

The Market for Labor

There are two types of decision makers in the economy: individuals and firms. Individuals supply factors of production, such as labor, to firms and demand output in return. Firms demand factors of production, use them to produce output, and supply that output to individuals.

In Chapters 3 and 4, we studied the demand for output by individuals and in Chapters 5 through 7, we studied the supply of output by firms. In Chapter 15 we studied the demand for inputs by firms; now in this chapter and the next, we will complete the picture by studying the supply of inputs by individuals.

In a competitive economy all prices and quantities are determined by the intersections of supply and demand curves. We know that the firm's supply of output and demand for inputs depend on available technology (encoded in marginal product curves, isoquants, and the like) and that the individual's demand for output depends on his tastes (encoded in indifference curves). Now we will discover that the individual's supply of inputs also depends on his tastes. It follows that ultimately all prices and quantities are determined by just two things: the technology available to firms and the tastes of individuals.

We will begin in Section 16.1 by studying individual labor supply curves, and then in Section 16.2 we will study equilibrium in the labor market. Sections 16.3 and 16.4 survey two special topics related to labor markets: Why do some people earn more than others? and What are the extent, causes, and effects of discrimination in labor markets?

16.1 Individual Labor Supply

Individuals supply labor to the market at a price called the wage rate of labor. We will begin by deriving an individual's labor supply curve.

Consumption versus Leisure

Each individual is endowed with 24 hours per day that he can allocate between labor and leisure. Labor consists of working in the marketplace for the going wage. **Leisure** consists of all other activities. Thus, leisure includes time spent on the beach, but it also includes time spent in productive activities such as going to school or looking for a better job.

Leisure

All activities other than labor.

Consumption

All goods other than leisure.

There are two goods relevant to the labor supply decision. One is leisure and the other is **consumption**. We use the word *consumption* to represent all the goods that can be purchased in the marketplace. Thus, *consumption* plays the same role that “all other goods” plays in the derivation of individual demand curves. Consumption stands for all goods other than leisure.

Consumption is often measured in dollars. We will find it more convenient to measure consumption in terms of the output good that the worker is producing. Thus, if he is a sausage maker, we will measure all consumption in terms of sausages.

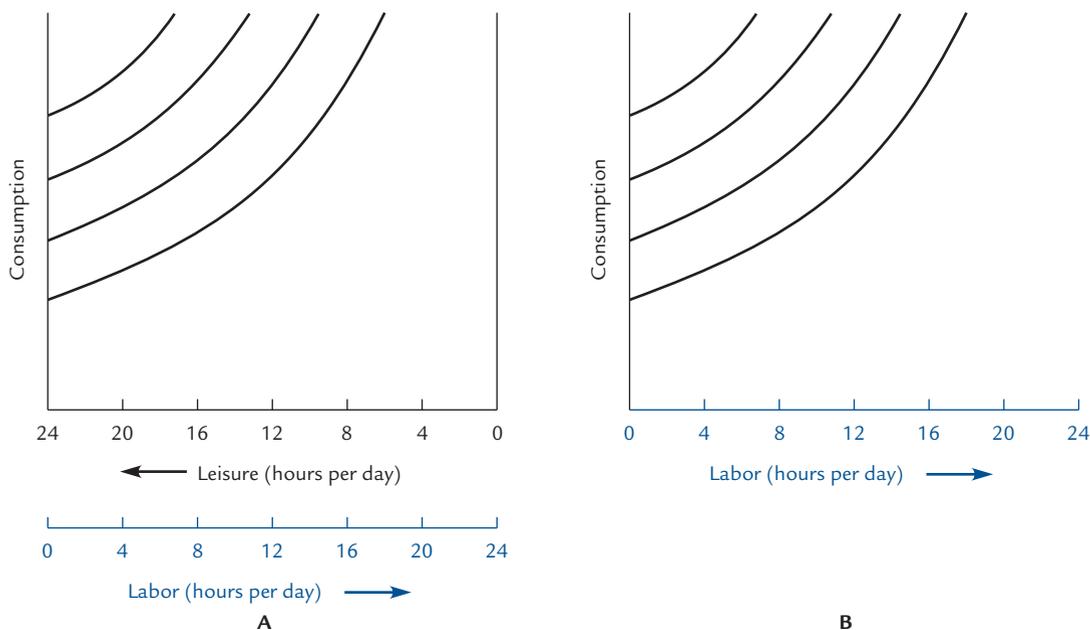
It is often useful to pretend that there is only a single consumption good in the economy, so that all workers receive their wages in the form of this single good.

Indifference Curves

We can draw indifference curves between leisure and consumption, and we will choose to draw them with the leisure axis running *from right to left*. This is pictured in panel A of Exhibit 16.1. Because it is not possible to have more than 24 hours of leisure per day, we have drawn a vertical barrier at the 24-hour mark. The number of hours that the individual devotes to labor is given by 24 minus the number of hours he devotes to leisure. This is indicated in the graph by the second row of labels on the horizontal axis.

EXHIBIT 16.1

Consumption versus Leisure



Panel A shows indifference curves between the two goods leisure and consumption, with the leisure axis running from right to left. Because of the reversed axis, the indifference curves appear to slope upward.

The alternate axis in panel A is the labor axis, since the amount of labor supplied per day is always 24 hours minus the amount of leisure taken. Panel B is a duplicate of panel A, with the leisure axis eliminated and only the labor axis shown.

In panel B of Exhibit 16.1 we have reproduced panel A without the right-hand vertical axis and with only the labor markings on the horizontal. This panel depicts the individual's indifference curves between labor and consumption. They are upward sloping, reflecting the fact that labor is considered undesirable. The slope of an indifference curve at any point is the amount of consumption needed to just compensate the worker for an additional hour of labor. It is the marginal value of leisure, measured in terms of consumption.

