sentence are gross complements or gross substitutes.

While the major policy interest is whether two inputs are *gross* complements or *gross* substitutes, obtaining credible estimates is challenging (because it is difficult to estimate scale effects). Therefore, most of the cross-wage empirical studies to date focus on whether two factors are substitutes or complements *in production*. These studies estimate the employment response for one category of labor to a wage or price change elsewhere, *holding output constant* (which in effect allows us to focus just on changes in the *mix* of factors used in production). The factors of production paired together for analysis in these studies are numerous and the results are not always clear-cut; nevertheless, the findings taken as a whole offer at least a few generalizations:[9]

1. Labor and energy are clearly substitutes in production, although their degree of substitutability is small. Labor and materials are probably substitutes in production, with the degree of substitutability again being small.
2. Skilled labor and unskilled labor are substitutes in production.[10]

---

[8]The share of the teenage wage bill in total costs influences the substitution effect as well as the scale effect in the example we are analyzing. For example, if teenage labor costs were a very large fraction of total costs, the possibilities for further substitution of teenagers for adults would be rather limited (this can be easily seen by considering an example in which teenagers constituted 100 percent of all production costs). Thus, while a larger share of teenagers in total cost would make for a relatively large scale effect, it also could reflect a situation in which the possibilities of substituting teenagers for adults are smaller than they would otherwise be.

[9]Hamermesh, *Labor Demand*, 105–127.

[10]James D. Adams, "The Structure of Firm R&D, the Factor Intensity of Production, and Skill Bias," *Review of Economics and Statistics* 81 (August 1999): 499–510; and Antonio Ciccone and Giovanni Peri, "Long-Run Substitutability between More and Less Educated Workers: Evidence from U.S. States, 1950–1990," *Review of Economics and Statistics* 87 (November 2005): 652–663.

3. We are not certain whether either skilled or unskilled labor is a substitute for or a complement with capital in the production process. What does appear to be true is that skilled (or well-educated) labor is more likely to be complementary with capital than is unskilled labor—and that if they are both substitutes for capital, the degree of substitutability is smaller for skilled labor.[11]
4. The finding summarized in 3 above suggests that skilled labor is more likely than unskilled labor to be a *gross* complement with capital. This finding is important to our understanding of recent trends in the earnings of skilled and unskilled workers (see chapter 15), because the prices of computers and other high-tech capital goods have fallen dramatically in the past decade or so.
5. The finding in 3 above also implies that if the wages of both skilled and unskilled labor were to rise by the same percentage, the magnitude of any employment loss associated with the substitution effect (as capital is substituted for labor) will be greater for the unskilled. Thus, we expect that, other things equal, *own-wage* labor demand elasticities will be larger in magnitude for unskilled than for skilled workers.

Policy Application: Effects of Minimum Wage Laws

History and Description

The Fair Labor Standards Act of 1938 was the first major piece of protective labor legislation adopted at the national level in the United States. Among its provisions were a minimum wage rate, below which hourly wages could not be reduced, an overtime-pay premium for workers who worked long workweeks, and restrictions on the use of child labor. When initially adopted, the minimum wage was set at $0.25 an hour and covered roughly 43 percent of all nonsupervisory wage and salary workers—primarily those employed in larger firms involved in interstate commerce (manufacturing, mining, and construction). Both the basic minimum wage and coverage under the minimum wage have expanded over time. Indeed, as of 2009, the minimum wage was set at $7.25 an hour and roughly 90 percent of all nonsupervisory workers were covered by its provisions.

It is important to emphasize that the minimum wage rate is specified in *nominal* terms and not in terms *relative* to some other wage or price index. As illustrated in Figure 4.3, the nominal wage rate has usually been raised only once every few years. Until the early 1980s, newly legislated minimum wage rates were typically at least 45 percent of the average hourly wage in manufacturing. During the years between legislation, productivity growth and inflation caused

[11]See Claudia Goldin and Lawrence Katz, "The Origins of Technology-Skill Complementarity," *Quarterly Journal of Economics* 113 (May 1998): 693–732, for citations to the literature.

Figure 4.3

Federal Minimum Wage Relative to Wages in Manufacturing, 1938–2009

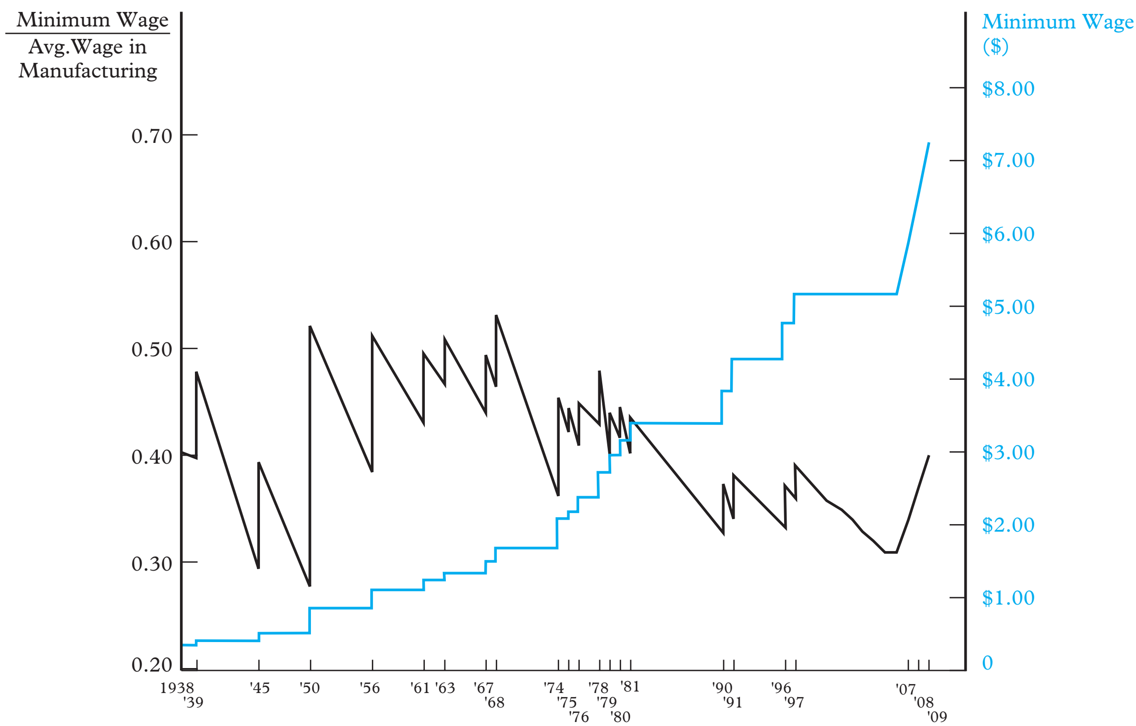

manufacturing wages to rise, with the result that the minimum wage has often fallen by 10 or more percentage points relative to the manufacturing wage before being raised again. In the last two decades, even the newly legislated minimums were below 40 percent of the average manufacturing wage. Under a law passed by Congress in 2007, which set the minimum wage at $5.85 and called for it to rise to $7.25 over a two-year period, the minimum wage in 2009 was again about 40 percent of the average manufacturing wage.

