CHAPTER 17

Wages, Rent, Interest, and Profit

How are wages, rent, interest, and profit determined?

Chapter Outline

Learning Objectives

- Investigate a worker's decision concerning how many work hours to supply.
- Examine the income and substitution effects of a higher wage rate and whether the net result of a wage increase involves a worker supplying more work hours.
- Analyze the general level of wage rates and why wages differ among jobs.
- Define what economists mean by the term *rent*.
- Explore selling or monopoly power in input markets and show how unions attempt to exercise such power in labor markets.
- Explain how the interest rate is determined through the interplay of the supply of and demand for capital.
- Overview why interest rates differ across specific credit markets.
- Describe the net present value method for analyzing the desirability of undertaking investment projects.

The general principles discussed in Chapter 16 apply to the analysis of the market for any type of input. However, some special issues arise in conjunction with particular input markets, and these deserve further attention. In this chapter, we extend the general analysis to specific input markets to see how wages, rent, interest, and profit are determined. Because labor earnings account for about 75 percent of total U.S. national income, we continue to emphasize labor markets.

17.1 THE INCOME-LEISURE CHOICE OF THE WORKER

In our discussion of consumer demand in Chapters 3 and 4, we assumed the consumer's income to be fixed. For most people, however, income is not fixed; among other things, it depends on the decision about how much time to spend working. To investigate the worker's decision concerning how many work hours to supply, we will assume that the individual worker is paid a fixed hourly wage and can work any number of hours desired at that wage.

LEISURE
the portion of a worker's time when he or she is not receiving compensation from an employer

This analysis utilizes the indifference curve–budget line technique developed in Chapters 3 and 4 for the analysis of consumer choice. In Figure 17.1, the vertical axis measures the individual worker's total weekly income, and the horizontal axis, from left to right, measures the worker's leisure time. The term **leisure** refers to the portion of the worker's time when he or she is not receiving compensation by an employer. Any worker has 168 hours a week available, 24 hours a day, 7 days a week. We divide this time into two mutually exclusive categories, work and leisure. Working time plus leisure time per week must equal 168 hours. In Figure 17.1, Z is the total time available, and the point L_1 indicates that the individual

FIGURE 17.1

Income-Leisure Choice of the Worker
Measuring leisure from left to right is the same as measuring hours worked from right to left from point Z. The budget line AZ has a slope equal to the hourly wage rate. The optimal point is E, with the individual working ZL_1 hours and earning \$800 per week.

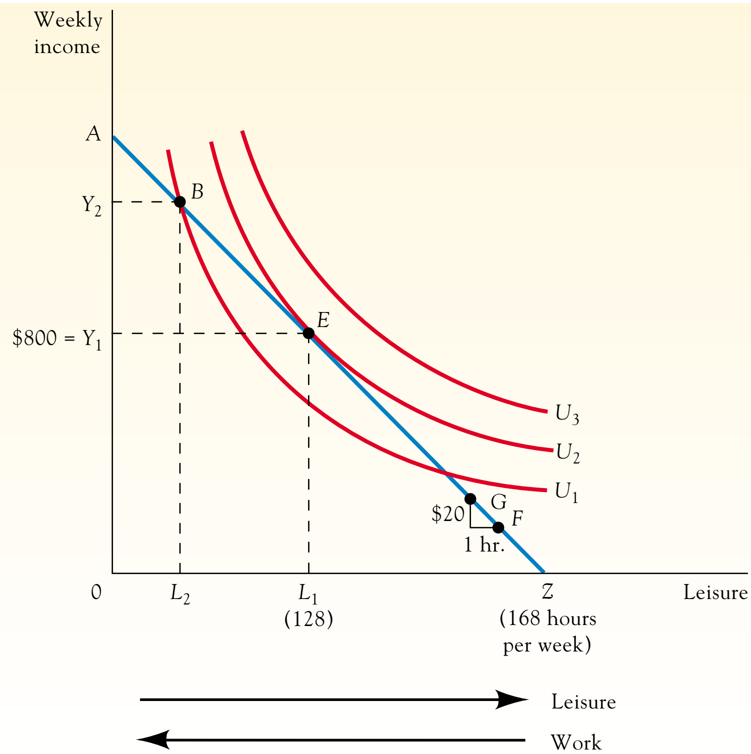

consumes L_1 hours of leisure and supplies the remaining time, ZL_1 hours, to an employer as work effort. Measuring leisure in this way, from left to right, is therefore equivalent to measuring work effort (hours of labor supplied) to the left from point Z.

With income and leisure measured on the axes, a worker's budget line reflects the wage rate received per hour of work provided. In the diagram, the budget line is AZ, and it shows the combinations of income and leisure available to the worker. Note that the more hours the individual works (that is, the less leisure time consumed), the higher the worker's income. For example, if ZL_1 hours are worked (leisure of L_1), income is Y_1, but if work effort increases to ZL_2 hours (leisure decreases to L_2), income rises to Y_2. If no work is done (point Z), income is zero since we assume that the worker has no nonlabor sources of income. Also observe that *the slope of this budget line equals the worker's wage rate*. For example, a movement from point F to point G indicates that the worker is providing one more hour of labor (giving up one more hour of leisure) and in return receives an additional $20 in wage income. Thus, the hourly wage rate is $20.

To the worker, both income and leisure are desirable economic goods. For a given amount of leisure, more income is preferred to less, and for a given amount of income, more leisure is preferred to less. As we pointed out in Chapter 3, whenever economic goods are measured on the axes, indifference curves have their normal shapes (downward-sloping, convex, and nonintersecting). The slope of an indifference curve relating income and leisure measures the willingness of the worker to give up leisure for more money income and therefore indicates the relative importance of these two goods to the individual. Figure 17.1 shows three of the worker's income-leisure indifference curves.

The optimal point for the worker is E since this point represents the most preferred combination of income and leisure from among those on the budget line. Work effort is ZL_1 (40 hours per week) and weekly income is Y_1 ($800). As usual, the optimal point involves a tangency between the budget line and an indifference curve. In this case the tangency indicates that the subjective marginal valuation of the worker's own leisure time is equal to the market valuation of the individual's work time, the wage rate. At point E the worker must be paid at least $20 to give up one more hour of leisure, since the marginal rate of substitution (*MRS*) between income and leisure is $20 per hour at that point.

Note that the optimal point in Figure 17.1 does not result in the worker's income being maximized. For example, the worker could earn a higher income by working longer hours, such as at point B. But the extra Y_1Y_2 income is worth less than the L_1L_2 hours of leisure time that must be sacrificed to earn the additional income, as shown by the fact that point B is on a lower indifference curve than point E. The worker thus is better off by forgoing the higher income (point B) for the sake of more leisure (point E). At point B, the marginal value of leisure exceeds the opportunity cost of leisure (forgone income); thus the worker will increase his or her utility by pursuing additional leisure until the *MRS* between income and leisure just equals the wage rate.

Is This Model Plausible?

This analysis rests on the assumption that the worker is able to choose how many hours to work. A common objection to the model is that workers don't really have the ability to vary their work hours. Most employment contracts specify that employees will work a certain length of time—perhaps 35 hours per week, for example. Even if an individual employee prefers a 30-hour week combined with a one-seventh reduction in pay, the employer may not provide that option. In short, the argument holds that the workweek is fixed by the employer and beyond the control of the individual worker. There is an element of truth to this point, but it is a less serious criticism than it might appear on first glance.

One way to justify the assumption of variable work effort is to recognize that most workers can, in fact, exercise some degree of control over how much they work, although perhaps not on a daily or weekly basis. Overtime, vacation leave, leaves without pay, moonlighting, sick leave, and early retirement are options available to many workers. At a more basic level, each person has a wide range of options in selecting a job in the first place. Some jobs entail long hours, some short, and some permit considerable variation in work effort. For example, it is not unusual for entrepreneurs to average 80-plus hours of work per week in start-up firms striving to bring new products to market. Many of the same entrepreneurs could opt for less demanding jobs in the nonprofit, public, or private sectors (at established firms, for example) that, while requiring fewer hours of work per week, are associated with less promising financial rewards.

Another justification for the variable work effort assumption is that at a more fundamental level the analysis may still be valid for many purposes, even if it is impossible for a worker to vary his or her workweek even slightly. Although the employer fixes the workweek, consider the economic factors that determine the level at which it is set. Employers are profit maximizers, and it is therefore in their interest to cater to the preferences of workers, just as they are led by profit motives to cater to the preferences of consumers. If a firm's employees prefer a 30-hour workweek, but the firm requires them to work 35 hours, it will lose workers to other firms that do a better job of satisfying the employees' preferences. Competition for workers thus leads firms to set workweeks that correspond to worker preferences. So, the assumption that workers can choose how much they will work should yield a reasonably correct analysis, although it does not precisely describe reality.

One qualification should be mentioned. Employers have an incentive to cater to workers' preferences on average, but not necessarily to each employee's preferences. The technology of production requires most firms to have a common workweek for all employees (though it can, and does, differ among firms). A fixed workweek means that workers with preferences different from the group average will not like the workweek schedule. Thus, our model may not be strictly accurate for any specific worker, but it does provide a basis for analyzing work effort decisions involving groups of workers.

17.2 THE SUPPLY OF HOURS OF WORK

Will workers work longer hours at a higher wage rate? The work-leisure choice model can help answer this question. Figure 17.2 examines the effect of a higher wage rate for a particular worker. When the hourly wage rate is $20, the budget line is *AZ*, and the worker's preferred point is *E* with ZL_1 work hours supplied. Remember that we measure work effort from right to left in the diagram. If the hourly wage rate rises to $25, the budget line rotates about point *Z* and becomes steeper. The new budget line is *A′Z* with a slope of $25 per hour. Note that at the higher wage rate income is greater for any level of work effort. Given the specific preferences of this worker, when confronted with the higher wage rate, the new optimal point *E′* involves an increase in hours of work, from ZL_1 to ZL_2.

Does a higher wage always lead a worker to work more? The answer is no, and we can see why by considering the income and substitution effects associated with a change in the wage.

The *substitution effect* of a higher wage rate encourages a worker to supply more hours of labor. When the hourly wage rate rises from $20 to $25, the sacrifice for leisure consumption is greater, since each hour of leisure now means giving up $25 in income instead of $20. Since leisure has become more costly in terms of income lost, the worker is encouraged to substitute away from leisure toward income—that is, to work more.

An *income effect* is also associated with a higher wage rate but has an opposite result on work effort from the substitution effect. A wage increase makes the individual better off, permitting the worker to reach a higher indifference curve. A higher real income tends to

FIGURE 17.2

Worker's Response to a Change in the Wage Rate
A higher wage rate pivots the budget line from *AZ* to *A′Z* and work effort increases from ZL_1 to ZL_2. We show the income and substitution effects of the change in the wage rate by using the hypothetical budget line *HH′* that is parallel to the *A′Z* budget line but just tangent to the initial indifference curve U_1. The substitution effect, L_1L_3, involves more work, but the income effect, L_3L_2, involves less. In this case the combined effect implies greater work effort at the higher wage rate.