In Exhibit 3.16, where we studied the effects of a head tax and an income tax, we ran the leisure axis from left to right, rather than from right to left as we will in the present chapter. The choice of a direction for the axis is purely a matter of convenience and does not affect the substantive analysis in any way.



Dangerous
Curve

Exercise 16.1 Use the observation of the preceding sentence to explain why the indifference curves become steeper as you move up and to the right.

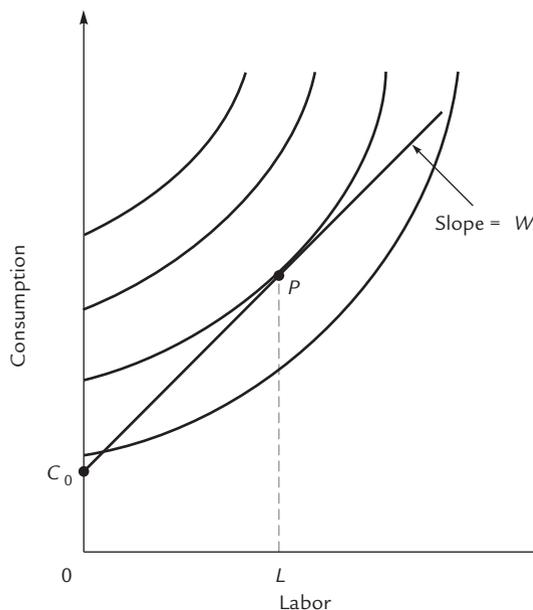
In Exhibit 16.2 we have added the budget constraint. When the individual does not work at all, he earns an income of C_0 . This **nonlabor income** is a return to some asset owned by the individual, such as an apple tree, a portfolio of stocks, a small business, or a pension. The slope of the budget line is equal to the wage, which we call W .

Nonlabor income

Income from sources other than wages.

EXHIBIT 16.2

The Worker's Optimum



The budget constraint is determined by C_0 , which is the worker's income from sources other than labor, and the wage rate W , which gives the slope of the budget line. The optimum is at P , where the worker supplies L units of labor. Here the wage rate (the slope of the budget line) is equal to the marginal value of leisure (the slope of the indifference curve).

If consumption is measured in sausages, then W is measured in sausages per hour. Each additional hour of labor yields W additional units of consumption.

The worker chooses his optimum point, which is at a tangency between an indifference curve and the budget line (point P in the exhibit). At the wage W the worker supplies L units of labor. At point P the wage rate (the slope of the budget line) is equal to the marginal value of leisure.

Exercise 16.2 Justify the worker's choice on economic grounds: If the wage were either more or less than the marginal value of leisure, how could the worker improve his position?

Changes in the Budget Line

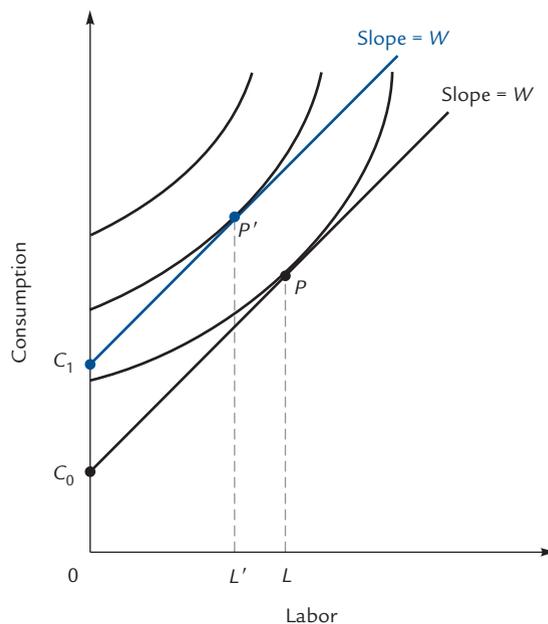
The worker's budget line changes if either his nonlabor income C_0 or his wage rate W changes. We will now study how the worker's optimum is affected by each of these possibilities.

Changes in Income

Exhibit 16.3 shows the effect of an increase in the worker's nonlabor income from C_0 to C_1 . The new optimum is at P' . If both consumption and leisure are normal (as opposed

EXHIBIT 16.3

An Increase in Nonlabor Income



When nonlabor income increases from C_0 to C_1 , the worker's budget line shifts upward parallel to itself. The new optimum is at point P' . If consumption and leisure are both normal (as opposed to inferior) goods, then P' lies above and to the left of P . Thus, an increase in nonlabor income leads to increased consumption and less labor supplied. The quantity of labor that this worker supplies falls from L to L' .

to inferior) goods, then the worker will choose more of each in response to his higher income; that is, P' will be above and to the left of P . Although it is logically possible for P' to be either below or to the right of P , we will assume that the income effects work in the expected directions, as in the exhibit. With this assumption:

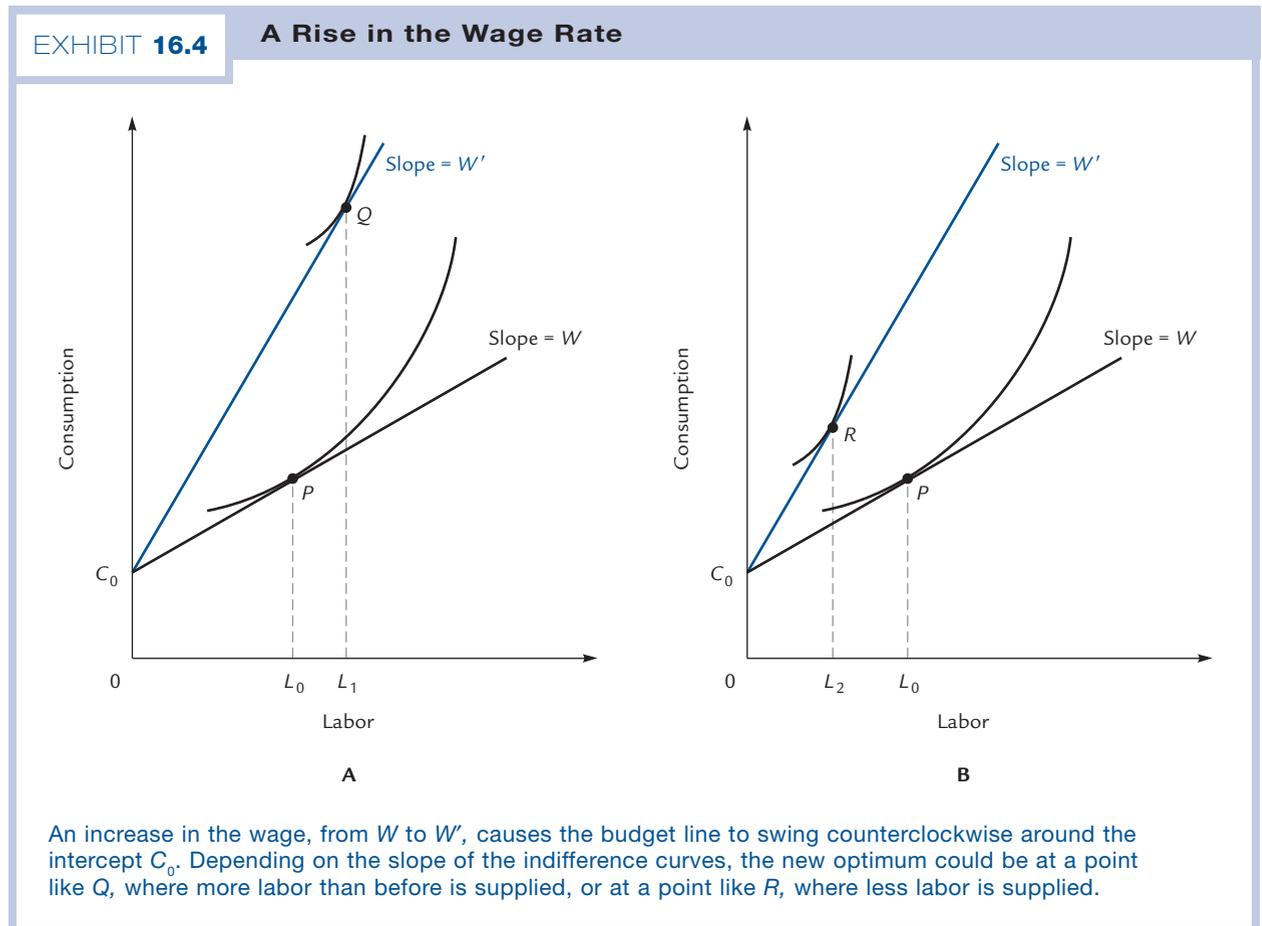
An increase in nonlabor income leads to a fall in the quantity of labor supplied.

An Increase in the Wage Rate

Suppose that the wage rises from W to W' while nonlabor income stays fixed. This has the effect of making the budget line steeper. Because there is no change in nonlabor income, the budget line swings through its intercept with the vertical axis. Exhibit 16.4 shows two possible outcomes. The optimum basket moves from P to Q in panel A of the exhibit or from P to R in panel B.

Income and Substitution Effects

When the wage goes up, there is both a substitution effect and an income effect. The substitution effect is that an additional hour of leisure is now more expensive in terms of forgone consumption. To say the same thing another way, additional consumption is now less expensive in terms of forgone leisure. In consequence of the substitution effect, the worker chooses more consumption and less leisure. Because he chooses less leisure, he supplies more labor.



The rise in the wage also has an income effect in that it makes suppliers of labor better off. As in Exhibit 16.3, we assume that both consumption and leisure are normal goods, so that the income effect leads the worker to choose more of both. Because the income effect leads the worker to choose more leisure, he supplies less labor.

Both the income and substitution effects lead to an increase in consumption (an upward movement in the consumer's optimum). These effects reinforce each other, and we can conclude that the new optimum (Q or R in the two panels of Exhibit 16.4) will be higher than the old optimum (P in either panel).

Regarding leisure, the income and substitution effects are at cross-purposes. The higher wage elicits more labor via the substitution effect, but it also makes the worker richer, eliciting more leisure (hence less labor) via the income effect. Either effect can dominate, so that the new optimum can be either to the right of P (as in panel A of Exhibit 16.4) or to the left of P (as in panel B). The worker might supply either more or less labor when the wage rate increases.

The Income and Substitution Effects via Geometry

We can use a graph to sort out the income and substitution effects. After the wage rises from W to W' , we imagine a downward adjustment in the worker's nonlabor income that just compensates for the wage increase, leaving him on the same indifference curve as before. This gives a compensated budget line, shown in color in each of the panels in Exhibit 16.5. We now imagine the movement to the new optimum as taking place in two steps: from P to Q' to Q in panel A or from P to R' to R in panel B. The first movement is the substitution effect and must be upward and to the right (it is a movement along an indifference curve to a steeper point). The second movement is the income effect, as in Exhibit 16.3, which is a movement upward and to the left.

In panel A, the substitution effect is greater than the income effect, while in panel B the reverse is true. Thus, in panel A the wage increase leads to more labor supplied, and in panel B the wage increase leads to less labor supplied.