Employment Effects: Theoretical Analysis

Since the minimum wage was first legislated, a concern has been that it will reduce employment, especially among the groups it is intended to benefit. In the face of downward-sloping labor demand curves, a policy that compels firms to raise the wages paid to all low-wage workers can be expected to *reduce employment opportunities* for the least skilled or least experienced. Furthermore, if the percentage loss of employment among low-wage workers is greater than the percentage increase in their wages—that is, if the demand curve for low-wage workers is *elastic*—then the *aggregate* earnings of low-wage workers could be made smaller by an increase in the minimum wage.

In evaluating the findings of research on the employment effects of minimum wages, we must keep in mind that good research must be guided by good theory. Theory provides us with a road map that directs our explorations into the real world, and it suggests several issues that must be addressed by any research study of the minimum wage.

Nominal versus Real Wages Minimum wage levels in the United States have been set in nominal terms and adjusted by Congress only sporadically. The result is that general price inflation gradually lowers the real minimum wage during the years between congressional action, so what appears to be a fixed minimum wage turns out to have constantly changing incentives for employment.

Also, the federal minimum wage in the United States is uniformly applied to a large country characterized by regional differences in prices. Taking account of regional differences in prices or wages, we find that the real minimum wage in Alaska (where wages and prices are very high) is lower than it is in Mississippi. Recognizing that there are regional differences in the real minimum wage leads to the prediction that employment effects of a uniformly applied minimum wage law generally will be most adverse in regions with the lowest costs of living. (Researchers must also take into account the fact that many states have their own minimum wage laws, many having minimums that exceed the federal minimum.)

Holding Other Things Constant Predictions of job loss associated with higher minimum wages are made *holding other things constant*. In particular, the prediction grows out of what is expected to happen to employment as one moves up and to the left along a *fixed* labor demand curve. If the labor demand curve were to shift at the same time that a new minimum becomes effective, the employment effects of the shift could be confounded with those of the new minimum.

Consider, for example, Figure 4.4, where, for simplicity, we have omitted the labor supply curve and focused on only the demand side of the market. Suppose that D_0 is the demand curve for low-skilled labor in year 0, in which year the real wage is W_0/P_0 and the employment level is E_0. Further assume that in the absence of any change in the minimum wage, the money wage and the price level would both increase by the same percentage over the next year, so that the real wage in year 1 (W_1/P_1) would be the same as that in year 0.

Now, suppose that in year 1, two things happen. First, the minimum wage rate is raised to W_2, which is greater than W_1, so that the real wage increases to W_2/P_1. Second, because the economy is expanding, the demand for low-skilled labor shifts out to D_1. The result of these two changes is that employment increases from E_0 to E_1.

Comparisons of observed employment levels at two points of time have led some investigators to conclude that minimum wage increases had no adverse employment effects. However, this simple before/after comparison is *not* the correct one if labor demand has shifted, as in Figure 4.4. Rather, we should ask, "How did the actual employment level in period 1 compare with the level that *would have*

Figure 4.4

Minimum Wage Effects: Growing Demand Obscures Job Loss

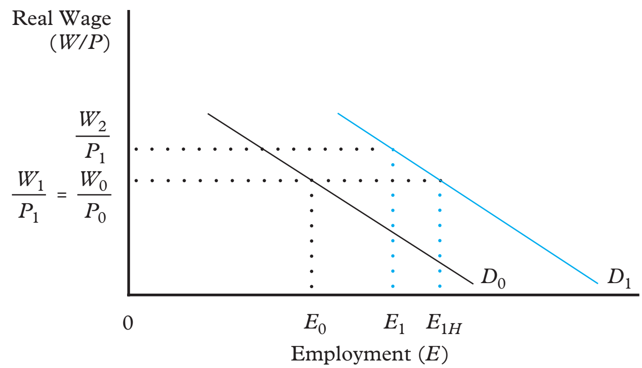

prevailed in the absence of the increase in the minimum wage?" Since demand grew between the two periods, this hypothetical employment level would have been E_{1H}. Because E_{1H} is greater than E_1, the actual level of employment in period 1, there is a loss of jobs ($E_{1H} - E_1$) caused by the minimum wage. In a growing economy, then, the expected effect of a one-time increase in the minimum wage is to *reduce the rate of growth of employment*. Controlling for all the "other things" besides wages that affect labor demand turns out to be the major difficulty in measuring employment changes caused by the minimum wage.

Effects of Uncovered Sectors The federal minimum wage law, like many government regulations, has an *uncovered* sector. Coverage has increased over the years, but the law still does not apply to some nonsupervisory workers (mainly those in small firms in the retail trade and service industries). Also, with millions of employers and limited resources for governmental enforcement, *noncompliance* with the law may be widespread, creating another kind of noncoverage.[12] The existence of uncovered sectors significantly affects how the overall employment of low-wage workers will respond to increases in the minimum wage.

Consider the labor market for unskilled, low-wage workers that is depicted in Figure 4.5. The market has two sectors. In one, employers must pay wages equal to at least the minimum wage of W_1; wages in the uncovered sector are free to vary with market conditions. While the total labor supply to both markets taken as a whole is fixed at E_T (that is, the total labor supply curve is vertical), workers can freely move from one sector to the other seeking better job offers. Free movement between sectors suggests that in the absence of minimum wage regulations, the wage in each sector will be the same. Referring to Figure 4.5, let

[12]Orley Ashenfelter and Robert Smith, "Compliance with the Minimum Wage Law," *Journal of Political Economy* 87 (April 1979): 335–350. A more recent study of noncompliance (among apparel contractors) can be found in David Weil, "Public Enforcement/Private Monitoring: Evaluating a New Approach to Regulating the Minimum Wage," *Industrial and Labor Relations Review* 58 (January 2005): 238–257.