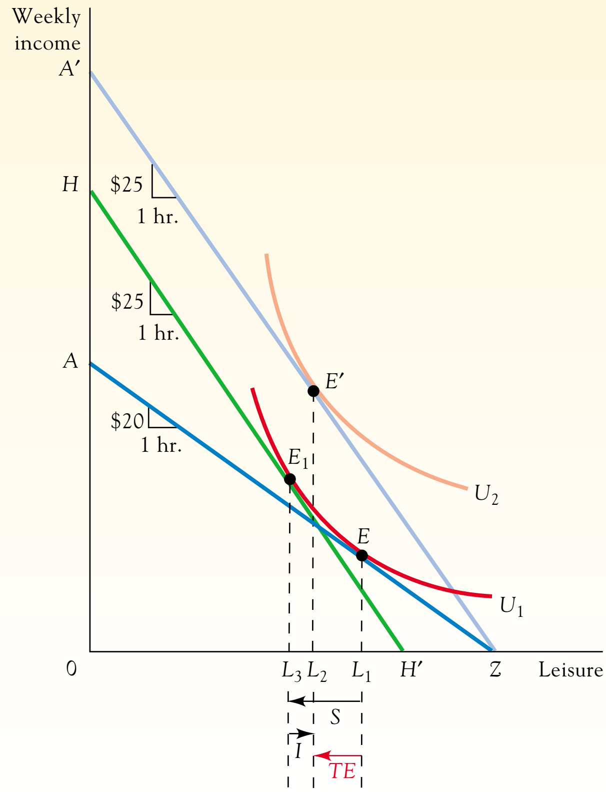

increase the consumption of all normal goods, and leisure for most people is a normal good. The income effect of a wage rate increase thus encourages the consumption of leisure and leads the worker to work less. Because of the higher wage, the worker can afford to work less; it is possible to work fewer hours and still achieve a higher money income than before the wage increase.

Figure 17.2 shows that *the substitution effect of a higher wage rate encourages more work, and the income effect encourages less work.* The hypothetical budget line *HH′* is drawn tangent to the worker's original indifference curve U_1. Its slope reflects the higher wage rate of $25. The substitution effect is shown as the movement along U_1 from *E* to E_1. Because leisure has become more costly, the worker consumes less, and work effort increases from ZL_1 to ZL_3. The income effect is shown as the movement from E_1 to *E′* when we allow the individual to move from the *HH′* budget line to the parallel *A′Z* budget line, reflecting the increase in real income associated with the rise in the wage rate. Since leisure is a normal good, the income effect involves more leisure, from L_3 to L_2, which is the same as saying it encourages less work—that is, ZL_2 instead of ZL_3. The total effect of the higher wage rate is the sum of the income and substitution effects. Although these effects operate in opposite directions, in this case the substitution effect is larger, so the total effect is an increase in work hours from ZL_1 to ZL_2.

Is a Backward-Bending Labor Supply Curve Possible?

For the worker whose preferences are depicted in Figure 17.2, the supply curve of work hours slopes upward, at least between wage rates of $20 and $25, since a higher wage leads to a greater quantity of labor supplied. This outcome need not always be the case, however. The income effect of a higher wage may be larger than the substitution effect, resulting in a reduction in work hours at higher pay. The intuition behind such an alternative outcome is straightforward. Beyond some point individuals may prefer to work a little less, take some time off, and enjoy the higher income made possible by a higher salary. For example, 30 hours a week at a wage of $25 per hour means more income and more leisure time with which to enjoy the income when compared to 40 hours at a wage of $5 per hour.

Figure 17.3a shows an individual who will choose to work longer hours when the wage increases from $20 to $25 but will then decide to work somewhat less when the wage rises again. When the wage increases from $20 to $25, work hours increase from ZL_1 to ZL_2, but at a wage of $30, hours worked falls to ZL_3. (We have not separated the income and substitution effects in the diagram, but you may wish to do so.) Figure 17.3b shows the same information plotted as a labor supply curve of weekly work hours. The supply curve slopes upward between wage rates of $20 and $25 but bends backward as the wage rate rises further.

A supply curve of work hours can thus be backward-bending beyond some wage rate. Note that a backward-bending labor supply curve does not depend on an unusual set of circumstances, as does an upward-sloping demand curve; rather, *it just requires that the normal income effect of a higher wage rate exceed the substitution effect*. Leisure is a normal good, so the income and substitution effects always work in opposing directions in the case of a labor supply curve because the worker is a *seller* of labor services. A rise in the price of something an individual sells (such as labor services) has a positive effect on the individual's income and thus leads to greater consumption of all normal goods, including leisure. (In contrast, when the price of something you purchase as a *buyer* increases, the effect on your income is negative. The negative income effect reinforces, rather than counteracts, the substitution

FIGURE 17.3

An Individual Worker's Weekly Supply of Work

(a) An individual's choices of how much to work at three different wage rates are represented by *E*, *E*′, and *E*″. (b) These labor supply choices are plotted as the supply curve of weekly work hours. When the income effect exceeds the substitution effect, the supply curve becomes backward-bending.

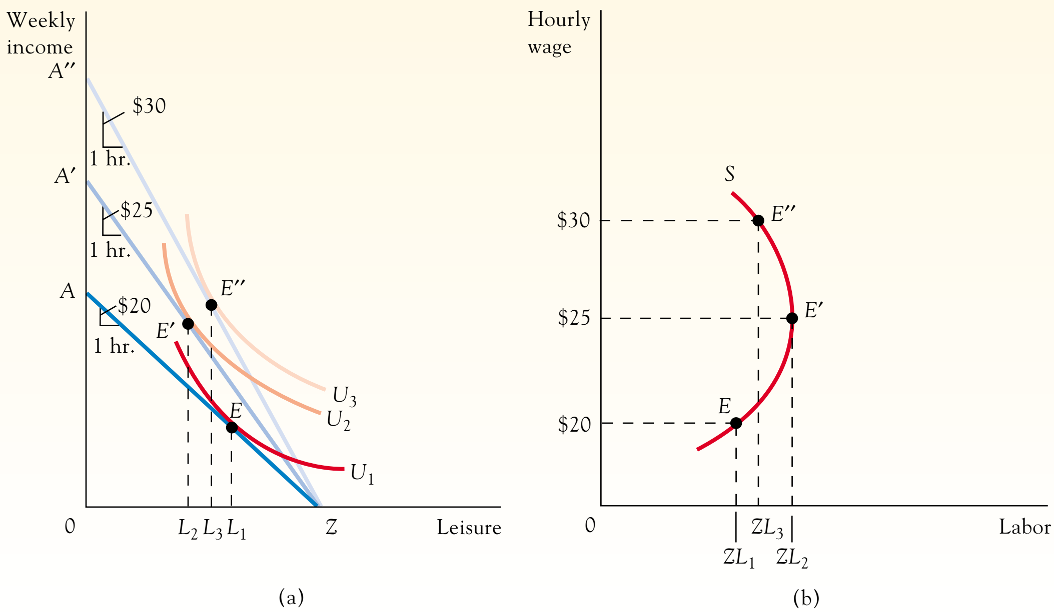

effect for a normal good.) Moreover, the income effect of a change in pay is likely to be large relative to the substitution effect since most income derives from providing labor.

The Market Supply Curve

To go from an individual's supply curve of work hours to the market supply curve, we need only horizontally sum the responses of all workers competing in a given labor market. Thus, the market supply curve can also slope upward, bend backward, or show a combination of the two, as shown in Figure 17.3b. Theoretical considerations alone do not permit us to predict the exact shape of the market supply curve. Empirical evidence, however, suggests that it slopes upward (at least for wage rates near present levels), with an elasticity somewhere between 0.1 and 0.3.[1] An elasticity of 0.1 means that a 10 percent increase in the wage rate increases the quantity of labor supplied by 1 percent. Such an inelastic supply curve slopes upward and is almost vertical.

A highly inelastic aggregate supply curve, as suggested by the empirical evidence, appears plausible. Casual observation suggests that the amount of time most people work stays the same over moderate periods of time despite changes in wage rates. If individual supply curves were sharply upward-sloping or backward-bending, we would see substantial changes in individuals' work hours in jobs where market forces have produced large wage rate changes. Work hours per worker in most jobs seem to be quite stable over time, suggesting that the effect of wage rate changes on work hours is not pronounced. This result does not mean that people are completely unresponsive to wage rate changes, only that the responses that occur are modest.

When should we use the aggregate labor supply curve? In the preceding chapter we emphasized that the labor supply curve *to a particular industry or occupation* is likely to slope upward and be quite elastic. We must distinguish that type of labor supply curve, however, from the one discussed here. In looking at the supply to an industry or occupation, we see that the quantity of labor services *can* increase sharply with a rise in the wage rate in that job, but the increase results mainly from an influx of workers from other jobs or industries and not from a change in work hours of current workers. The supply curve of labor to a specific job or industry therefore depends mainly on how the number of workers varies with wage rates in those specific occupations.

The aggregate labor supply curve is used in cases when the movement of workers between jobs is not likely to be significant but the possible change in the hours supplied by workers in their current jobs is. For example, how would a 10 percent increase in all wage rates affect the total quantity of labor supplied? Since all jobs will pay proportionately more, people will have little incentive to change jobs. The only way the total labor supply will increase, therefore, is if people work longer hours (which includes the possibility that some people will enter the labor force for the first time at the higher wage rate). For this type of aggregate analysis, which involves the total quantity of labor supplied across all industries, the market supply curve of work hours is appropriate. We discuss an example using this supply curve in the next section and offer more examples in Chapter 18.

17.3 The General Level of Wage Rates

In analyzing the determination of wage rates, it is convenient to divide the subject into two parts: determination of the general level of wage rates and consideration of why they differ among jobs. In this section we focus on the factors that influence the level of real wages, or

[1]For a representative survey of the empirical evidence see U.S. Congressional Budget Office, *An Analysis of the Roth-Kemp Tax Cut Proposal* (Washington, D.C.: U.S. Government Printing Office, 1978).

the average wage rate; we defer until the next section a discussion of the factors that cause a variation in wage rates around the average.

Supply and demand are still the applicable concepts for investigating the level of wage rates. The supply curve of labor indicates the total quantity of labor supplied by all persons at various wage levels. The appropriate supply concept is therefore the aggregate supply curve of work hours discussed in the previous section. This supply curve is probably quite inelastic, like curve S_1 in Figure 17.4.

The aggregate demand curve for labor reflects the marginal productivity of labor to the economy as a whole. Indeed, it is convenient to think of the wage rate as being paid in units of national output (each unit composed of the combination of goods consumed by the average person) to emphasize that we are dealing with the level of *real* wage rates. In constructing the demand curve relevant for a particular time period, the following factors are held constant: capital (including land, buildings, and equipment), technology, and the skills, knowledge, and health of the labor force. If these factors are fixed, an increase in the total quantity of labor is subject to the law of diminishing marginal returns. Consequently, the aggregate marginal product curve slopes downward and is the aggregate demand curve for labor.

At any particular time, if the supply curve is S_1 and the demand curve is D_1 as in Figure 17.4, then the (average) real wage rate is w_1 and employment is L_1. At the high degree of aggregation used in this analysis, the model is necessarily abstract and ignores a multitude of factors that could influence the positions of the demand or supply curves. Yet it highlights the importance of the productivity of labor in determining the level of wages. National output and national income are two sides of the same coin, and with labor receiving about 75 percent of national income (output), it is clear that the factors that determine the output level produced from a given quantity of labor play a central role in the analysis. These factors are primarily (but not exclusively) technology, the skill level of the labor force, and the amounts of other inputs, which in this example we refer to as *capital*.

This concept explains why real wages are so much higher in the United States than in less developed countries: the (marginal) productivity of labor is greater. Marginal productivity is higher because of the factors determining the position of the demand curve: capital, technology, and skills. In U.S. manufacturing industries the amount of capital per worker is about $125,000, contributing to high average and marginal productivity of labor. Technological knowledge is superior in this country, and the U.S. labor force is well educated and

FIGURE 17.4

Determination of the General Wage Level

The aggregate demand curve for labor interacts with the aggregate supply curve to determine the general wage level. Over time, normally both supply and demand increase. If demand increases faster than supply, wage rates tend to rise over time.