Comparing the Two Effects

Which is larger, the income effect or the substitution effect? First, consider the situation when the wage is very low. In that case, the worker supplies very little labor (for example, if the wage is zero, there is no incentive to work at all!). Therefore, a change in the wage has little effect on the worker's income, so the income effect is negligible. It follows that at low wages, the substitution effect dominates the income effect, as in panel A of Exhibit 16.5. Therefore:

When the wage is very low to begin with, an increase in the wage leads to an increase in labor supplied.

When the wage rate is high, both the income and the substitution effects can be substantial. Therefore, at high wage rates there is no way to tell which effect will dominate.

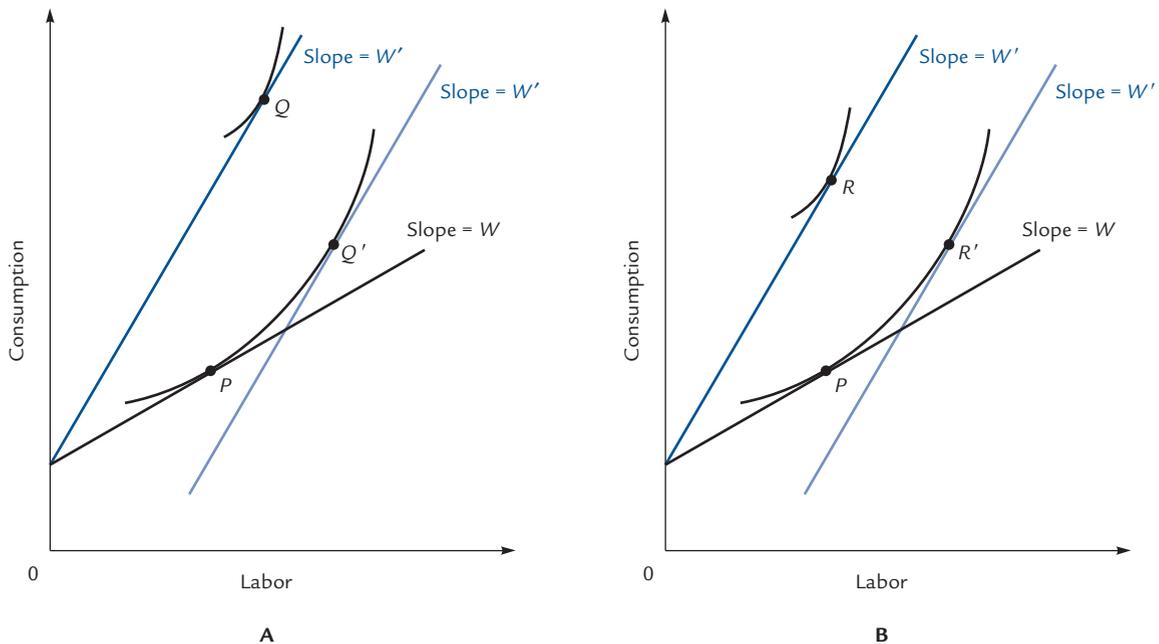
The Worker's Supply of Labor

Deriving the Labor Supply Curve

From graphs like those in Exhibit 16.4, we can derive the labor supply curves of individuals. Exhibit 16.6 depicts the labor supply curves of the two individuals whose indifference curves appear in Exhibit 16.4. Both curves slope upward at low wages,

EXHIBIT 16.5

Income and Substitution Effects



The effect of a wage increase can be decomposed into a substitution effect followed by an income effect. When the wage goes up, we pretend that the worker loses just enough nonlabor income to keep him on his original indifference curve. In either panel this yields the lightly colored budget line. The substitution effect is from P to Q' in panel A or from P to R' in panel B; it is a movement along the indifference curve and leads to more labor supplied.

The income effect is from Q' to Q in panel A or from R' to R in panel B. It leads to less labor supplied.

In panel A the substitution effect dominates the income effect, so that more labor is supplied after the wage increase. In panel B the opposite is true.

reflecting the dominance of the substitution effect over the income effect. The second curve “bends backward” at higher wages, to reflect the fact that for this individual the income effect eventually comes to dominate the substitution effect. An individual’s labor supply curve might or might not be backward-bending.

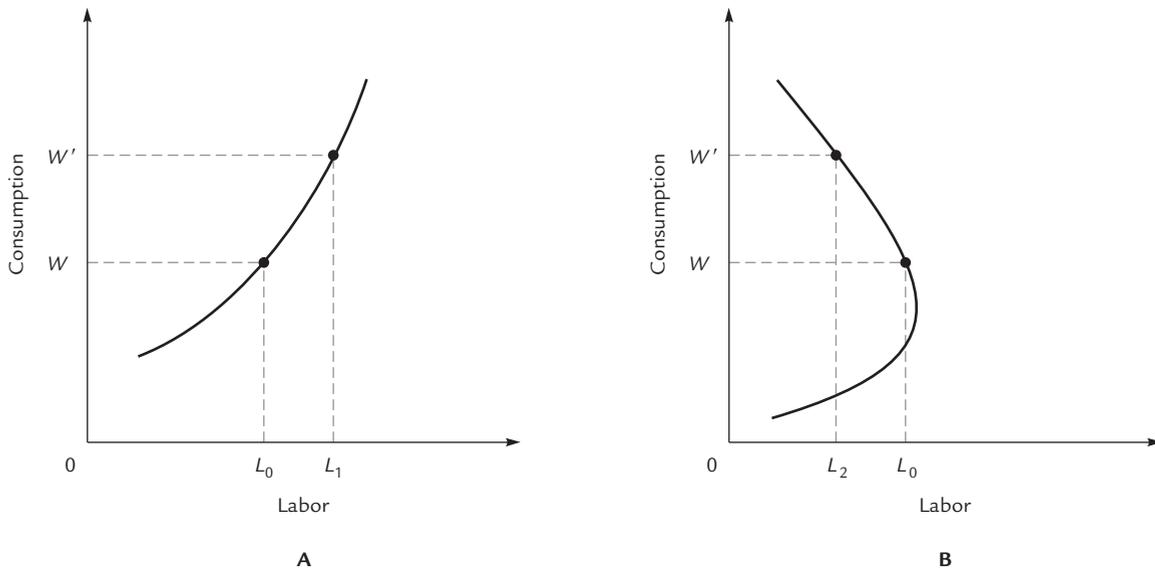
Using the Labor Supply Curve

Changes in wage rates correspond to movements along the labor supply curve, whereas changes in other things, such as nonlabor income, correspond to shifts of the curve.