Figure 4.5

Minimum Wage Effects: Incomplete Coverage Causes Employment Shifts

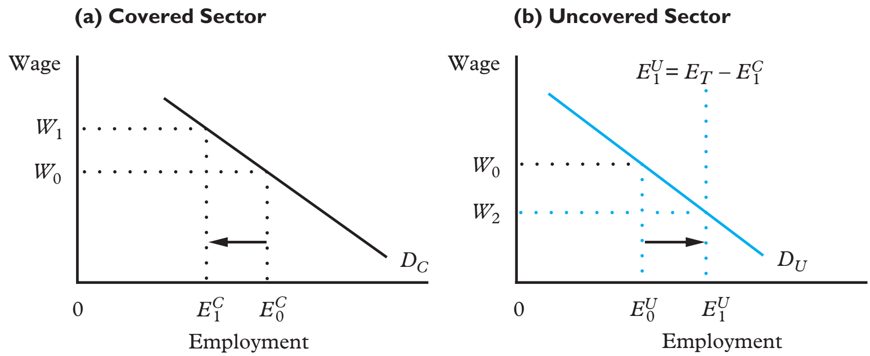

us assume that this "pre-minimum" wage is W_0 and that total employment of E_T is broken down into E_0^c in the covered sector plus E_0^u in the uncovered sector.

If a minimum wage of W_1 is imposed on the covered sector, all unskilled workers will prefer to work there. However, the increase in wages in that sector, from W_0 to W_1, reduces demand, and covered-sector employment will fall from E_0^c to E_1^c. Some workers who previously had, or would have found, jobs in the covered sector must now seek work in the uncovered sector. Thus, to the E_0^u workers formerly working in the uncovered sector are added $E_0^c - E_1^c$ other workers seeking jobs there. Hence, all unskilled workers in the market who are not lucky enough to find "covered jobs" at W_1 must now look for work in the uncovered sector,[13] and the (vertical) supply curve to that sector becomes $E_1^u = E_0^u + (E_0^c - E_1^c) = E_T - E_1^c$. The increased supply of workers to that sector drives down the wage there from W_0 to W_2.

The presence of an uncovered sector thus suggests the possibility that employment among unskilled workers will be rearranged, but not reduced, by an increase in the minimum wage. In the above example, all E_T workers remained employed after the minimum was imposed. Rather than reducing overall employment of the unskilled, then, a partially covering minimum wage law might serve to shift employment out of the covered to the uncovered sector, with the further result that wages in the uncovered sector would be driven down.

The magnitude of any employment shift from the covered to the uncovered sector, of course, depends on the size of the latter; the smaller it is, the lower are the chances that job losers from the covered sector will find employment there. Whatever the size of the uncovered sector, however, its very presence means that

[13]Under some circumstances, it may be rational for these unemployed workers to remain unemployed for a while and to search for jobs in the covered sector. We shall explore this possibility of "wait unemployment"—which is discussed by Jacob Mincer in "Unemployment Effects of Minimum Wage Changes," *Journal of Political Economy* 84 (August 1976): S87–S104—in chapter 13. At this point, we simply note that if it occurs, unemployment will result.

the *overall* loss of employment is likely to be less than the loss of employment in the covered sector.

Intersectoral Shifts in Product Demand The employment effects of a wage change are the result of scale and substitution effects. Substitution effects stem from changes in how firms choose to produce, while scale effects are rooted in consumer adjustments to changes in product prices. Recall that faced with a given increase (say) in the minimum wage, firms' increases in costs will generally be greater when the share of low-wage labor in total costs is greater; thus, the same increase in the minimum wage can lead to rather different effects on product prices among different parts of the covered sector. Furthermore, if these subsectors compete with each other for customers, it is possible that scale effects of the increased wage will serve to *increase* employment among some firms in the covered sector.

Suppose, for example, that convenience stores sell items that supermarkets also carry and that a minimum wage law raises the wages paid to low-skilled workers in both kinds of stores. If low-skilled labor costs are a higher fraction of total costs in convenience stores than they are in supermarkets, then, other things equal, the minimum wage law would raise costs in convenience stores by a greater percentage. With prices of items increasing more in convenience stores than in supermarkets, consumers would tend to shift some of their convenience store purchases to supermarkets. Thus, the minimum wage increase could have an ambiguous effect on employment in supermarkets. On the one hand, increased costs of unskilled workers in supermarkets would create scale and substitution effects that cause employment to decline. On the other hand, because they may pick up business formerly going to convenience stores, supermarkets may experience a scale effect that could work to increase their demand for labor.

Employment Effects: Empirical Estimates

While the initial employment effects of adopting a minimum wage in the United States were readily observed (see Example 4.2), the effects of more recent increases are not as obvious—and must therefore be studied using sophisticated statistical techniques. The demographic group for which the effects of minimum wages are expected to be most visible is composed of teenagers—a notoriously low-paid group!—but studies of how mandated wage increases have affected their employment have produced no consensus.

Widely reviewed and replicated studies of employment changes in the fast-food industry, for example, disagree on whether employment was affected at all by minimum wage increases in the early 1990s.[14] A study that reviewed and

[14]See David Neumark and William Wascher, "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Comment," and David Card and Alan B. Krueger, "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply," both in *American Economic Review* 90 (December 2000): 1362–1420. These studies contain references to earlier studies and reviews on this topic.

EXAMPLE 4.2

The Employment Effects of the First Federal Minimum Wage

When the federal minimum wage first went into effect, on October 24, 1938, it was expected to have a substantial impact on the economy of the South, where wages were much lower than in the rest of the country. An examination of one of the largest manufacturing industries in the South, seamless hosiery, verifies these predictions.

It is readily apparent that the new minimum wage was binding in the seamless hosiery industry. By 1940, nearly one-third of the labor force earned within 2.5 cents per hour of the minimum wage (which was then 32.5 cents per hour). A longitudinal survey of 87 firms shows that employment, which had been rising, reversed course and started to fall, even though overall demand for the product and production levels were rising. Employment fell by 5.5 percent in southern mills but rose by 4.9 percent in northern mills. Even more strikingly, employment fell by 17 percent in mills that had previously paid less than the new minimum wage, while it stayed virtually the same at higher-wage mills.