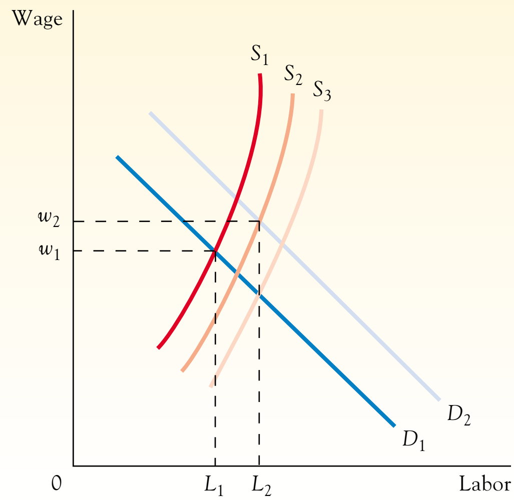

highly skilled by international standards. Consequently, output per worker in the United States is much higher. Other factors, such as climate, efficiency of the economic system in allocating the available resources, the degree of political stability, and attitudes toward work, play a role in determining national real wage levels. We can analyze their influence on the productivity of the labor force by examining the way they affect the positions of the demand or supply curves.

Over time, both the demand and supply curves shift. Because of saving and investment, the amount of capital tends to grow over time, and this growth shifts the demand curve outward: more capital per worker at each employment level increases marginal productivity. Similarly, technological progress and improvements in the labor force's skill level also increase marginal productivity. On the supply side, the supply curve shifts to the right as the population grows over time. Whether these changes lead to a rising or falling real wage depends on how much demand shifts relative to supply. In Figure 17.4, if demand increases relatively more, to D_2, while supply increases to S_2, the real wage goes up to w_2. Such an outcome has been the experience in most industrial countries over the past two centuries. Because of capital accumulation, technological progress, and the availability of workers with greater skills, the demand for labor has increased faster than the supply, pulling real wages and living standards up in the process. Unfortunately, this happy outcome is not inevitable.

APPLICATION 17.1 THE MALAISE OF THE 1970S

To most U.S. consumers, who had become accustomed to uninterrupted economic progress since the 1930s, the decade of the 1970s was a shock. Foremost among the decade's economic woes was a failure of many families' incomes to keep pace with inflation. Median family income in constant dollars increased by 33 percent in the 1950s and again in the 1960s. In the 1970s, however, real median family income increased less than 1 percent before taxes. After taxes were paid, many families lost ground.

What went wrong with the U.S. economy in the 1970s? Observers point to a variety of possible culprits: inflation, higher energy prices, reduced national saving and investment, Japanese imports, increased government taxes and regulation, and others. The basic problem was a failure of real wage rates to rise as much as they had in previous decades. Thus, an explanation of the situation should focus on what happened to the supply of and demand for labor. Focusing on supply and demand conditions doesn't rule out the possibility that some of the factors mentioned above could have caused wages to grow less rapidly, but it suggests that a major role may have been played by a factor that most commentators fail to recognize.

During the 1970s the labor force in the United States grew dramatically. Following increases of 12 percent in the 1950s and 19 percent in the 1960s, the labor force grew by 29 percent in the 1970s—the largest increase on record in a single decade. Referring to Figure 17.4, we can see the consequences of such a large increase in the labor force. If the normal increase in supply due to labor force growth over a decade is from S_1 to S_2 (by about 15 percent, for example), then the increase during the 1970s would be shown as the much greater shift in labor supply from S_1 to S_3. With the same increase in demand from D_1 to D_2, the abnormally large increase in labor supply means that over that period wage rates will not rise as much as usual, or might not rise at all.

Table 17.1 provides data on the growth in the labor force and real hourly earnings for the 1950 to 1980 period. A large increase in the labor force in the 1970s was associated with a 2.3 percent decline in real wages over the entire decade. Simple supply and demand forces appear to have been at work.

The increase in the labor force in the 1970s was primarily the result of two factors. First, the number of young persons reaching working age during the decade was unusually large. The "baby boom" generation of the

TABLE 17.1

SIZE OF LABOR FORCE AND WAGE RATES, 1950–1980

	Labor Force (Millions)	Increase over Previous Period (Percentage)	Index of Real Hourly Earnings (1971 = 100)	Increase over Previous Period (Percentage)
1950	62.2	–	64.0	–
1960	69.6	11.9	81.4	27.2
1970	82.8	19.0	95.7	17.6
1980	106.9	29.1	93.5	−2.3

Source: *Economic Report of the President*, 1982.

1950s had grown up. Second, the share of females in the labor force increased from 43 percent in 1970 to 52 percent in 1980. (The labor force participation rate of females had never exceeded 40 percent before 1966.) Each of these factors was significant in itself, but together they spelled an unusually large increase in labor supply. Labor markets, however, adjusted to accommodate this influx of workers, and employment increased by millions more than it had in any previous decade. The relatively large shift in supply meant that wage rates rose less than in previous decades.

What roles did inflation, energy prices, decreased saving, and the other factors so frequently mentioned play in affecting real wage rates? Several caused the demand for labor to grow less rapidly than it had in previous decades and thus reinforced the tendency for wage rates to rise more slowly. For example, a reduced rate of investment and a higher price of oil—an input purchased in large quantities from other countries—both depressed the rate of increase in labor demand. Theory suggests, however, that some of the other factors frequently mentioned, like inflation and higher imports of consumer goods from Japan, would have little effect, since they do not significantly affect productivity. The exact quantitative contribution of each of these factors to the slowdown in the growth of real wages is an unresolved issue and has been the subject of some debate among economists.

The role of other factors notwithstanding, the massive increase in the number of workers in the labor force appears to have had a significant impact on real wages during the 1970s in the United States. Moreover, the supply–demand model appears to offer the correct approach to the question of what caused the decline in real wages during this period.

17.4 WHY WAGES DIFFER

From the market forces that determine the general level of wage rates, we turn now to the question of why there is such wide variation in the wages received by different individuals. In Chapter 16, we explained why wage rates across firms or industries tend to equalize. That analysis depended on the assumptions that workers were identical and that they evaluated the desirability of the jobs only in terms of the money wage rates. Dropping these assumptions, as we must for a fuller understanding of labor markets, suggests that wage rates can differ among jobs and among people employed in the same line of work. People are different in the type of work they are both able and willing to perform, and these differences on the supply side of labor markets produce differences in wage rates.

Perhaps the best way to see what is involved is to take a hypothetical, but plausible, example. Figure 17.5a shows the labor market for clerks, and Figure 17.5b shows the labor market for engineers. Under competitive conditions, the intersection of supply and demand curves in each market yields a wage rate for engineers that is twice that for clerks. We suggest that these markets are in full equilibrium with no tendency for the wage rates to equal-

FIGURE 17.5

Equilibrium Wage Differences

(a) The labor market for clerks results in wage rate w_C. (b) The labor market for engineers results in wage rate w_E. Although the wage rate for engineers is higher, there is no tendency for these wage rates to equalize: they are equilibrium wage differences.

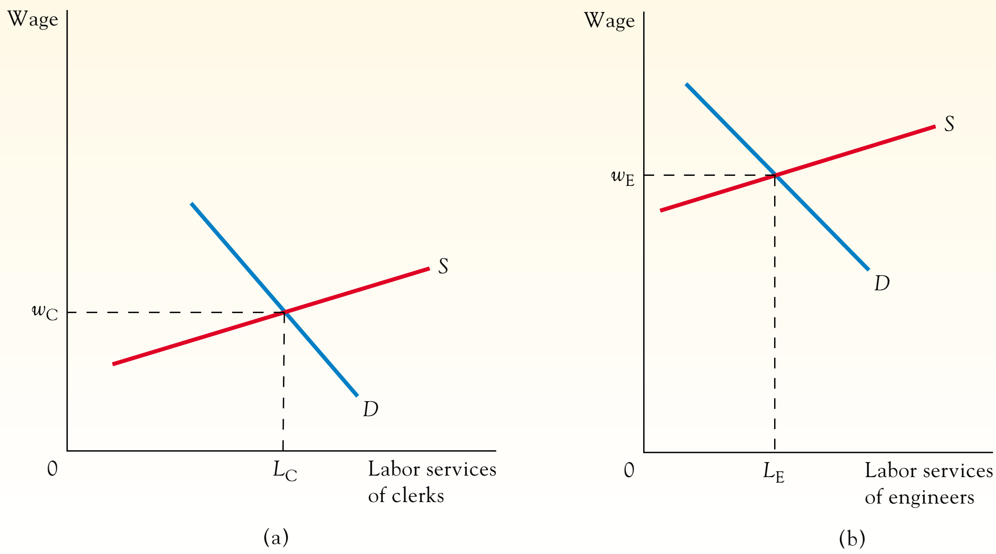

ize. Is this result possible, and if so, why? Why don't some clerks leave their low-paying jobs and become engineers, a movement that would tend to equalize wage rates in the two occupations? Why don't employers of engineers seek out workers presently employed as clerks and offer them engineering jobs at better pay (but at wage rates slightly below w_E) since their demand curve indicates they would be willing to hire more engineers at a lower wage?

Thinking about these questions suggests several possible answers. First, workers currently employed as clerks may prefer their jobs despite the financial difference; that is, they don't want to work as engineers. Second, acquiring the skills to become an engineer may have a significant cost. The wage for engineers may not be sufficiently high to compensate clerks for the training costs they would have to bear to become engineers. Third, even if there were no training costs, clerks may not have the aptitude for science and mathematics necessary to work as engineers.

Training costs as well as differences in workers' abilities and/or preferences for particular jobs can lead to *equilibrium differences in wage rates* among persons and jobs; there is no tendency toward adjustments that would wipe out wage differentials due to such factors. Now let's take a somewhat more detailed look at these factors.

Compensating Wage Differentials

Monetary compensation is not the only factor, and sometimes not even the most important factor, influencing individuals' job choices. People routinely make decisions to take jobs paying less in monetary terms than they might earn elsewhere. Many academic economists, for example, could earn 50 percent more working for government or industry but choose to remain in a scholarly environment. Similarly, some people might agree to work at the same job for less money if they could live in the Sun Belt instead of the Northeast.

When workers view some jobs as intrinsically more attractive than others, the forces of supply and demand produce differences in the wages paid. These differences are called **compensating wage differentials** because the less attractive jobs must pay more to equalize the real (monetary and nonmonetary) advantages of employment across jobs.

COMPENSATING WAGE DIFFERENTIALS
differences in wages paid that are created by the forces of supply and demand when workers view some jobs as intrinsically more attractive than others

We can illustrate the implications of differences in job attractiveness. Suppose a certain number of potential workers are identical in their abilities to work as police officers or fire fighters. At equal wages, they would all prefer to be fire fighters, perhaps because they think police work is a thankless task. Only if the police wage is at least 25 percent higher than the fire fighting wage will they choose to enter the police force. If market conditions determine wages, then wages for police officers will be 25 percent higher than those for fire fighters. Only then would the *real* wage, in the eyes of the workers, be the same for the two jobs.

Differences in money wages are necessary to equate the quantity of labor supplied and demanded in different occupations when the nonmonetary attractiveness of jobs differs. Some cities have ignored this fact and set the same wages for police officers and fire fighters. The outcome? A surplus of fire fighter recruits and a shortage of police force applicants.

APPLICATION 17.2 TWELVE HOURS' PAY FOR TEN MINUTES' WORK

Each year hundreds of "jumpers" or "glow boys" are hired by public utilities to fix steel pipes in aging power plants.[2] The jumpers work for only 10 minutes at a time and are paid the equivalent of 12 hours of average maintenance wages in addition to free travel and per diems. The only catch is the job site: the pipes are located inside nuclear power plants where radiation is so intense that a jumper can stay for only a limited period of time before, in industry parlance, "burning out." The maximum amount of radioactivity to which a jumper may be exposed is 5,000 millirems per year (the equivalent of 250 chest X-rays). The Nuclear Regulatory Commission, which sets the limit, estimates that if workers are exposed to even that level for 30 years, 5 percent will die of cancer as a result. The high pay earned by jumpers thus illustrates a compensating wage differential: riskier jobs pay more.