Since the early days of the Industrial Revolution, wage rates have increased substantially and, at the same time, the quantity of labor supplied has decreased. The 60-hour workweeks that were common for unskilled laborers 100 years ago are uncommon today. This evidence is consistent with a backward-bending labor supply curve. However, there is an alternative explanation. Along with the increase in wages has come a substantial increase in nonlabor income. As you can see from Exhibit 16.3, an increase in nonlabor income leads to less labor supplied at any given wage; that is, it causes the labor supply curve to shift leftward. Thus, the fall in hours worked might be explained by an upward-sloping labor supply curve that has shifted leftward, as in Exhibit 16.7.

EXHIBIT 16.6

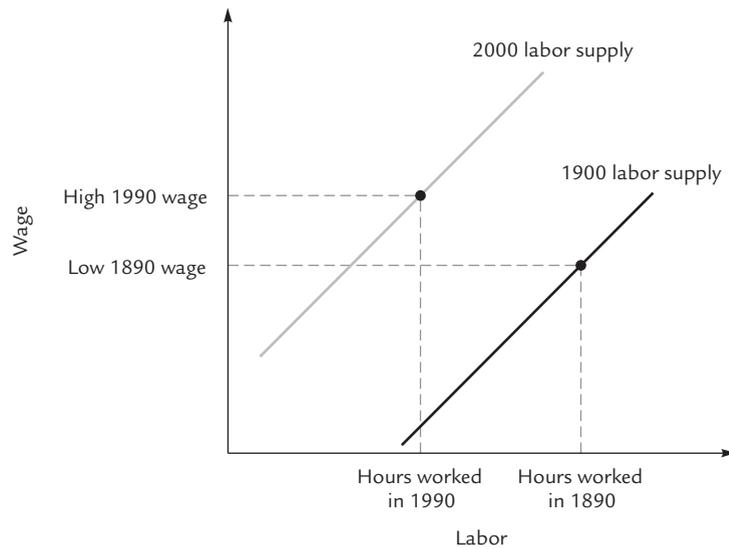
The Individual's Labor Supply Curve



The graphs show the labor supply curves of the two individuals whose indifference curves are depicted in Exhibit 16.4. The enlarged points here are derived from the points P , Q , and R in that exhibit.

EXHIBIT 16.7

A Rise in the Wage Accompanied by a Rise in Nonlabor Income



Over the last 100 years both wage rates and nonlabor income have increased. The rise in nonlabor income causes the labor supply curve to shift leftward, as shown. This could explain the observed fall in the quantity of labor supplied.

International Differences in Labor Supply

The average American works 25 hours a week; the average Frenchman, 18; the average Italian, a bit more than 16 1/2. Even the hardest-working Europeans—the British who put in an average of 21 1/2 hours—work far less than their American cousins.

Several factors influence these averages. First, Europeans are more likely than Americans to be unemployed (and hence to work 0 hours a week). Next, even for the employed, workweeks are shorter in Europe; employed Americans put in about 3 hours more per week than employed Frenchmen. More importantly, Europeans take much longer and more frequent vacations. The average employed American takes less than 6 weeks of vacation per year; the average Frenchman takes 12. The world champion vacationers are the Swedes, who average 16 1/2 weeks of vacation per year.

Why do Americans choose to work so much? Or, if you prefer, why do Europeans choose to work so little?

The answer is unlikely to involve cultural differences, because all of this is quite recent. Just 30 years ago, Europeans worked slightly *more* than Americans. So the right question is not “Why is Europe different?” but rather “What changed?”

According to Nobel laureate Edward Prescott¹, the answer is tax policy. Thirty years ago, European and U.S. marginal tax rates were comparable—and so were European and U.S. labor supplies. Between the 1970s and the 1990s, the U.S. marginal rate stayed fixed at about 40%, while the French rate rose to 59% and the Italian rate to 64%. On a country-by-country basis, steeper marginal tax hikes are closely correlated with shrinking workweeks and expanding vacations.

Case closed? Not quite. The problem is that a 20-percentage-point increase in your marginal tax rate is essentially equivalent to a 20% pay cut. We have lots of data on how strongly people react to 20% pay cuts, and by and large those responses are nowhere near as dramatic as the changes we’ve seen in Europe. So although marginal tax rates do a good job of explaining *relative* changes (the countries with the biggest tax hikes have the biggest labor supply contractions), they do a poor job explaining *absolute* changes—that is, they can’t explain why labor supplies across Europe have fallen so far.

One trio of economists² offers the following theory: When your own wages are cut by 20%, you’ll take more vacations. But when your *friends’* wages are cut by 20%, you’ll take even *more* vacations, because vacations are more fun when you have friends to share them with. So a 20% across-the-board tax hike, which affects both you and your friends, yields a more dramatic response than a 20% cut in your own wages.