Before the passage of the minimum wage, there had been a slow movement from the use of hand-transfer to converted-transfer knitting machines. (A converted-transfer machine had an attachment to enable automated production for certain types of work.) The minimum wage seems to have accelerated this trend. In the first two years of the law's existence, there was a 23 percent decrease in the number of hand-transfer machines, a 69 percent increase in converted-transfer machines, and a 10 percent increase in fully automatic machines. In addition, the machines were used more intensively than before. A night shift was added at many mills, and these workers did not receive extra pay for working this undesirable shift. Finally, total imports of seamless hosiery surged by about 27 percent within two years of the minimum wage's enactment.

Data from: Andrew J. Seltzer, "The Effects of the Fair Labor Standards Act of 1938 on the Southern Seamless Hosiery and Lumber Industries," *Journal of Economic History* 57 (June 1997): 396–415.

updated prior estimates of how *overall* teenage employment has responded to increases in the minimum wage, however, found negative effects on employment. Once account is taken of the extent to which minimum wage increases raised the average wage of teenagers, the implications of this latter study are that the elasticity of demand for teenagers is in the range of −0.4 to −1.9.[15]

A recent estimate of how increases in the minimum wage affects employment for *all* low-wage workers, not just teenagers, suggests an own-wage labor demand elasticity that is considerably lower. This study looked at the employment status of those who were at or near the minimum wage right before it increased and then looked at their employment status a year later. The estimated

[15]The reviews cited in footnote 2 suggest that the elasticity of teenage employment with respect to changes in the minimum wage generally falls into the range of −0.2 to −0.6. Dividing these elasticities by the estimated elasticity of response in the average teen wage to changes in the minimum wage (the percentage change in the average teen wage divided by the percentage change in the minimum wage was in the range of 0.32 to 0.48) yields estimates of the elasticity of the labor demand curve for teenagers. A recent study suggests that most of the effects of minimum wages on teenage employment are observed in temporary jobs or new hires; see Jeffrey P. Thompson, "Using Local Labor Market Data to Re-Examine the Employment Effects of the Minimum Wage," *Industrial and Labor Relations Review* 62 (April 2009): 343–366.

decline in the probability of employment implied that the labor demand curve facing these workers has an own-wage elasticity of roughly −0.15.[16]

With some studies estimating no effect on employment, and with many of those that do estimating an own-wage labor demand elasticity well below unity (the average we saw in Table 4.1), we remain notably uncertain about how employment among low-wage workers responds to increases in the minimum wage. We will come back to this issue in chapter 5 and offer a possible reason for why *mandated* wage increases might have a smaller and more uncertain effect on labor demand than wage increases generated by market forces.

Does the Minimum Wage Fight Poverty?

Aside from the potentially adverse effects on employment opportunities for low-wage workers, two other reasons suggest that the minimum wage is a relatively ineffective instrument to reduce poverty. First, many who live in poverty are not affected by the minimum wage, either because they are not employed or because their wages, while low, are already above the minimum. For example, one study of the minimum wage increases in 1990–1991 divided the distribution of family incomes into 10 equally sized groups (deciles). Among adults in the lowest decile, 80 percent were below the poverty line (given the size of their families), yet only about one-quarter of them worked; of those who did work, less than one-third earned wages that were less than the new minimum![17] Thus, even without any loss of employment opportunities, less than 10 percent of those in the lowest income decile stood to benefit from the 1990–1991 increases in the minimum wage.

Second, many of those most affected by the minimum wage are teenagers, who may not reside in poor families. The study cited earlier found that only 19 percent of the estimated earnings increases generated by the higher minimum wage went to families with incomes below the poverty line, while over 50 percent of the increases went to families whose incomes were at least twice the poverty level.

[16]David Neumark, Mark Schweitzer, and William Wascher, "The Effects of Minimum Wages Throughout the Wage Distribution," *Journal of Human Resources* 39 (Spring 2004): 425–450. This study finds that the elasticity of employment with respect to changes in the minimum wage for those at the minimum is −0.12, while the elasticity of their wages to changes in the minimum is +0.8; dividing −0.12 by 0.8 equals the estimated demand elasticity of −0.15.

[17]Richard V. Burkhauser, Kenneth A. Couch, and David C. Wittenburg, "'Who Gets What' from Minimum Wage Hikes: A Re-Estimation of Card and Krueger's Distributional Analysis in *Myth and Measurement: The New Economics of the Minimum Wage*," *Industrial and Labor Relations Review* 49 (April 1996): 547–552. For a summary of studies on how minimum-wage laws (and other social policies) affect poverty, see David Neumark, "Alternative Labor Market Policies to Increase Economic Self-Sufficiency: Mandating Higher Wages, Subsidizing Employment, and Increasing Productivity," working paper no. 14807, National Bureau of Economic Research (Cambridge, Mass., March 2009); also see Joseph J. Sabia and Richard V. Burkhauser, "Minimum Wages and Poverty: Will a $9.50 Federal Minimum Wage Really Help the Working Poor?" *Southern Economic Journal* 76 (January 2010): 592–623. For a different view of minimum-wage laws, see Bruce E. Kaufman, "Institutional Economics and the Minimum Wage: Broadening the Theoretical and Policy Debate," *Industrial and Labor Relations Review* 63 (April 2010): 427–453.

"Living Wage" Laws

Perhaps because the federal minimum wage is relatively low and has not been changed very often, roughly 100 cities, counties, and school districts in the United States have adopted "living wage" ordinances. These ordinances apply to a subset of employers within their jurisdictions and impose wage floors that are higher than either federal or state minimum wages on these employers. The affected employers are generally those performing contracts with the local government, although in some cases, the ordinances also apply to employers receiving business assistance from the city or county. Living wage levels usually relate to the federal poverty guidelines, which in 2007 were $17,170 for a family of three and $20,650 for a family of four in the continental United States (it takes wages of $8.50 to $10 per hour to reach these poverty lines). In 2007, typical wage levels specified by living wage laws were in the range of $8 to $12 per hour.

The potentially beneficial effects of living wage ordinances for low-wage workers are obviously limited by the rather narrow group of employers to which they apply. The benefits are also reduced, of course, if these laws cause the affected employers to either reduce their employment levels or move their operations to cities that do not have living wage regulations.