[2] "Ten Minutes' Work for 12 Hours' Pay? What's the Catch?," *Wall Street Journal*, October 12, 1983, pp. 1 and 21.

Differences in Human Capital Investment

Our ability to perform useful services can be augmented by training, education, and experience. People can become more productive workers, and more productive workers receive higher wage rates. The process by which people augment their earning capacity is sometimes called **human capital investment.** In these terms, human beings are viewed as capable of generating a flow of productive services over time, much as capital assets can. When they bear the costs of training or education themselves, they are investing in their own earning capacity by attempting to increase the productive services they can provide as workers.

HUMAN CAPITAL INVESTMENT
the process by which people augment their earning capacity

Education and training are much like other investments. Initially, people incur costs. For college students, for example, the explicit and implicit costs include tuition, fees, and forgone income. The payoff to the investment comes several years later; then, you find out how profitable the investment was. If earning capacity has increased sufficiently, the higher earnings in later years cover the initial investment cost. No student needs to be told that this form of investment, like many others, is risky.

Jobs requiring large human capital investments tend to pay higher wages. The reason is simple: if the wages weren't higher, few people would be willing to incur the training costs. The higher wages associated with highly skilled work are, in part, the returns on past investments in human capital. According to this view, it is no accident that college graduates on average earn more than high school graduates or that neurosurgeons earn more than typists.

In terms of our labor supply–demand model, the supply curve tells us that the amount of labor supplied to jobs requiring large investments in human capital will be forthcoming only at higher wage rates. Thus, in Figure 17.5, the supply curve of engineers is positioned higher (vertically) than the supply curve for clerks.

APPLICATION 17.3 THE RETURNS TO INVESTING IN A BA AND AN MBA

Although investing in a college education involves some risk, the average return, net of any costs, appears to be greater than ever.[3] As of 2000, according to the U.S. Census Bureau, workers with a college degree made 80 percent more than workers with only a high school diploma. As recently as 1988, college degree holders made just 48 percent more than workers with only a high school diploma. Researchers estimate that as much as 30 to 50 percent of the relative wage gains made by college graduates over the past decade reflect the spread of computer technology and the fact that college-educated workers tend to be more proficient at using computers than their non-college-educated counterparts.

The Census Bureau does not track the median annual income for holders of an MBA degree. According to *Business Week*, however, pursuing an MBA at one of the top 30 schools appears to be a profitable investment. Graduates of these schools landed jobs with median starting salaries of $111,100 in 2002. Prior to enrolling in an MBA program, the same graduates had median earnings of $53,700 per year.

[3]This application is based on "Advanced Degrees Result in Higher Earnings," *Chronicle.com*, July 18, 2002; "The Payoff from Computer Skills," *Business Week*, November 3, 1997, p. 30; and "The Best Business Schools," *Business Week*, October 21, 2002.

Differences in Ability

Workers' productive capacities depend not only on their training and experience (human capital investment) but also on certain inherited traits. The relative importance of these two factors is greatly disputed. For years people have debated whether genetic or environmental factors are more important in explaining IQs. There is no doubt, however, that inherited traits play a significant role in determining what we are capable of doing. No amount of training could turn all of us into nuclear physicists, basketball stars, entertainers, models, politicians, or business executives. People differ in strength, stamina, height, mental ability, physical attractiveness, motivation, creativity, and numerous other traits. These traits, or the lack of them, influence the work we are capable of accomplishing.

Similarly, possessing abilities that are scarce is no guarantee of a high wage. For example, the ability to wiggle your ears may be uncommon but is unlikely to generate a high income. What matters is the supply of persons with abilities required to perform certain jobs relative to the demand for their services. If consumers were unwilling to pay to watch athletes throw balls through a hoop, being able to slam-dunk in 12 different ways would not command a million-dollar salary. Obviously, then, some human abilities are in limited supply relative to demand, and, as a result, those endowed with such abilities can command higher wages.

APPLICATION 17.4 WHAT'S ON THE OUTSIDE ALSO COUNTS

A recent study indicates that beauty tends to be rewarded by the labor market.[4] Using interviewers' ratings of respondents' physical appearance, the study found that good-looking people earn higher wages than average-looking people, who in turn earn higher wages than plain-looking people. The wage premium paid by the labor market for good looks is roughly 5 to 10 percent relative to average looks and 10 to 20 percent relative to plain (below average) looks. The wage premium for beauty holds across occupations and for both women and men (if anything, the premium is larger for men than for women). Thus, contrary to the well-known saying, it's not just what's on the inside that counts when determining the wages earned by different workers.

[4]Daniel S. Hamermesh and Jeff E. Biddle, "Beauty and the Labor Market," *American Economic Review*, 84 No. 5 (December 1994), pp. 1174–1194.

17.5 ECONOMIC RENT

ECONOMIC RENT
that portion of the payment to an input supplier in excess of the minimum amount necessary to retain the input in its present use

In ordinary usage, *rent* refers to payments made to lease the services of land, apartments, equipment, or other durable assets. Economists use the term differently. **Economic rent** is defined as that portion of the payment to an input supplier in excess of the minimum amount necessary to retain the input in its present use. The economic rent accruing to suppliers in input markets is thus analogous to the concept of producer surplus in output markets.

In the history of economics, the term *rent* was originally associated with the payments to landowners for the services of their land. To see why, let's suppose that the supply of land is fixed, so its supply curve is vertical. (This assumption is a slight exaggeration since the supply of usable land is not absolutely fixed; without care land can erode, become overgrown, or lose its fertility.) A vertical supply curve means that landowners will place the same quantity on the market regardless of its price. Even at a zero price the same amount of land would be available. Thus, the minimum payment necessary to retain land in use is zero, and any actual payment above zero exceeds the minimum amount necessary to call forth the supply. *All the payments to landowners therefore satisfy the definition of economic rent.*

Figure 17.6 illustrates this point. The vertical supply curve of land interacts with the demand curve to determine the equilibrium price. The price and quantity are specified per month to indicate that we are concerned not with the sale price of the land but rather with the price for the services yielded by the use of land. The shaded area indicates the monthly income received by land suppliers.[5] The shaded area also equals the economic rent received by landowners, because all payments for the services of land exceed the (zero) amount necessary to have a supply of Q. The position of the demand curve determines the amount of economic rent. If demand increases, the price of land services goes up and the shaded area becomes larger, but the larger shaded area would still be all rent.

Because nineteenth-century economists regarded the supply of land as fixed, they viewed payments to landowners as economic rent. Today economists recognize that suppliers of other inputs may receive economic rent as well. The prices received by the owners of any

[5]In practice, owners often do not receive the income from the use of their land in monetary form. Homeowners, for example, secure the services of the land on which their houses rest. Since they own the land, the income from its use is in the form of these services directly received, but these services have a monetary value determined by supply and demand.

FIGURE 17.6

Economic Rent with a Vertical Supply Curve
When the input supply curve is vertical, the entire remuneration of the input (the shaded area) represents economic rent, since the same quantity would be available even at a zero price.

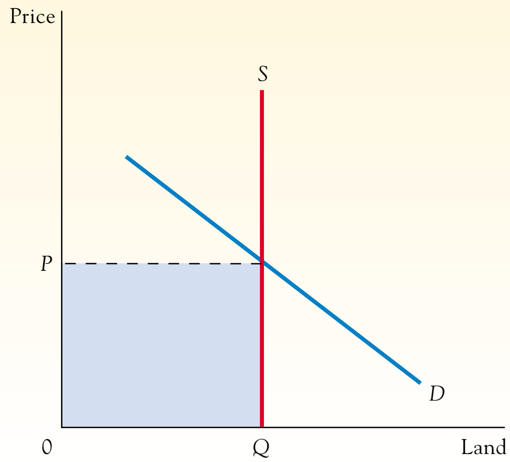

input in fixed supply are entirely rent, and part of the prices received by inputs with upward-sloping supply curves are also rent. Consider the supply of university professors. For clarity, let's examine the discrete case involving a small number of persons. The supply curve is then the step-like relationship in Figure 17.7. In equilibrium, five professors are working: individuals *A*, *B*, *C*, *D*, and *E*. We assume that they are identically productive as professors but not identical in their abilities to perform other jobs.

Individual *A* will be *willing* to work as a professor if paid area A_1. Even at a very low wage he will opt for the academic life because his next best employment opportunity is as a dishwasher, where he would earn very little. Individual *B* has a better alternative; she could sell used cars. She will work in academe if she receives at least area B_1. And so we progress up the supply curve. The supply curve slopes upward because to attract more people, colleges must bid them away from increasingly more attractive alternative employments.

FIGURE 17.7

Economic Rent with an Upward-Sloping Supply Curve
With an upward-sloping input supply curve, part of the payment to input owners represents rent. In this case individuals *A*, *B*, *C*, and *D* receive rent equal to areas A_2, B_2, C_2, and D_2, respectively.

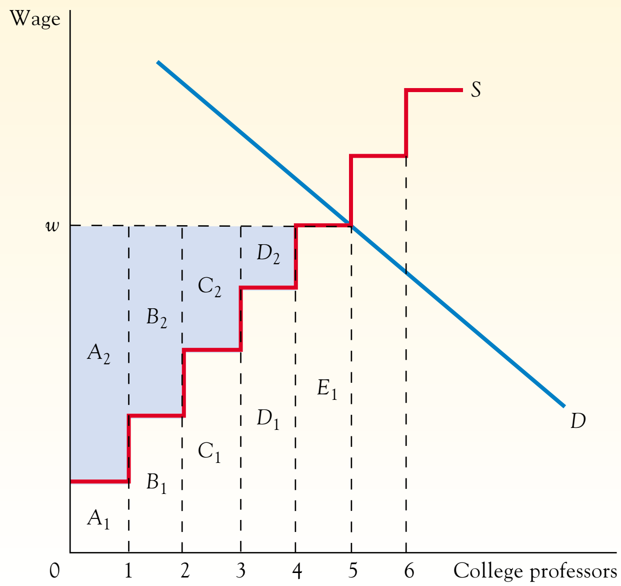

When this labor market is in competitive equilibrium, individuals *A* to *E* are paid a wage of *w*. A wage of *w* means that individual *A* receives an amount greater than the minimum necessary to induce him to work as a professor. Recall that the minimum amount is area A_1, but he is being paid A_1 plus A_2, so A_2 represents an economic rent. Individuals *B*, *C*, *D*, and *E* receive rents equal to B_2, C_2, D_2, and zero, respectively. Individual *E* is paid just enough to induce her to supply her services as a professor. If the wage were any lower, she would work elsewhere, so none of her earnings represents rent.

In this example, part of the earnings of professors (except for *E*) are economic rent, and the remainder is the payment required to keep the individuals from leaving their academic jobs and working elsewhere. The total rent received is the sum of the colored areas. Whenever the supply curve of any input slopes upward, part of the payment to inputs will be rent, as in this example. The more inelastic the supply curve, the larger the rent as a fraction of the total payment to an input. In the extreme case of perfectly inelastic supply (Figure 17.6), all the payment represents rent.

17.6 MONOPOLY POWER IN INPUT MARKETS: THE CASE OF UNIONS

In Chapter 16 we examined the hiring of inputs by firms operating in either perfectly competitive or monopoly output markets. We also explored the effects of buying or monopsony power in input markets. We did not, however, address selling or monopoly power in input markets. We do so here by focusing on labor unions. If effectively organized, labor unions are an input's sole supplier of the labor services of the union's members and thus have some influence over the wages union members are paid. Much as a monopolist in a product market seeks a price and an output level so as to maximize profit, a labor union seeks a wage and employment level so as to maximize the economic rent accruing to its members.

Consider Figure 17.8, where the demand and supply of labor to a market are given by the curves *D* and *S*. These curves depict the quantity of labor services demanded and supplied, respectively, at various possible wage rates absent any monopsony or monopoly power. Under competitive conditions, L_C workers are hired at a wage rate of w_C.