16.2 Labor Market Equilibrium

We have now constructed a single individual’s labor supply curve. Repeating this construction for each individual separately and then adding up, we can construct a market labor supply curve. At any given wage, we read each individual’s quantity supplied from his own supply curve; then we add these quantities to get the quantity of labor supplied to the market.

There is also a market demand for labor, which we know from Chapter 15 coincides with the MRP_L curve. Putting the supply and demand curves together, we can find

¹ E. Prescott, “Why Do Americans Work So Much More Than Europeans?,” *Federal Reserve Bank of Minneapolis Quarterly Review* 28, July 2004.

² A. Alesina, E. Glaeser, and B. Sacerdote, “Work and Leisure in the U.S. and Europe: Why So Different?,” *Harvard Institute of Economic Research Discussion Papers* 2068, April 2005.

the point of labor market equilibrium. Now we can use the machinery of supply and demand to analyze the effects of some simple changes.

To carry out the exercises of this section, we will assume that the labor supply curve is upward sloping.

Changes in Nonlabor Income

Part A of Exhibit 16.8 illustrates what happens when a single worker experiences an increase in nonlabor income, such as an unexpected inheritance. The labor market (depicted in the left-hand panel) is in equilibrium at the wage W . The individual worker, who is a competitive supplier in the labor market, then faces a flat demand curve for his services at that wage. His initial supply curve is s in the right-hand panel and he supplies L units of labor.

When he learns about his inheritance, the worker feels wealthier and decides to work less, so his labor supply curve shifts back to s' . Because he represents an insignificant part of the market, the market curves in the left-hand panel do not move, and neither does the flat demand curve d . The worker now supplies L' units of labor instead of L .

Part B of the same exhibit illustrates what happens when many workers experience a simultaneous increase in their nonlabor income. This could happen for a variety of reasons. Perhaps a lasting peace in the Middle East brings down the price of oil (which effectively increases the wealth of oil consumers); perhaps the same lasting peace leads to a reduction in U.S. military expenditures and a consequent reduction in taxes; perhaps (as happened in 1991) Rocky and Bullwinkle are released on videotape, yielding a widespread improvement in standards of living. Now the individual supply curve in the right-hand panel shifts back from s to s' as before. The new wrinkle is that the shift in supply occurs for every individual, not just one, and therefore, the market supply curve shifts back as well, from S to S' in the left-hand panel. The market wage rises from W to W' . The total quantity of labor supplied to the market certainly falls, but a given worker, moving from the intersection of s and d to the intersection of s' and d' , can either increase or decrease the quantity he supplies.

To sum up, a marketwide increase in nonlabor income can lead to either an increase or a decrease in any one individual's working hours, but the *average* or "representative" worker must decide to work less. We know this because the total quantity of labor supplied to the marketplace must fall.

Changes in Productivity

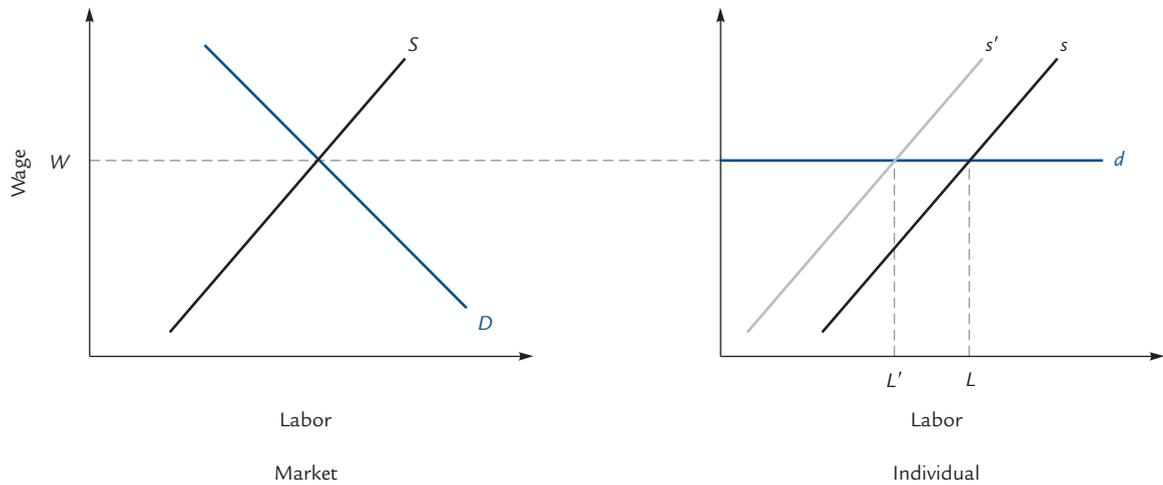
Workers can be made more productive in many ways, including technological advances (like faster computers, which improve the productivity of office workers), improvements in the weather (which improves the productivity of agricultural workers), or unexpected disasters (which can improve the productivity of medical personnel).

Panel A of Exhibit 16.9 shows the effect of an increase in marginal productivity. The labor demand curve, which coincides with the MRP_L curve, shifts rightward from D to D' in the left-hand panel. The market wage rises from W to W' , and the flat demand curve for the services of an individual worker rises accordingly, from d to d' in the right-hand panel. Each individual worker supplies more labor than before.