Estimating the employment effects of adopting living wage laws, however, requires more than merely comparing employment changes in cities with and without such regulations, because the two groups of cities may have fundamentally different employment or wage trends. Cities with rapidly expanding employment opportunities, for example, may decide differently about adopting a living wage law than cities with stagnant or declining opportunities. Because these laws are relatively new, and because the best way to evaluate their employment effects is subject to debate, there is currently no consensus about how living wage ordinances affect employment.[18]

Applying Concepts of Labor Demand Elasticity to the Issue of Technological Change

Technological change, which can encompass the introduction of new products and production techniques as well as changes in technology that serve to reduce the cost of capital (for example, increases in the speed of computers), is frequently viewed as a blessing by some and a curse by others. Those who view it positively point to

[18] One promising way to estimate the employment effects is to compare employment changes in cities that implemented living wage laws with those in cities that passed such laws but saw them derailed by some outside force (the state legislature or a court decision). This approach is included in Scott Adams and David Neumark, "The Effects of Living Wage Laws: Evidence from Failed and Derailed Living Wage Campaigns," *Journal of Urban Economics* 58 (September 2005): 177–202. For estimated employment effects of a *city* minimum wage, see Arindrajit Dube, Suresh Naidu, and Michael Reich, "The Economic Effects of a Citywide Minimum Wage," *Industrial and Labor Relations Review* 60 (July 2007): 522–543.

the enormous gains in the standard of living made possible by new technology, while those who see technological change as a threat often stress its adverse consequences for workers. Are the concepts underlying the elasticity of demand for labor useful in making judgments about the effects of technological change?

Product Demand Shifts There are two aspects of technological change that affect the demand for labor. One is product demand. *Shifts* in product demand curves will tend to shift labor demand curves in the same direction, and changes in the *elasticity* of product demand with respect to product price will tend to cause qualitatively similar changes in the own-wage elasticity of labor demand. The invention of new products (personal computers, for example) that serve as substitutes for old ones (typewriters) will tend to shift the labor demand curve in the older sector to the left, causing loss of employment in that sector. If greater product substitution possibilities are also created by these new inventions, the introduction of new products can increase the *elasticity* of product—and labor—demand. This increases the amount of job loss associated with collectively bargained wage increases, and it reduces the power of unions to secure large wage increases in the older sector. While benefiting consumers and providing jobs in the new sectors, the introduction of new products does necessitate some painful changes in established industries, as workers, unions, and employers must all adjust to a new environment.

Capital-Labor Substitution A second aspect of technological change is often associated with automation, or the substitution of capital for labor. For purposes of analyzing its effects on labor demand, this second aspect of technological change should be thought of as reducing the cost of capital. In some cases—the mass production of personal computers is one example—a fall in capital prices is what literally occurs. In other cases of technological change—the miniaturization of computer components, for example, which has made possible new production techniques—an invention makes completely new technologies available. When something is unavailable, it can be thought of as having an infinite price (it is not available at any price); therefore, the availability of a new technique is equivalent to observing a decline in its price to some finite number. In either case, with a decline in its cost, capital tends to be substituted for labor in the production process.

The *sign* of the cross-elasticity of demand for a given category of labor with respect to a fall in the price of capital depends on whether capital and the category of labor are gross substitutes or gross complements. If a particular category of labor is a substitute in production for capital, *and* if the scale effect of the reduced capital price is relatively weak, then capital and the category of labor are gross substitutes and automation reduces demand for workers in this category. For categories of labor that are not close substitutes for the new technology, however, the scale effect may dominate, and the two can be gross complements. Thus, the effect of automation on the demand for *particular* categories of labor can be either positive or negative.

Clearly, whether capital and a given type of labor are gross substitutes depends on several factors, all of which are highly specific to particular industries and production processes. Perhaps the most that can be said generally is that unskilled labor and capital are more likely to be substitutes in production than are skilled labor and capital, which some studies have identified as complements in production. Because factors of production that are complementary must be gross complements, technological change is more likely to increase the demand for skilled than for unskilled labor.[19]

Before concluding that technological change is a threat to the unskilled, however, we must keep three things in mind. First, even factors that are substitutes in production can be gross complements (if scale effects are large enough). Second, substitution of capital for labor can destroy some jobs, but accompanying scale effects can create others, sometimes in the same industry.

Finally, although the fraction of all workers who are unskilled laborers has declined over the course of the last 100 years, this decline is not in itself convincing evidence of gross substitutability between capital and unskilled labor. The concepts of elasticity and cross-elasticity refer to changes in labor demand caused by changes in wages or capital prices, *holding all else constant*. That is, labor demand elasticities focus on the labor demand curve at a particular point in time. Actual employment outcomes over time are also influenced by labor *supply* behavior of workers. Thus, from simple observations of employment levels over time, it is impossible to tell anything about own-wage demand elasticities or about the signs or magnitudes of cross-elasticities of labor demand.

Overall Effects of Technological Change From the analysis above, it is clear that technological innovations affect the demand for labor through both the scale and substitution effects. In many public discussions of technological change, however, scale effects are overlooked, and the focus is placed on the substitution effect—sometimes in frightening words. For example, in a book titled *The Collapse of Work*, published in 1979, the authors referred to technological change as creating a "jobs holocaust" and called for policies designed to cope with "ever-increasing unemployment."[20] Because of the fears created by technological change, we need to pause and use economic analysis to consider whether technological change creates, *for an entire society*, more harm than good.

Fortunately, the fear that technological change creates a "jobs holocaust" has proven groundless. When *The Collapse of Work* was published, for example, 60 percent of adults in the United States were working, and among all those who

---

[19]See David Autor, Lawrence Katz, and Alan Krueger, "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (November 1998): 1169–1213. For a study indicating that capital and labor exhibit gross complementarity in the manufacturing sector of developing countries, but that this complementarity is probably larger for skilled workers, see Peter Blair Henry and Diego Sasson, "Capital Market Integration and Wages," working paper no. 15204, National Bureau of Economic Research (Cambridge, Mass., July 2009).

[20]Clive Jenkins and Barrie Sherman, *The Collapse of Work* (Fakenham, England: Eyre Methuen, 1979), chapter 20.

wanted to work, 5.8 percent were unemployed. In 2008, after three decades of rapid technological innovation, the unemployment rate also stood at 5.8 percent (a bit above the average for the years 2000–2009), but 62 percent of American adults were working!

Technological change, however, does impose costs on *some* workers—those who face decreased demand for their services and must therefore bear the costs of changing jobs. These costs may involve wage loss, temporary unemployment, or the need to invest in learning new skills. But because technological innovation also enhances the demand for *other* workers and results in lower costs or greater product variety for *consumers*, it is natural to ask if there is a way to analyze whether the *overall net effects* of technological change are positive or negative. Put differently (and in the context of the normative principles outlined in chapter 1), can economic theory be used to tell us whether, within a society, the gainers gain more from technological change than the losers lose?