Now suppose that workers unionize and effectively cartelize the supply side of the market depicted in Figure 17.8. By organizing, the union's members acquire monopoly power and can exercise that power by selecting any wage–employment level combination on the demand curve for labor services confronting the union.[6]

What point on the demand curve for their services should a union's members select if they are interested in maximizing their economic rent? Analogous to the case of monopoly power in output markets, the union needs to take into account the marginal revenue curve, *MR*, associated with the demand curve for its services. The marginal revenue curve represents the additional wages earned by the union membership as a whole when an incremental worker is hired. The marginal revenue curve lies below the demand curve at all employment levels since the hiring of an additional worker requires the union to lower the wage level paid to all previously employed workers. The reduced payments to previously employed workers must be subtracted from the wage earned by the incremental worker (the height of the demand curve) to derive the additional wages earned by the

[6]For ease of discussion we assume that the union cannot engage in price discrimination, securing different wages for its various members. When an input supplier can price discriminate, an analysis analogous to the one presented in Chapter 12 on price discrimination in output markets is applicable.

FIGURE 17.8

The Effect of an Input Market Monopoly

An effectively organized union represents a monopoly in the market for labor services. To maximize economic rent, the union selects the employment level where the marginal revenue, in terms of additions to the union's wage bill, from having an additional worker hired (the height of the *MR* curve) equals the opportunity cost needed to induce the worker to work (the height of the *S* curve). The union employment level is L_U at a wage rate of w_U. The economic rent accruing to the union equals area w_UABF.

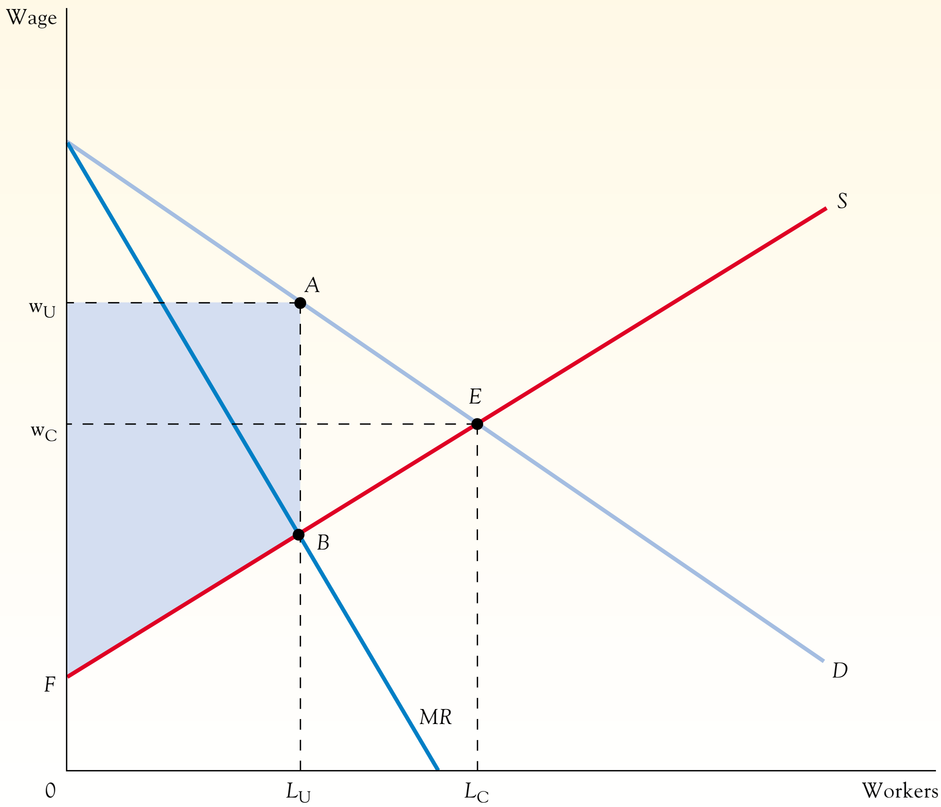

union as a whole (the height of the marginal revenue curve) when the incremental worker is hired.

To maximize economic rent, the union compares the additional wages resulting from the hiring of another worker (*MR*) with the bare minimum the worker needs to be paid to work in this market—the worker's opportunity cost. The height of the supply curve represents the bare minimum that needs to be paid to induce an incremental worker to work at various possible total employment levels and thus is effectively the union's marginal cost curve. Much as profit in an output market is maximized by setting marginal revenue equal to marginal cost, the economic rent earned by a labor union's members is maximized by selecting an employment level, L_U, where *MR* intersects *S* in Figure 17.8. Up to L_U, every worker contributes more to the union's total wage payments than it costs to induce him or her to work. The wage rate set by the union-monopoly is w_U and the economic rent accruing to the union membership as a whole is depicted by area $w_U ABF$—the excess of the wage rate over workers' opportunity costs from zero to the quantity of labor services sold, L_U, by the union-monopoly.

Note that in Figure 17.8 fewer workers end up being employed as a result of the union than under competitive conditions. This is because the union is presumed to be interested in maximizing the economic rent accruing to its membership rather than the number of workers having jobs. The economic rent from the monopoly outcome of L_U and w_U (area $w_U ABF$) exceeds the economic rent associated with the competitive outcome of L_C and w_C (area $w_C EF$).

By restricting output in the market for labor and raising the prevailing wage rate, the union illustrated in Figure 17.8 produces a deadweight loss similar to the one resulting from monopoly in output markets. Triangular area *AEB* is the deadweight loss associated with the labor union. This can be seen by noting that in order to maximize economic rent, the union restricts labor employment to L_U from L_C. Each worker between L_U and L_C can generate more in terms of the value of output produced (the height of the *D* curve) than it costs to induce the worker to work (the height of the *S* curve). These potential net gains are not realized, however, because a labor union interested in maximizing economic rent does not care about the total value of the output produced by an incremental worker (the height of the demand curve). Rather, such a union cares about the incremental worker's contribution to the union's total wage bill (the height of the *MR* curve).

While acknowledging the potential deadweight loss associated with unions, some economists have argued for a more benign evaluation of their role since they may be organized to counteract monopsony power possessed by input-buying firms.[7] For example, in this view, the United Auto Workers union arose out of the need to protect workers' rights and wages from the actions of the "Big Three" automobile companies. Furthermore, to the extent that unions set up effective grievance procedures and give workers a "voice" with their employers, they may actually make workers more satisfied and productive on the job (the height of the labor demand curve shifts upward if workers become more productive, all else equal).

There is some empirical evidence that, controlling for other factors, union workers not only receive higher wages but also are more productive than nonunion workers. Most economists, however, are skeptical that unions actually improve workers' productivity. One would expect the value of the final output produced by an additional worker to increase the fewer the number of workers hired. Fewer workers have more of the other, nonlabor inputs to share between themselves and thus, if the law of diminishing returns holds, have higher marginal products.

In Figure 17.8, for example, the height of the labor demand curve reflects the value of the final output produced by an incremental worker and could be used as a measure of worker productivity. According to such a measure, workers are more productive at the union employment level of L_U than at the competitive level of L_C. This does not imply, however, that the union has made each worker more productive. The height of the demand curve at an employment level L_U is just as high under a union-monopoly as it is with perfect competition. Once again it is important to distinguish between a movement along a given demand curve (in this case a labor demand curve) and a potential shift of the entire demand curve.

[7]See, for example "Unions Need Not Apply," *New York Times*, July 26, 1999, pp. C1 and C14; and Richard B. Freeman and James L. Medoff, *What Do Unions Do?* (New York: Basic Books, 1984).

APPLICATION 17.5 THE DECLINE AND RISE OF UNIONS

Labor unions in the United States have undergone dramatic membership changes over the past few decades.[8] In the private sector, the percentage of nonagricultural wage and salary workers who are union members has plummeted from 38 percent in the early 1950s to 9 percent in the late 1990s. By contrast, over the same time period, public sector union membership has skyrocketed from 10 to 37 percent.

What accounts for the different fortunes of private and public sector unions in the United States over the past few decades? Public sector unions have grown in relative importance due to changes in labor laws favorable to the unionization of government workers at all levels of government throughout the 1960s and 1970s. Prior to the 1960s, collective bargaining by most government workers was illegal.

In the private sector, the rapid growth of the high technology sector explains some of the overall decline in union membership in the United States. Only a small fraction of the 5 million employees who work for computer semiconductor and software companies are unionized. This is due to demand for such workers growing much more quickly than the supply and companies promising generous salaries and stock options to recruit and retain workers—rewards that diminish the relative benefits workers could obtain through unionization.

Furthermore, growing international competition in certain industries, such as automobile production, and the enhanced competitiveness of other deregulated industries, such as trucking and airlines, have made the output demand curves for firms operating in these industries more price elastic. As we saw in the previous chapter, one of the factors that determines the elasticity of demand for an input (and hence the monopoly power of sellers of that input) is the elasticity of the demand for the final product produced by the input. For example, the monopoly power possessed by the United Auto Workers has been lessened as American car manufacturers have faced growing foreign competition. In the 1950s, when foreign competition was more limited, the demand for General Motors (GM) cars was more price inelastic and any wage increase demanded by the United Auto Workers from GM could more readily be passed on to final consumers and so agreed to by GM.

[8]"Public and Private Unionization," *Journal of Economic Perspectives*, 2 No. 2 (Spring 1988), pp. 59–110.

17.7 BORROWING, LENDING, AND THE INTEREST RATE

INTEREST RATE (1) the price paid by borrowers for the use of funds, and (2) the rate of return earned by capital as an input in the production process

This chapter has so far focused on labor markets. We now turn to the market for another critical input, capital. We begin our discussion with a term, **interest rate**, that has two apparently different meanings in economics. *Interest rate* sometimes refers to the price paid by borrowers for the use of funds. When a person borrows $100 this year and must return $110 to the lender a year later, the additional $10 is interest, and the ratio of the interest to the principal amount, 10 percent, is the annual interest rate. *Interest rate* can also refer to the rate of return earned by capital as an input in the production process. When a person purchases a machine for $5,000, and the use of that machine generates $500 in income each year thereafter, the rate of return is 10 percent. Economists designate both the return on loaned funds and the return on invested capital as *interest rates* because there is a tendency for these returns to become equal.

To explore the determinants of the interest rate, let's start with a simple example. Suppose that no capital investment is taking place; that is, no investors are borrowing to finance building construction or equipment purchases. Although there is no borrowing for investment purposes, there are still people who wish to lend money and others who wish to borrow. Households whose current incomes are high in comparison with their expected future incomes may be willing to lend to others and use the repayment of the loan to augment their otherwise lower ability to consume in the future. Other households may wish to borrow in order to consume more than their current income in the present, repaying the loan in the future. These households are the potential suppliers and demanders of consumption loans, and their interaction determines a price, or interest rate, for borrowed funds.

The level of the interest rate affects both the quantity supplied and quantity demanded of loanable funds. The suppliers of funds are saving—that is, consuming less than their current income allows. As noted in Chapter 5, a higher interest rate has opposing income and substitution effects on the amount saved. The substitution effect encourages less present consumption (more saving) because each dollar saved returns more in the future. The income effect, however, increases the real incomes of savers, encouraging more present consumption and less saving. On balance, a higher interest rate may lead to more or less saving. The supply of loanable funds from savers is likely to slope upward at low interest rates but can become backward bending at sufficiently high interest rates. This curve is analogous to the supply curve of hours of work.

On the other side of the market, the demand curve for funds from borrowers must slope downward. To a borrower, a higher interest rate has income and substitution effects, both of which reduce the level of desired borrowing. The substitution effect reflects the fact that a higher interest rate makes it more expensive to finance increased consumption from borrowed funds, inhibiting borrowing. The income effect reflects the fact that a higher interest rate reduces the real income of borrowers, so they cannot afford to borrow as much.