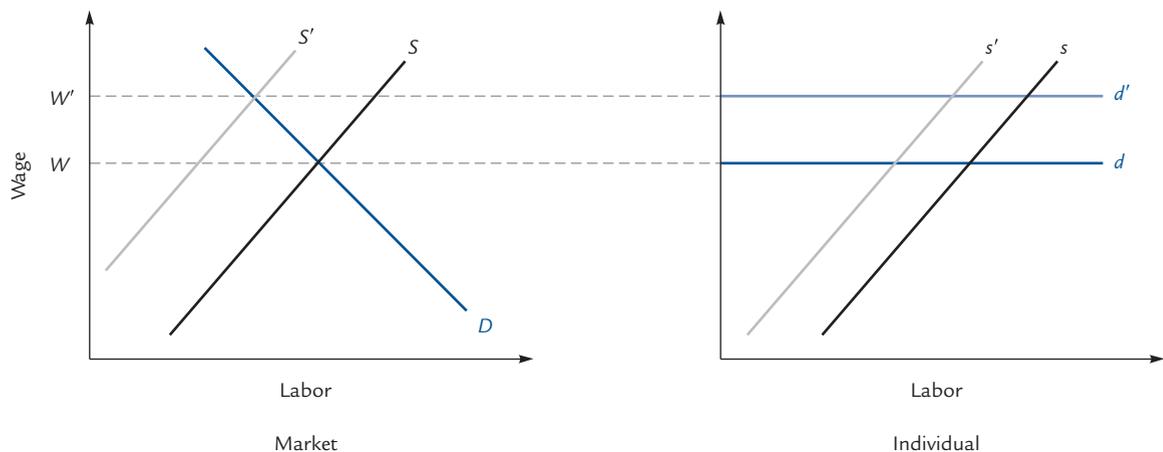
That analysis holds workers' nonlabor income fixed. But when there is an economy-wide increase in productivity, such an analysis is likely to be incomplete, because workers' nonlabor income is likely to rise. Here's why: When workers become more productive, the value of capital increases. (For example, a factory employing highly

EXHIBIT 16.8

An Increase in Nonlabor Income



A. An Increase in One Worker's Nonlabor Income



B. An Increase in All Workers' Nonlabor Income

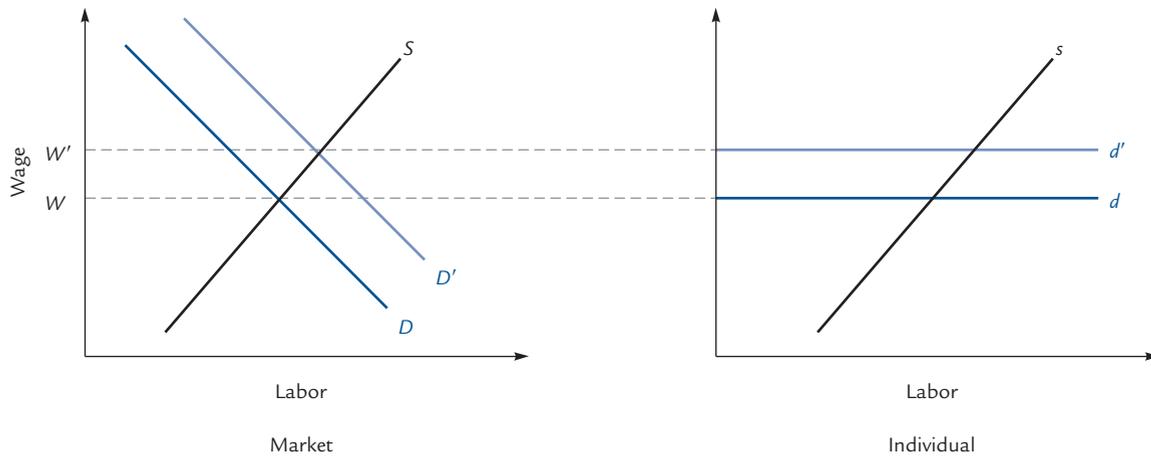
Part A depicts the effect of an increase in a single worker's nonlabor income. The market is in equilibrium at a wage of W . The worker faces a flat demand curve at that wage. When his income increases, the worker's supply curve shifts back from s to s' and his quantity of labor supplied falls from L to L' .

Part B depicts the effect of a simultaneous increase in all workers' nonlabor income. The individual supply curve shifts back from s to s' as before. Since *all* workers' supply curves shift, the market supply curve shifts back also, from S to S' . The wage rises from W to W' . The quantity of labor supplied to the market falls; the quantity supplied by a given individual can either rise or fall.

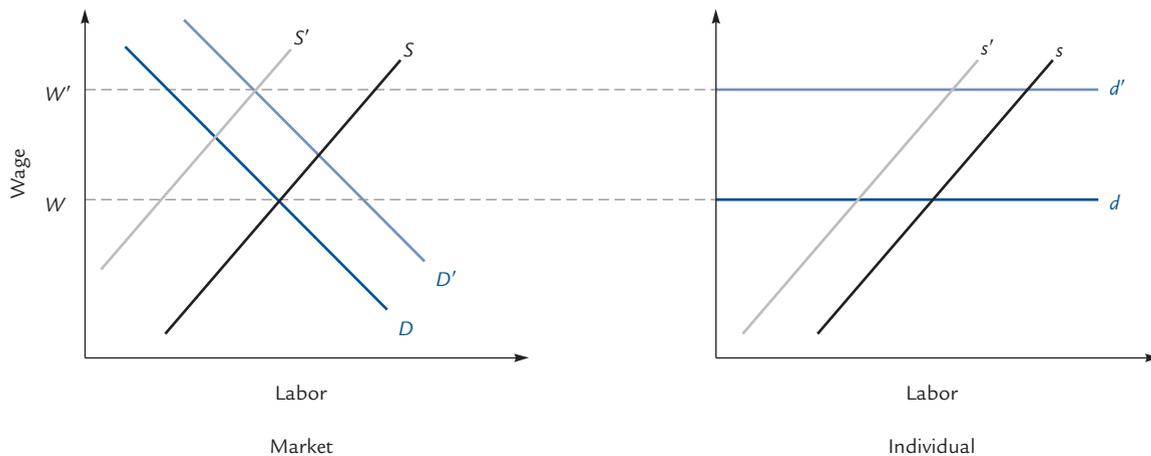
productive workers is worth more than a factory employing less productive workers.) Therefore, the owners of capital experience an increase in nonlabor income. But in many cases, the owners of capital include the workers themselves: Farmers own tractors; plumbers own plumbing tools; and many workers own stock in corporations that in turn own all sorts of capital equipment.