To begin our analysis, let us consider a society that has a fixed amount of labor and capital resources, and for the sake of simplicity, let us assume that these resources can be used to produce two goods: food and clothing. Figure 4.6 summarizes the production possibilities we assume for this simple society. If all labor and capital inputs were devoted to the production of food, 200 million units of food (and no clothing) could be produced (see point *Y*). Similarly, if all resources were devoted to the production of clothing, 100 million units of clothing (and no food) could be produced (point *X*). If, say, 50 percent of the resources were devoted to food and 50 percent to clothing, the society could produce 100 million units of food and 50 million units of clothing (point *A*). Limits on the combinations of food and clothing this society could produce are displayed in Figure 4.6 along line *XY*, which is called a "production possibilities curve."[21] All combinations along or below (southwest of) *XY* are possible; combinations above *XY* (to the northeast of it) cannot be produced.

In complex, modern societies, the *actual* mix of food and clothing produced can be decided by the government, by the market, or by some combination of the two. At one extreme, a centralized governmental bureaucracy could mandate how much food and clothing are to be produced; at the other, the decision could arise from the market interactions between consumers (demand) and producers (supply). Of course, even in a market setting, government could influence the mix of food and clothing produced—through taxes, subsidies, or regulations that alter the cost or methods of producing food and/or clothing.

Whatever the decision-making process, we normally assume that a society would want to choose a mix of food and clothing that lies *on* the production possibilities curve rather than a mix that lies *below* the curve. If, for example, a society were to choose the combination of food and clothing represented by point *M* in

[21]The production possibilities "curve" in Figure 4.6 is a straight line, which reflects the simplifying assumption that the ratio at which food can be "transformed" into clothing, and vice versa, never changes. This assumption is not necessary to the argument but does make it a bit easier to grasp initially.

Figure 4.6, it would not be producing as much food or clothing as it could, given its technology and resources. In short, its resources would be under-utilized, and its consumers would not have available to them all the goods these resources would allow. Let us start our analysis, then, by supposing that the society depicted in Figure 4.6 chooses point *A* along *XY* and produces 100 million units of food and 50 million of clothing.

Now, imagine that someone invents a device that doubles the speed of the sewing process, making it possible to produce twice as much clothing with any level of inputs. Thus, if all resources were devoted to the production of clothing, this new device would permit the production of 200 million units of clothing (point *Z*)—a large increase over the old level of 100 million units. However, the new device does nothing to enhance the production of food, so if all resources were devoted to the production of food, this society could still produce only 200 million units of food. The new set of production possibilities is depicted by the blue line (*ZY*) in Figure 4.6.

Looking at Figure 4.6, it is obvious that the new technological invention expands the consumption possibilities for those in this society. They might choose to keep half of their resources allocated to food production and half to clothing production; if so, they could consume the same 100 million units of food but increase their clothing consumption from 50 to 100 million units (see point *B* in Figure 4.6). Alternatively, they could choose to keep clothing consumption at 50 million units, which with the new device now would require only 25 percent of society's resources to produce, and devote 75 percent of their inputs to food; food production would then increase from 100 to 150 million units (see point *C* in the figure). Finally, instead of keeping the production of one good constant and increasing the other, they could choose to allocate inputs so that *more of both goods are produced* (see all the points between *B* and *C*).

Obviously, choosing any point other than *B* involves a reallocation of labor and capital between the food and clothing industries. Even if society were to continue allocating half of its resources to each industry, however, the new sewing technology might change the occupational requirements in the clothing industry—requiring that workers in that industry learn new skills or accept different employment conditions. The faster and more smoothly these inter- and intra-industry changes occur, the faster the move from the initial point on *XY* to a new point on *ZY*. For a society to *actually* obtain the increased production made possible by technological change, then, it must have policies or institutions that promote (or at least permit) the mobility of capital and labor.

To this point, our analysis of the effects of technological change has demonstrated that such change makes it possible for society to obtain more goods and services from its limited resources, thus potentially increasing average consumption per capita.[22] But would greater *average* consumption levels be enough to

[22]For ease of illustration, we have confined our analysis to the two goods of food and clothing—but the analysis and its conclusions are unaffected if we consider a society in which people can consume many goods or services, including *leisure* (see chapter 6).

Figure 4.6

The Production Possibilities for a Hypothetical Society

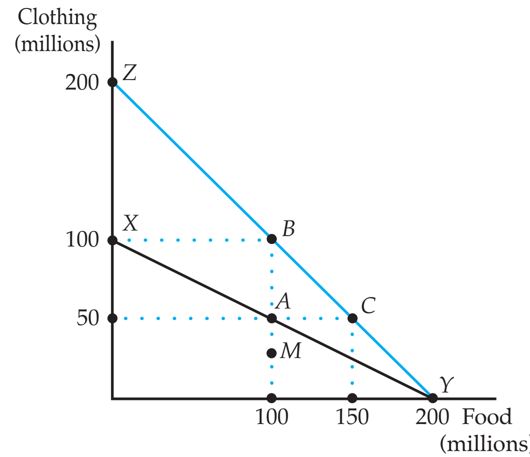

guarantee that society *as a whole* gains from technological change? To answer this question, we must return to some principles of both positive and normative analysis introduced in chapter 1.

Economic theory assumes that individuals, as both workers and consumers, are attempting to maximize their utility. Furthermore, we usually assume utility is enhanced when individuals are able to consume more goods or services (including leisure; see footnote 22). Thus, one might think that when technological change increases the *average* consumption per capita, economic theory leads us to say that society has been made better off—but this is not completely correct.

Consider an (admittedly extreme) case in which the sole beneficiary of technological change is society's richest person, who makes $100 billion per year, and the costs fall on one million low-wage workers, who each make $16,000 per year. If the rich person gains $5 billion from technological change, while costs of $4,000 fall on each of the one million low-wage workers (for a total of $4 billion in costs), society as a whole gains $1 billion in overall consumption. However, as explained below, this $1 billion gain could be associated with *a loss in overall utility* in society.[23]

The gain to the rich person in our example represents 5 percent of his or her annual income, and with such a huge income to begin with, the addition of $5 billion may not add much to this person's utility. The loss of $4,000 per worker for each of one million workers is equal to 25 percent of their annual income, and the associated loss of utility may—in the aggregate—be larger than the relatively small gains

[23]See Richard Layard, "Happiness and Public Policy: A Challenge to the Profession," *Economic Journal* 116 (March 2006): C24–33, for a discussion of recent psychologically based findings that people in wealthy economies are *loss-averse*; that is, their gains in utility from increases in income are smaller than their losses in utility from reduced income.