Figure 17.9 illustrates the interaction of borrowers and lenders in the market for consumption loans. The supply of saving (funds to lend) is S_S, and the demand for these funds by borrowers for consumption purposes is D_C. Equilibrium occurs when lenders provide $10 million to borrowers in the present time period, with borrowers agreeing to repay the loans at an interest rate of 5 percent.

In this simple example the funds saved are not used in a way that increases the output of goods through capital accumulation. Yet saving is productive. It provides the means

FIGURE 17.9

A Borrowing-Lending Equilibrium
If there is no investment demand for funds, the demand for consumption loans and the supply of saving determine the interest rate. Here people borrow $10 million to finance present consumption with the commitment to repay it plus 5 percent interest.

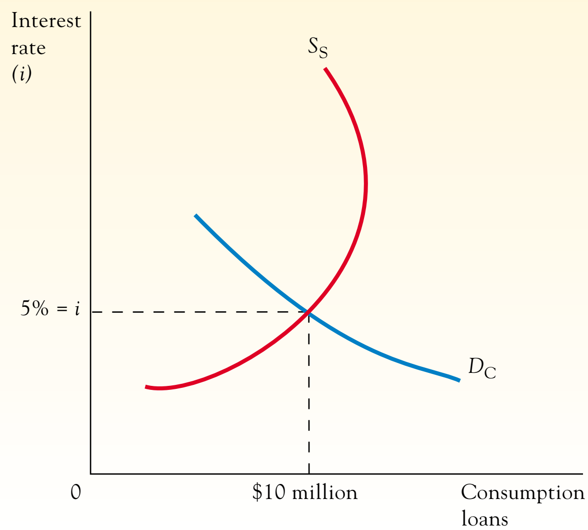

for borrowers to have a consumption pattern over time that is more to their liking than having to live within their incomes each year, and these preferences explain why borrowers are willing to pay for the use of borrowed funds. A positive interest rate emerges because lenders must receive compensation for sacrificing their use of the funds for present consumption.

17.8 Investment and the Marginal Productivity of Capital

In the consumption loan example, the only outlet for saver-supplied funds is borrowing by other households to finance consumption. Now let's expand the analysis to account for the fact that saving also provides funds used to finance investments. Firms may borrow money for the purpose of enlarging their stock of capital equipment, and they must compete for the limited supply of saving with households that are borrowing to finance consumption.

Why are firms willing to incur interest costs to finance investments in capital? Basically, the reason is that capital contributes sufficiently to production to repay the interest costs. For example, if Robinson Crusoe fishes by hand, he may be able to catch 20 fish per week. If he takes one week off to weave a net, he can then catch 25 fish a week with the net until it wears out in 10 weeks. Unfortunately, Crusoe is going to get very hungry if he doesn't eat for a week. However, if he borrows 10 fish from his Man Friday, on the condition that he pays back the 10 fish plus an extra 5 fish, he can catch an additional 5 fish a week.

gross marginal productivity
the total addition to productivity that capital investment contributes

net marginal productivity
the total addition to productivity that capital investment contributes, less the cost of capital

The 5 additional fish caught per week for 10 weeks (50 fish) are a measure of the **gross marginal productivity** of the net. Whether the net (the capital) is productive depends on whether the gain in output outweighs the cost of not fishing for a week so he could weave a net. Since Crusoe could have caught 20 fish that week, the cost of the net is a sacrifice of 20 fish plus the 5 fish paid as interest on the 10 borrowed fish. (The 10 borrowed fish are not a cost; he borrowed 10 and returned 10, so the cost is zero.)

By sacrificing 25 fish, Crusoe gains 50 fish over the life of the net, representing a gain of 25 fish. In this example, capital—the net—is productive in the sense that its **net** (no pun intended) **marginal productivity** is 25 fish. When the capital's net marginal productivity is positive, investment—the act of adding to the amount of capital—allows Crusoe to produce a larger output even after "netting" out the cost of the capital. (Note as a brief aside that Man Friday also benefits, as he picks up an extra 5 fish for his trouble.)

Next, we need a measure for the net productivity of investment in capital that allows us to compare the productivities of various projects. The annual percentage rate of return is convenient for this purpose. Suppose a machine costs $100 to construct this year and its use next year will add $120 to output. For simplicity we will let the machine wear out after one year's use. Then the net gain is $20 in one year, an annual rate of return on the initial $100 investment of 20 percent. The 20 percent figure is the net marginal productivity of investment, and it measures how much the capital investment will add to output one year hence per unit of present cost. It is essentially the rate of return on the investment. Denoting this rate of return measure of productivity by g, we can calculate it from the formula:

$$C = R/(1 + g) \qquad \text{or}$$

$$\$100 = \$120/(1 + g); \qquad (1)$$

where C is the initial cost and R is the resulting addition to output (capital's gross marginal value product) the next year. When capital equipment yields services over more than one year, the normal case, the principle is the same but the formula is slightly more complicated.

Suppose the equipment lasts for 2 years before wearing out and adds \$60 to output after one year (R_1) and \$81 in the second year (R_2). Then, we calculate the rate of return from:

$$C = [R_1/(1 + g)] + [R_2/(1 + g)^2] \qquad or$$

$$100 = [\$60/(1 + g)] + [\$81/(1 + g)^2]. \qquad (2)$$

Given the initial cost of the equipment and the contribution to output in each year, we can solve this expression for g. In this example the rate of return is 0.25, or 25 percent per year.[9]

The Investment Demand Curve

The net productivity of a capital investment, g, depends on the value of additional output generated (the R terms) and on the initial cost (the C term). In turn, the C and R terms depend on many factors, among them the size and skills of the labor force, the amount and costs of other inputs such as natural resources, technology, and the degree of political stability. When we hold these other factors constant, the net productivity of investment also depends on the rate of investment undertaken per time period. The greater the rate of investment, the lower the net *marginal* productivity of still more investment.

Consider the rate of investment for the economy as a whole. As investment increases, each additional dollar's worth of capital adds less to output than the previous dollar's worth because of the law of diminishing marginal returns; more capital applied to a given labor force and quantity of land causes its marginal product to decline. This factor reduces the R terms in the formula and contributes to a lower rate of return. In addition, an increase in investment also means more demand for machinery, vehicles, buildings, and other types of capital, which tends to raise their prices. This factor increases the initial cost of investment, the C term, and also contributes to a lower rate of return.

INVESTMENT DEMAND CURVE the relationship between the rate of return generated and various levels of investment

Taken together, these factors imply that a lower rate of return will be associated with a greater amount of investment, other things being equal. The downward-sloping D_I curve in Figure 17.10 illustrates this point. When investment is at a level of I_0, g is equal to 15 percent. If investment increases to I_1, the rate of return on invested capital falls to 10 percent.

The D_I curve, indicating the rate of return generated by different investment levels, is the **investment demand curve.** This idea is easiest to understand if we suppose that firms and individuals finance their investments by borrowing. If the interest rate is 10 percent, in-

FIGURE 17.10

Investment Demand Curve

The D_I curve shows the rate of return generated by alternative investment levels. It is the investment demand curve, and it shows the amount invested at various interest rates.

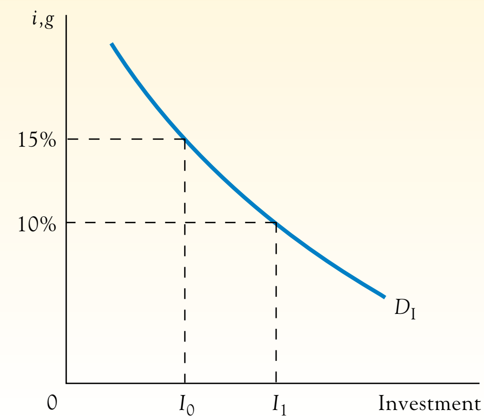

[9]In the general case where equipment lasts n years, the formula is: $C = [R_1/(1 + g)] + [R_2/(1 + g)^2] + [R_3/(1 + g)^3] + \ldots + [R_n/(1 + g)^n]$.

vestment expands to I_1, where the return on the investment just covers the cost of the borrowed funds. At any lower investment level, the rate of return on investment will be higher (15 percent at I_0, for example), yielding a pure economic profit to firms and individuals who are only paying 10 percent for the funds they are investing. An expansion in investment will occur as long as the rate of return is greater than the cost of borrowed funds, and such an expansion causes the rate of return to fall. Equilibrium results when investments yield a return just sufficient to cover the interest rate on borrowed funds, at I_1 when the interest rate is 10 percent. *Thus, the rate of return on capital investment* (g) *tends to equal the interest rate for borrowed funds* (i).

Even if investment is not financed by borrowing, I_1 still is the equilibrium investment level if the interest rate is 10 percent. For example, a firm with $1 million in retained earnings can use this sum to finance an equipment acquisition, but it will not do so unless the investment yields at least 10 percent. Why? Because the firm can lend the $1 million at the 10 percent interest rate; if the investment project yields less than 10 percent, the firm can do better by lending funds rather than investing them. The opportunity cost of investing is the same whether the funds are borrowed or acquired in any other way, so the 10 percent interest rate guides investment decisions in either case.

Our discussion so far assumes that firms and individuals know the rate of return generated by investments. In most cases, however, investors do not know, since rates of return depend on future outcomes. The investment demand curve reflects what investors *expect* the outcome of investment projects to be. After the fact, some investments expected to yield 10 percent fail to do so, and others do better. For given expectations on the part of actual and potential investors, the demand curve still slopes downward. A change in expectations causes the entire curve to shift. If investors expect nationalization of industry or violent revolution to occur in a country, the demand shifts far to the left, one reason firms are reluctant to invest in countries with unstable political regimes.

17.9 SAVING, INVESTMENT, AND THE INTEREST RATE

Households with a demand for consumer loans, and firms and persons with investment projects, compete for saver-supplied funds. So far we have discussed these elements separately; Figure 17.11 brings them together. The consumer loan demand curve is D_C, and D_I is the investment demand curve. The horizontal summation of these curves yields D_T, the total demand for funds supplied by savers. The total demand in conjunction with the supply of savings determines the interest rate. At the market-determined interest rate of 10 percent, consumer loans are L_1 and investment is I_1, with the sum equal to T. As drawn, the investment demand is much greater than consumer loan demand. This tends to be the real-world case.

In this model, the gross saving of households and firms, T, is not equal to investment, I_1; consumer loans account for the difference. Economists sometimes define saving as net of consumer loans ($T - L_1$), and in that definition net saving equals investment. (We have also ignored the government's demand for funds to help finance budget deficits. Including that factor would create another demand for funds supplied by savers.)

The Figure 17.11 analysis corresponds to the aggregate labor demand and supply model since we are talking about the supply and demand for funds aggregated across all saving-lending and borrowing-investment markets. It is important to recognize that there is not just one market that determines the interest rate but many closely interrelated markets that we are summarizing. Bond markets, stock markets, mortgage borrowing, credit card loans, bank deposits and loans, investment of retained earnings by firms, and other markets are all involved in this process of allocating funds, provided by a multitude of sources, among

FIGURE 17.11

The Equilibrium Levels of Saving, Investment, Consumer Loans, and the Interest Rate
The total demand for funds supplied by savers, D_T, is the sum of the investment demand curve, D_I, and the demand for consumer loans, D_C. The intersection of D_T and S_S determines the interest rate. Investment is I_1, saving is T, and consumer loans are L_1.

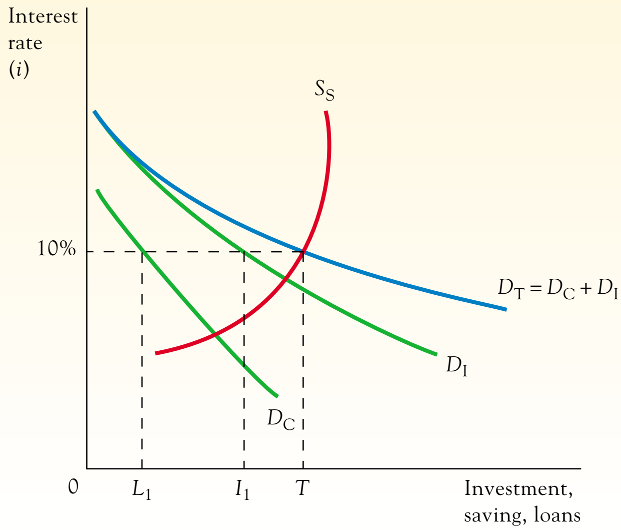

different competing uses. We lose some detail by using an aggregated model (as we did in the labor market model), but we gain the important advantage of emphasizing the common underlying factors affecting all the interrelated markets: namely, the willingness to consume less than current income and the existence of profitable investment opportunities.

Investment in excess of the amount required to replace worn-out capital adds to the stock of productive capital, which, in turn, increases the productive capacity of the economy in subsequent periods. Figure 17.12 illustrates this effect. In the present year, a society's production possibility frontier (*PPF*) relating the attainable output of consumer goods and capital goods is *FF′*. The production of capital goods requires the use of resources (inputs) that can otherwise be used to produce consumer goods, so if more capital goods are produced, fewer consumer goods are available, and vice versa. The market forces summarized in Figure 17.11 determine where we are located on this frontier. Investment of I_1 in Figure 17.11 means an addition to the stock of capital of that amount and corresponds to a point such as K_1 in Figure 17.12. With more capital available the following year, the productive capacity of the economy increases. When K_1 is invested in year 1, the *PPF* becomes F_1F_1' in year 2. If investment were greater in year 1—for instance, at K_2 (point *E′*)—then the *PPF* would move out even farther, to F_2F_2'. The effect of capital investment on the position of the *PPF* in subsequent periods is due to the productivity of capital discussed in Section 17.8. (There are also other reasons why the *PPF* shifts over time—most importantly, growth in the labor force, investment in human capital, and technological progress.)

Investing more in the present time period thus means that incomes and consumption are greater in the future. Still, we should avoid drawing the conclusion that an ever-increasing expansion in investment is desirable. To invest more, less must be consumed; that is, for an increase in investment from K_1 to K_2, consumption must fall from C_1 to C_2. Thus, the cost of greater investment (yielding higher future income) is increased saving (reduced consumption) in the present. The saving supply curve shows how much households must be paid in return for providing funds for investment purposes. At the Figure 17.11 optimal point, households will supply more saving than *T* only if they receive an interest rate higher than 10 percent. Investors, however, will be unwilling to pay savers more than 10 percent for ad-

FIGURE 17.12

The Level of Investment and Productive Capacity Greater investment in the present—E' rather than E—means less current consumption—C_2 versus C_1—but greater capacity to produce goods in the future. When K_1 is invested in year 1, the production frontier in year 2 is F_1F_1'; if K_2 is invested, it is F_2F_2'.

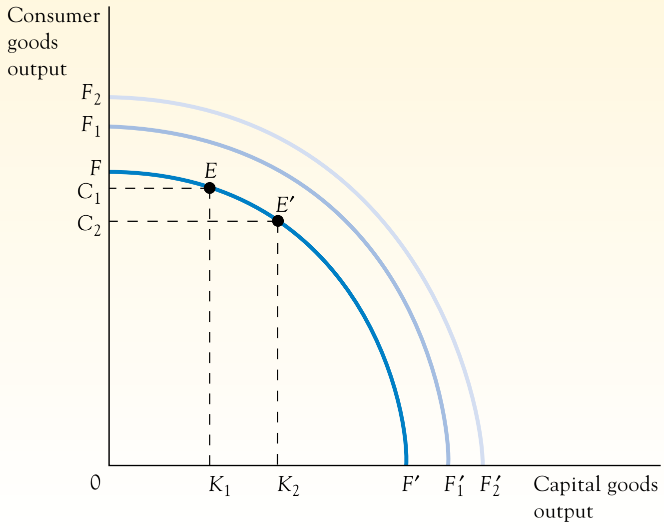

ditional funds since the increased investment yields less than a 10 percent return. As a result, the market is in equilibrium with investment of I_1.

Equalization of Rates of Return

There is a tendency for capital to be allocated across firms and industries so that the rate of return is equal everywhere. This parallels the tendency of labor to be allocated so that wages are equal. In the case of labor, however, we pointed out a number of reasons why wages differ because of differences in the productivities and preferences of people. With capital, fewer qualifications are necessary. A firm purchasing a machine is indifferent to the source of funds for financing; one person's money can purchase as much capital as another's. Moreover, a person supplying funds usually doesn't care whether the funds are used to finance a computer in the aerospace or construction industry; lenders care only about the rate of return earned on the saving. There are, of course, some reasons why rates of return won't be exactly equal, such as differing degrees of risk, but they are generally not quantitatively as important as the factors that produce differences in wages.

The process by which rates of return tend to equalize is much like the process for labor. Let's say that two industries, *E* (for entertainment) and *S* (for steel), are initially in equilibrium, earning the same rates of return on invested capital. Next, let an unexpected shift in demand occur. Suppose consumers' demand for *E* increases while demand for *S* falls. As explained in Chapter 9, the short-run effects are economic profits in industry *E* and losses in industry *S*. In terms of the return to capital, capital now earns an above-normal rate in industry *E* but a below-normal rate in industry *S*. (A normal rate of return signifies the rate of return investors can expect to earn on capital invested elsewhere and thus represents the opportunity cost of capital.) The owners of capital in industry *S* have an incentive to shift their investments to industry *E*, where the return is higher. As capital (and other resources) moves from *S* to *E*, the output of industry *E* expands, and price falls until the industry earns a normal return (zero profit). The opposite happens in industry *S*: the industry contracts, investors withdraw capital, price rises, and once again the industry earns a normal return. This process describes how industries adjust in response to a change in consumer demand from the perspective of the market for the input, capital.

17.10 WHY INTEREST RATES DIFFER

Although we have been discussing *the* interest rate and its relationship to *the* rate of return on capital, you should know that this description is a simplification, just as it was when we discussed *the* wage rate in the aggregate labor supply–demand model. There is, in fact, a range of interest rates, and *the* interest rate in an aggregate model is best thought of as shorthand for "the general level of interest rates." In that sense the model developed in the preceding sections helps to pinpoint important determinants of the level of interest rates.

Differences can exist in specific interest rates in equilibrium, although the differences are less pronounced than differences in wage rates. The four most important reasons for differences in interest rates are as follows:

1. *Differences in risk.* There is always the possibility that a loan will not be repaid. The greater the risk that a borrower will default, the higher the interest rate that a lender will charge. If there is a one-in-five chance of a default, for example, the lender will have to charge an interest rate about 25 percent higher than for no-risk loans to receive the same expected return. For this reason corporate bonds pay higher interest rates than do government bonds, and loans secured by collateral (like home mortgages) involve lower interest rates than installment credit.
2. *Differences in the duration of the loan.* Borrowers will generally pay more for a loan that does not have to be repaid for a long time since it gives them greater flexibility. Usually, lenders must also receive a higher interest rate to part with funds for extended periods. This is one reason why savings accounts, where the funds can be withdrawn on short notice, pay lower interest rates than six-month certificates of deposit.
3. *Cost of administering loans.* A small loan usually involves greater bookkeeping and servicing costs per dollar of the loan than a large loan. Loans repaid in frequent installments, such as automobile loans, also involve higher administrative costs. When greater costs are associated with administering loans, the borrower must cover these costs and thus pay a higher interest rate.
4. *Differences in tax treatment.* The way the tax system treats interest income and investment income is very complex, and it sometimes leads to divergences in interest rates that would otherwise be more nearly equal. For example, the interest paid by state and local governments on municipal bonds is not taxable under the federal income tax, but the interest on otherwise comparable corporate bonds is. The after-tax returns are what guide lenders' decisions, and they will tend to be brought into equality, implying a difference in the before-tax (market) rates of interest. The result is that state and local governments can borrow at lower rates of interest than corporations.

Although these factors lead to divergences in interest rates in specific credit markets, they should not obscure the common factors that affect all interest rates. For example, if the public decides to save less (a leftward shift in the S_S curve), all these rates would go up, as would the rate of return on invested capital.

17.11 VALUING INVESTMENT PROJECTS

Our discussion of interest rates suggests a way for profit-maximizing firms or individuals to determine if undertaking a particular investment is worthwhile. To see how investment projects may be valued, let's use a very simple example: namely, let's discard any differences between specific interest rates due to risk, tax considerations, and the like. There is only one interest rate, *i*, reflecting the annual rate at which one can borrow or lend money. Furthermore, let's assume that there is no change in the price level over time, so that *i* reflects both the nominal and real rate of interest. In such a setting, consider the case of a hospital administrator contemplating a $1.4 million investment in a CAT (computerized axial tomography)-scanning machine. The machine will generate $500,000 in annual net revenue in each of the first three years before wearing out. Should the administrator make the investment?

As with most other questions in economics, the correct answer is, "It depends." In our specific example, the answer depends on the level of the interest rate, *i*—the opportunity cost of using money over time. If *i* equals zero and there is no opportunity cost to using money over time, then the investment in the CAT-scanning machine is clearly worthwhile. In such a case, a payoff of $1 a year from now is worth the same as a dollar in the bank today; $1 saved in the bank (or invested in capital, for that matter) generates zero interest and consequently still equals $1 at the end of the year. The $500,000 payoff generated by the CAT-scanning machine a year into the investment thus is equivalent to $500,000 in the bank at the present time. The additional $500,000 generated 2 years into the investment is also equivalent to $500,000 in the bank at the present time (since the interest rate or opportunity cost of using money for each of the first 2 years is zero percent). The same logic applies to the third-year payoff. In total, if the interest rate is zero, the machine provides a payoff equivalent to $1.5 million in terms of present dollars ($500,000 in each of the first 3 years) and therefore generates more in net revenue than the $1.4 million investment cost confronting the administrator.

PRESENT DISCOUNTED VALUE OF $1
the amount of money, X_t, placed in the bank today that will yield $1 *t* years from now at interest rate *i*

If the interest rate is greater than zero, the answer is less clear-cut. Suppose, for example, that the annual interest rate, *i*, is 10 percent. One dollar invested in the bank today generates $\$1(1 + i) = \$1(1 + 0.1) = \$1.10$ a year from now. Moreover, the **present discounted value** (*PDV*) of $1 paid a year from now is the amount of money, X_1, placed in the bank today that will yield $1 a year from now at the interest rate, *i*. Since we are looking for the sum X_1 that when multiplied by $(1 + i)$ will yield $1 a year from now, we can solve for X_1 as follows:

$$X_1 = PDV \text{ of } \$1 \text{ paid a year from now} = \$1/(1 + i).$$

At an interest rate of 10 percent, the present discounted value of $1 paid a year from now thus equals $\$1/1.1 \approx \$0.91$ paid today; $1 paid a year from now is worth only $0.91 today when the opportunity cost of using money over time is 10 percent. This relationship makes intuitive sense. When money can be put to productive use over time (when the interest rate is positive), then $1 a year from now is worth less than $1 in hand today. The dollar in hand today can be invested and will yield more than $1 a year from now.

The foregoing formula can readily be employed to determine the present discounted value of $500,000 paid a year from now. The *PDV* is simply $500,000 discounted by the factor $1/(1 + i)$ or $\$500{,}000/(1 + i)$. If the interest rate is 10 percent, the *PDV* of $500,000 paid a year from now is thus $\$500{,}000/1.1 \approx \$454{,}545.45$.