EMPIRICAL STUDY

Estimating the Labor Demand Curve: Time Series Data and Coping with "Simultaneity"

When a proposed labor market policy increases the cost of labor, we frequently want economists to tell us more than "It will reduce employment." We want to know *how much* employment will be affected! Thus, for practical purposes, it is very helpful to have estimates of the elasticity of demand for labor.

Estimating the elasticity of demand for labor is actually very difficult, which helps account for how few studies of demand elasticity were cited in Table 4.1. First, we can only obtain credible estimates if we have data on wages and employment for groups of workers who are reasonably homogeneous in terms of their job requirements, their substitutability with capital, and the characteristics of product demand facing their employers. Given the diversity of firms that hire workers in a given occupation (security guards, for example, are hired by retailers, schools, and movie stars), homogeneity often requires analyzing groups so narrow that data are very difficult to obtain.

A second problem in estimating labor demand curves is that wages and employment are determined *simultaneously* by the interaction of supply and demand curves, and both curves show a (different) relationship between wages and employment. If we gather data just on wage and employment levels, we will not be able to tell whether we are estimating a demand curve, a supply curve, or neither! Consider Diagrams #1 and #2, which show wage (W) and employment (E) outcomes in the market for an occupation.

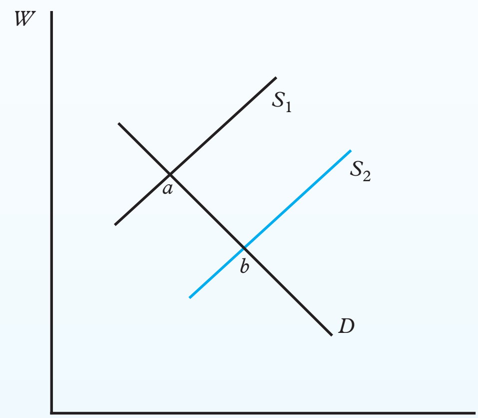

Diagram #1

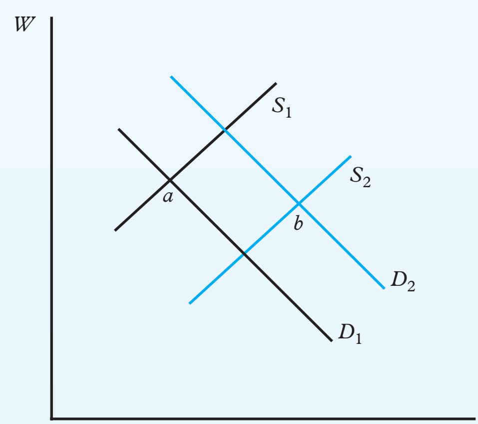

Diagram #2

What we hope to do is illustrated in Diagram #1. There, the labor demand curve remains unchanged, but the supply curve shifts for some reason. All that is observed by the researcher are points a and b, but connecting them traces out the demand curve (of course, credible estimates would require many more than two observations).

Thus, if the demand curve is not shifting, we can "identify" it if we can observe a shifting supply curve. In reality, however, both supply and demand curves can shift over time (see Diagram #2). When both shift, drawing a line between points *a* and *b* traces out neither a supply nor a demand curve. How can we identify the demand curve when both are likely to be shifting?

First, we must have access to variables that cause the demand curve to shift; if we can *control* for factors that shift the demand curve over time, we—in a statistical sense—can shift it back to its original position and create a situation like that in Diagram #1.

Second, for the condition in Diagram #1 to be met, we must also find at least one variable that *shifts the supply curve but does not affect demand*. (Some variables, like real income levels, can theoretically affect both labor demand *and* labor supply curves. If all our "shift" variables are expected to affect both curves, we are back in the situation depicted by Diagram #2, where we cannot distinguish between the two curves!)

A study of the demand for coal miners in Britain (cited at the bottom of Table 4.1) offers an example of how to estimate a labor demand curve for a specific occupation. The occupation is found in one industry, which is very homogeneous in terms of product demand and employer technology, and time-series data on wages and employment were available for several years (the study used data from the 1950–1980 period). The researchers were able to gather data on factors that were expected to shift the labor demand curve (the price of oil, for example, which is a substitute for coal in generating electricity). They also had access to data on variables that were expected to shift the supply curve—including those (such as wages in alternative jobs miners might choose) that were expected to shift *only* the supply curve. The researchers were thus able to identify the labor demand function, and their use of regression analyses suggested that the labor demand elasticity (of employment changes with respect to wage changes) in British coal mining was −1.0 to −1.4.

Source: Alan A. Carruth and Andrew J. Oswald, "Miners' Wages in Post-War Britain: An Application of a Model of Trade Union Behaviour," *Economic Journal* 95 (December 1985): 1003–1020.

to the rich person. The only way to ensure that society as a whole gains (in terms of *utility*) in this case is to require the gainer to compensate all the losers. If the person who gained were required to distribute $4 billion of the gains to those who bore the costs of change, the workers would end up being no worse off, and the gainer would still be ahead because of the $1 billion he or she gets to keep. Thus, after the compensation of losers takes place, a normative condition put forth in chapter 1 would hold: some would gain from technological change, and no one would lose.

To restate the normative principle outlined in chapter 1, *we can be sure that society gains from any economic transaction—technological change in this case—only when all those who lose from it are fully compensated*. Because most technological change occurs through decisions made by the millions of firms in the marketplace, what is needed to compensate those who lose jobs as a result of these decisions is a broad set of social insurance policies that can assist displaced workers. Unemployment insurance, for example, can support workers during their search for new jobs, and wage supplements of one form or another can minimize their loss of income if they have to take a lower-wage job; training programs can assist with the acquisition of new skills; government employment centers or online "job banks" can help the workers locate job openings; public wage subsidies can be paid to firms that add new workers; and in some countries, the government operates as the "employer of last resort"—putting to work those who cannot find new jobs. While we analyzed wage subsidies to employers in chapter 3, we will analyze the effects of many of these other programs later in this text.