The logic we have just employed can also be used to determine the *PDV* of $1 paid two years from now, X_2. Since the dollar amount that needs to be placed in the bank today, X_2, to equal $1 two years from now will yield $X_2(1 + i)$ a year from now and $X_2(1 + i)(1 + i)$ two years from now, we know that $X_2(1 + i)^2 = \$1$ paid two years from now. Consequently:

$$X_2 = PDV \text{ of } \$1 \text{ paid two years from now} = \$1/(1 + i)^2.$$

At an interest rate of 10 percent, the *PDV* of $1 paid 2 years from now is $\$1/(1.1)^2 \approx \$0.83$; $1 paid two years from now is worth only $0.83 in the bank today. Note that with a payment of $1 that is two years versus one year away, we need less money in the bank today ($0.83

versus \$0.91) to be able to equal the payment. Furthermore, the *PDV* of \$500,000 paid two years from now can be derived by simply multiplying \$500,000 by the discount factor of $1/(1.1)^2$. With a 10 percent interest rate, \$500,000 paid two years from now equals $\$500{,}000/(1.1)^2 \approx \$413{,}223.14$.

The *PDV* of \$1 paid three years from now, X_3, can be derived in an analogous fashion, as can the *PDV* of \$1 paid *n* years from now, X_n. The correct formulas are as follows:

$$X_3 = PDV \text{ of } \$1 \text{ paid three years from now} = \$1/(1+i)^3, \text{ and}$$

$$X_n = PDV \text{ of } \$1 \text{ paid } n \text{ years from now} = \$1/(1+i)^n.$$

Based on the formula for X_3, the \$500,000 payoff three years into the CAT-scanner investment with an interest rate of 10 percent has a present discounted value of $\$500{,}000/(1.1)^3 \approx \$375{,}657.40$.

We can now put all the foregoing pieces together to determine a general rule for whether the stream of payments associated with an investment project such as the CAT scanner is worth the investment cost. Specifically, the **net present value** (*NPV*) of a project takes into account the initial investment cost (a negative present value) and the *PDV* of any future payoffs. The *NPV* of the CAT scanner is thus:

NET PRESENT VALUE (NPV)
the present discounted value of any future payoffs from an investment minus the initial investment cost of a project

$$NPV = -\$1{,}400{,}000 + \$500{,}000/(1+i) + \$500{,}000/(1+i)^2 + \$500{,}000/(1+i)^3.$$

With an interest rate of 10 percent, the *NPV* of the CAT scanner turns out to be:

$$NPV = -\$1{,}400{,}000 + \$454{,}545.45 + \$413{,}223.14 + \$375{,}657.40$$
$$= -\$156{,}574.01.$$

This implies that, based on the opportunity cost of capital over time, investing in the CAT scanner is not profitable when the interest rate is 10 percent. The same \$1.4 million that it would take to purchase the CAT scanner yields a higher return if invested in the bank than if invested in the CAT scanner when the interest rate is 10 percent. *As a general rule, if the net present value of an investment project is negative, then the funds required by the project could be better employed elsewhere. When the* NPV *of a project is positive, then the project represents a profitable employment of capital at capital's prevailing opportunity cost.*

While extremely straightforward, the mechanism just outlined for valuing investment projects has wide-ranging applicability. Among other things, it can be employed to price stocks and bonds that entail an up-front purchase price but promise future returns such as a stream of dividends and coupon payments, determine the value of lost earnings when an individual is involved in an accident, and evaluate the cost of a home mortgage or car loan.

APPLICATION 17.6 WHY LOTTERY WINNERS MAY NOT BE "MILLIONAIRES"

Most state lotteries offer jackpots of \$1 million or more. Thus, the lucky winners are often called "millionaires."[10] Because of the manner in which winnings are paid out by state lottery commissions, however, winning \$1 million does not make you a millionaire in present discounted value terms. The California State Lottery Commission, for example, awards all million-dollar prizes over a 20-year period: the winner

[10]This application is based on *www.calottery.com/winnersgallery/stories.html*.

is paid $50,000 immediately and $50,000 in each succeeding year for the next 19 years. By paying off on the installment plan, the Commission pays a $1-million-winner much less than $1 million in present discounted value terms. If the prevailing interest rate is 10 percent, for instance, the present discounted value of $1 million paid out in equal installments of $50,000 each over 20 years is slightly less than $426,000.

Summary

- The aggregate labor supply curve identifies the total number of hours of work that will be supplied to the economy as a whole at different wage rate levels. It is derived from the income-leisure choices of individual workers.
- Since the income and substitution effects operate in opposite directions, the aggregate labor supply curve is likely to be quite inelastic and probably bends backward above some wage rate.
- A major application of the aggregate labor supply curve, in connection with the aggregate labor demand curve, is in analyzing the determination of the general level of real wages. The aggregate demand curve is essentially the marginal product of labor curve. Its position depends on the amount of other inputs available (such as capital), the level of technology, and the productive characteristics of the labor force.
- Wage rates can differ among jobs and workers for several reasons, with no tendency for the rates to become equal. Equilibrium wage differences can be due to compensating wage differentials, differences in human capital investments, and/or differences in ability.
- Economic rent arises when an input supplier is paid more than the minimum amount necessary to retain the input in its present use. Supply and demand analysis shows that part of an input's price represents rent whenever the input is in less than perfectly elastic supply.
- Borrowing-lending and saving-investment markets interact to determine both the interest rate on loaned funds and the rate of return on invested capital. In equilibrium, the interest rate on loaned funds and the rate of return on invested capital tend to be equal, ignoring any differences in risk, tax treatment, and the like. The interest rate thus acts to equalize the willingness of people to give up present consumption for future consumption (marginal rates of substitution) and the real net marginal productivity of investment.
- The equilibrium interest rate may be used to determine the net present value of various types of investment projects. Profit maximization dictates that projects with a positive net present value should be undertaken.

Review Questions and Problems

Questions and problems marked with an asterisk have solutions given in Answers to Selected Problems at the back of the book (pages 583–584).

17.1. Edie chooses to work 90 hours per week when the wage rate is $16 per hour. If she is offered time-and-a-half ($24 per hour) for "overtime work" (that is, hours in excess of 90 per week), will she choose to work longer hours? Support your results with a diagram.

17.2. If workers differ widely in their preferred hours of work at a given wage rate, would you expect to see all firms set the workweek at the average preferred level? Why or why not?

***17.3.** Jeeves works 40 hours a week at a $7.50 wage rate. He unexpectedly inherits a trust fund that pays him $200 per week. How does inheritance of the trust fund affect his budget line? Will he continue to work 40 hours a week? Will his total income rise by $200 a week?

***17.4.** Can a person choose to work so much less at a higher wage that her total earnings decline? Show the income and substitution effects implied by that outcome.

17.5. "Kevin (a highly skilled businessman) earns $150,000 a year. If he were to return to his native New Guinea (where there are more limited business opportunities), Kevin would be able to earn only $5,000 a year. This proves both that Kevin doesn't deserve his high salary and that wages are arbitrarily set in the United States." Evaluate this statement.

17.6. In general, the aggregate demand for and the supply of labor increases over time. Can we predict what will happen to real wage rates and employment over time? What factors are responsible for the shift in the demand curve? In the supply curve?

***17.7.** "Since only a few people have the ability to become nuclear physicists, the long-run supply curve of nuclear physicists is vertical." True or false? Explain.

17.8. Discuss the three reasons for equilibrium wage rate differences given in the text. Which one, or more, accounts for differences in wage rates between engineers and elementary school teachers? College professors and high school teachers? Basketball superstars and basketball coaches? Doctors and lawyers?

17.9. If workers had identical preferences and were equally able to perform any job, would all wage rates be equal?

17.10. If Wayne Gretzky's supply curve of hours of work to professional hockey is vertical, does this imply that all of his labor earnings represent economic rent?

17.11. What do we mean when we say that capital is productive? How do we measure this productivity? What does productivity have to do with the investment demand curve?

17.12. "In equilibrium, the interest rate equates the willingness of people to give up present goods in return for future goods and the ability of the economy to transform present goods into future goods." Explain.

17.13. How does the interest rate serve to equalize the desired rate of saving and the desired rate of investment? Would this function be served if the government placed an upper limit on the interest rate lenders could charge? How would such a law affect the amount of investment undertaken?

***17.14.** In Figure 17.12, will the aggregate demand for labor in year 2 be higher if investment is K_1 or K_2 in year 1? Why?

17.15. If the union shown in Figure 17.8 wishes to maximize the total wage payments made to its members, what wage and employment levels should be selected? Explain your answer.

17.16. In the example of the hospital administrator and the CAT scanner discussed in Section 17.11, is the net present value positive if the interest rate is 5 percent?

17.17. Suppose that you are considering purchasing a corporate bond promising to pay a $1,000 "coupon" in each of the next five years. If the interest rate is 10 percent, what will be the price for such a bond if the bond market is competitive? Will the bond's price rise or fall if the interest rate rises to 20 percent? What do your answers imply about the nature of the relationship (inverse or direct) between bond prices and the interest rate?

17.18. A speeding car swerves to avoid a rattlesnake slithering across a road but accidentally runs over and kills a lawyer walking along the side of the road. The lawyer was expected to make $90,000 per year over the next and final 10 working years of her life before her untimely demise. If the interest rate is 10 percent and the lawyer's family sues the driver, for how much can the family sue in terms of lost earnings due to the accident?

17.19. A Fijian island resort offers customers the opportunity to buy 30 vacation days at current prices with the opportunity to spend the purchased days on the island anytime over the next 5 years. The money on any days not spent on the island, plus 10 percent annual interest, is refundable at the end of the 5 years. If the interest rate over the same period on a typical money market fund is expected to be only 5 percent, is investing in the vacation days a good idea? Explain your answer.

17.20. Why might a medical doctor who has established a lucrative practice reduce her patient load and spend more time on the golf course or tennis court? Explain using a graph of a labor supply curve.

17.21. In 1999, tobacco companies agreed to settle a lawsuit brought by attorneys general from 46 states. According to the terms of the settlement, the companies are required to pay $206 billion to the states over 25 years in equal annual installments. Explain why the net present value of the settlement to the states is likely to be far less than $206 billion.

17.22. At any one point in time the supply of lawyers in the United States is relatively fixed. Since lawyers are an input to structuring new companies and business transactions, explain why the growth of the Internet economy generated some economic rents for lawyers. (*Note*: Whereas law firms historically had taken their cues regarding salaries from established New York partnerships, by the late 1990s Silicon Valley firms had become the trend-setters. In 1999, for example, several Silicon Valley law firms began offering starting salaries of $125,000 to first-year attorneys in contrast to the $91,000 offered by established New York partnerships. The New York partnerships had to play catch-up with their Silicon Valley counterparts.)

17.23. High-tech industries are not extensively unionized in the United States. Explain why this is the case, knowing that software programming, semiconductor manufacturing, and even customer service representative positions can be outsourced by companies such as Intel, Microsoft, and Amazon to locations in India, Singapore, and Malaysia. Rely on the model of input market monopoly developed in Section 17.6 in your answer.

17.24. On average, television news reporters are better looking than newspaper reporters. Explain why this might be the case, relying on the discussion in Section 17.4 regarding wage differences.

17.25. Workers in countries such as Sweden and Germany work fewer hours per week, on average, than do workers in the United States. Does this mean that Swedes and Germans are lazier than Americans? Explain how laziness could be interpreted in the context of the work–leisure choice model developed in this chapter. What other information might you want to gather to determine whether Swedes or Germans were indeed less work-oriented than Americans?

17.26. The biotech start-up Genentech has to pay a much higher interest rate on any funds it borrows relative to established pharmaceutical giant Eli Lilly and Company. Why would this be the case if both firms have access to the same capital markets and lending institutions?