Review Questions

1. Suppose that the government raises the minimum wage by 20 percent. Thinking of the four Hicks–Marshall laws of derived demand as they apply to a particular industry, analyze the conditions under which job loss among teenage workers in that industry would be smallest.
2. California employers of more than 50 workers are now required to offer paid family leave for workers with newborn children. Under this law, businesses with more than 50 workers are required to hold a job for a worker who goes on paid leave for up to six weeks. When on leave, workers receive 55 percent of their normal pay. What are the likely responses on the demand (employer) side of the labor market? Include in your analysis a consideration of factors that would affect the *size* of these responses.
3. The federal government, in an effort to stimulate job growth, passes a law that gives a tax credit to employers who invest in new machinery and other capital goods. Applying the concepts underlying cross-elasticities, discuss the conditions under which employment gains in a particular industry will be largest.
4. The public utilities commission in a state lifts price controls on the sale of natural gas to manufacturing plants and allows utilities to charge market prices (which are 30 percent higher). What conditions would minimize the extent of manufacturing job loss associated with this price increase?
5. Many employers provide health insurance for their employees, but others—primarily small employers—do not. Suppose that the government wants to ensure that all employees are provided with health insurance coverage that meets or exceeds some standard. Suppose also that the government wants employers to pay for this coverage and is considering two options:

   Option A : An employer not voluntarily offering its employees acceptable coverage would be required to pay a tax of X cents per hour for each labor hour employed. The funds collected would support government-provided health coverage.

   Option B : Same as option A, except that the government-provided coverage would be financed by a tax collected as a fraction of the employer's total revenues.

   Compare and contrast the labor market effects of each of the two options.
6. In 1942, the government promulgated regulations that prohibited the manufacture of many types of garments by workers who did the sewing, stitching, and knitting in their homes. If these prohibitions are repealed so that clothing items may now be

made either by workers in factories or by independent contractors doing work in their homes, what effect will this have on the labor demand curve for *factory workers* in the garment industry?

7. Briefly explain how the following programs would affect the elasticity of demand for labor in the steel industry:
   a. An increased tariff on steel imports.
   b. A law making it illegal to lay off workers for economic reasons.
   c. A boom in the machinery industry (which uses steel as an input)—causing production in that industry to rise.
   d. A decision by the owners of steel mills to operate each mill longer than has been the practice in the past.
   e. An increase in the wages paid by employers in the steel industry.
   f. A tax on each ton of steel produced.

Problems

1. Suppose that the demand for dental hygienists is $L_D = 5{,}000 - 20W$, where L = the number of dental hygienists and W = the daily wage. What is the own-wage elasticity of demand for dental hygienists when W = \$100 per day? Is the demand curve elastic or inelastic at this point? What is the own-wage elasticity of demand when W = \$200 per day? Is the demand curve elastic or inelastic at this point?
2. Professor Pessimist argues before Congress that reducing the size of the military will have grave consequences for the typical American worker. He argues that if 1 million individuals were released from the military and were instead employed in the civilian labor market, average wages in the civilian labor market would fall dramatically. Assume that the demand curve for civilian labor does not shift when workers are released from the military. *First*, draw a simple diagram depicting the effect of this influx of workers from the military. *Next*, using your knowledge of (i) the definition of the own-wage elasticity of labor demand, (ii) the magnitude of this elasticity for the economy as a whole, and (iii) the size of civilian employment in comparison with this flood from the military, graph these events and estimate the magnitude of the reduction in wages for civilian workers as a whole. Do you concur with Professor Pessimist?
3. Suppose that the demand for burger flippers at fast-food restaurants in a small city is $L_D = 300 - 20W$, where L = the number of burger flippers and W = the wage in dollars per hour. The equilibrium wage is \$4 per hour, but the government puts in place a minimum wage of \$5 per hour.
   a. How does the minimum wage affect employment in these fast-food restaurants? Draw a graph to show what has happened, and estimate the effects on employment in the fast-food sector.
   b. Suppose that in the city above, there is an uncovered sector where $L_S = -100 + 80W$ and $L_D = 300 - 20W$, before the minimum wage is put in place. Suppose that all the workers who lose their jobs as burger flippers due to the introduction of the minimum wage seek work in the uncovered sector. What happens to wages and employment in that sector? Draw a graph to show what happens, and analyze the effects on both wages and employment in the uncovered sector.

4. The following table gives the demand for labor at Homer's Hideaway, a motel in a small town.

Wage ($)	Number of Hours
10	2
8	3
6	4
4	5
2	6

   a. Draw the demand for labor curve.
   b. Calculate the wage elasticity of demand at points along the demand curve. Indicate whether the elasticity is elastic, inelastic, or unitary elastic.
   c. As you slide down along the demand curve, does the demand curve become more or less elastic?

5. Union A faces a demand curve in which a wage of $4 per hour leads to demand for 20,000 person-hours, and a wage of $5 per hour leads to demand for 10,000 person-hours. Union B faces a demand curve in which a wage of $6 per hour leads to demand for 30,000 person-hours, whereas a wage of $5 per hour leads to demand for 33,000 person-hours.
   a. Which union faces the *more* elastic demand curve?
   b. Which union will be more successful in increasing the total income (wages times person-hours) of its membership?

6. Calculate the own-wage elasticity of demand for occupations a, b, and c below. E_D and W are the original employment and wage. E'_D and W' are the new employment and wage. State whether the demand is elastic, inelastic, or unitary elastic.
   a. $\%\Delta E_D = 5$, $\%\Delta W = -10$
   b. $E_D = 50$, $W = 7$
      $E'_D = 40$, $W' = 8$
   c. $E_D = 80$, $W = 8$
      $E'_D = 100$, $W' = 6$

7. When the cost of dough-making machines fell by 10 percent, the demand for assistant bakers fell by 15 percent. What is the cross-wage elasticity of demand for assistant bakers in this case? Are assistant bakers and dough-making machines gross substitutes or gross complements?

Selected Readings

Card, David, and Alan B. Krueger. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton: N.J.: Princeton University Press, 1995.

Cunningham, Wendy. *Minimum Wages and Social Policy: Lessons from Developing Countries*. Washington, D.C.: World Bank, 2007.

Hamermesh, Daniel S. *Labor Demand*. Princeton, N.J.: Princeton University Press, 1993.

Kennan, John. "The Elusive Effects of Minimum Wages." *Journal of Economic Literature* 33 (December 1995): 1950–1965.

Neumark, David, and William L. Wascher. *Minimum Wages*. Cambridge, Mass.: MIT Press, 2010.

"Review Symposium: *Myth and Measurement: The New Economics of the Minimum Wage*, by David Card and Alan B. Krueger." *Industrial and Labor Relations Review* 48 (July 1995).