EXHIBIT 16.9

An Increase in Marginal Productivity



A. An Increase in Marginal Productivity



B. An Increase in Productivity That Increases Workers' Wealth

Part A shows the effect of an increase in marginal productivity. The market demand for labor moves out from D to D' ; the wage increases from W to W' ; and the quantity of labor supplied, both by the market and by the individual, increases.

Part B shows the effect of an increase in marginal productivity that causes an increase in workers' non-labor income (by increasing the value of the capital that they own). Demand in the left-hand panel moves out as in part A. Supply in both panels moves back as in part B of Exhibit 16.8. The wage increases from W to W' , and the quantity of labor supplied, both by the entire market and by the individual, moves ambiguously.

Therefore, a general increase in productivity is likely to yield an increase in workers' nonlabor income. In that case, Panel A in Exhibit 16.9 must be replaced by panel B. Here the demand for labor shifts out, just as in panel A, but at the same time, the market and individual labor supply curves shift leftward, as in Exhibit 16.8. The wage rate increases, but the quantity of labor supplied, both by the entire market and by the individual, moves ambiguously.

Temporary Changes in Productivity: Intertemporal Substitution

In March 1989, the *Exxon Valdez* oil tanker ran aground in Alaska, creating an oil spill of historic proportions. Clean-up operations were urgent, so the marginal productivity of Alaskan workers quickly jumped up.

According to panel A of Exhibit 16.9, there should have been increases in both the wage rate and the quantity of labor supplied by each worker. Both these predictions were borne out. Wages quickly rose (from about \$9 an hour to about \$10.60 an hour), and at the same time, the average workweek shot up from about 35 hours a week to about 49 hours a week.

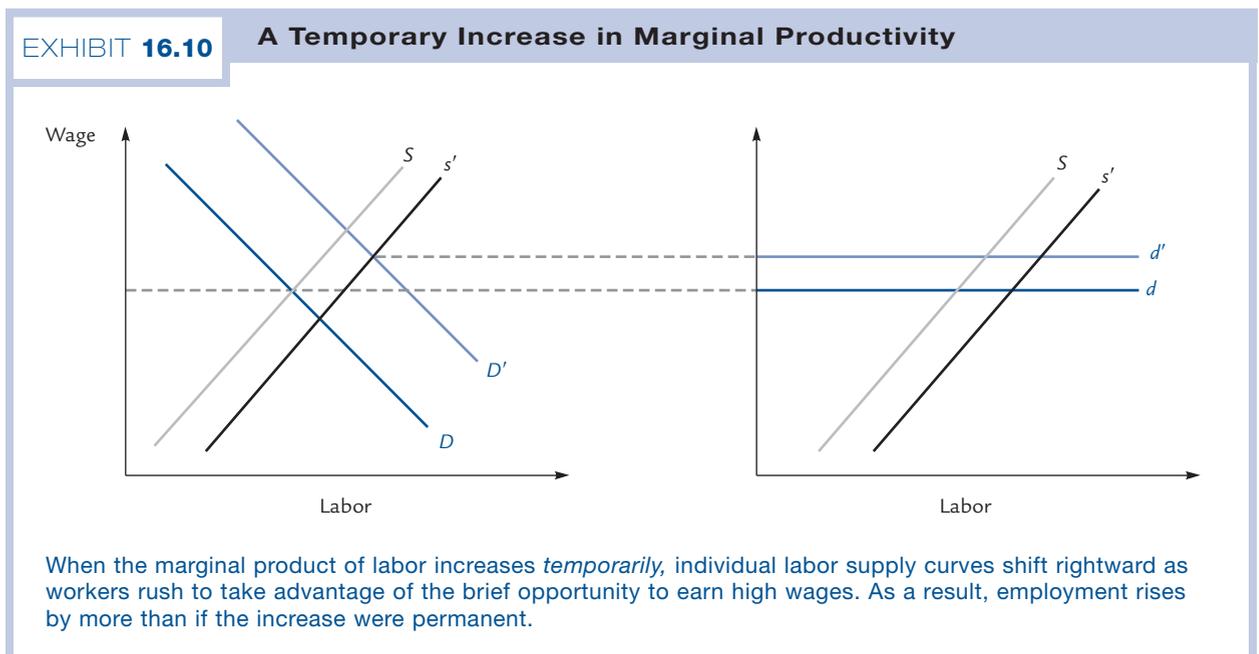
What is surprising here is not that the average workweek increased, but that it increased so dramatically. In fact, Alaskan wages had been as high as \$10.60 an hour just a few years earlier, but at that time average workweeks were just slightly more than 40 hours a week. To put this another way, the average worker's labor supply curve, which had recently passed through (\$10.60, 40) now passed through (\$10.60, 49). So the labor supply curve must have shifted far to the right in 1989. What could have caused that shift?

The answer is that the high productivity of Alaskan workers was *temporary*. When there is a temporary opportunity to earn high wages, workers often rush to take advantage of it, working extra hard during the brief window of opportunity and postponing leisure time. (For example, people who sell Christmas trees tend to work very hard in December and compensate by relaxing in January.) Such behavior is called **intertemporal substitution**.

Intertemporal substitution leads to a rightward shift in the labor supply curve, as shown in Exhibit 16.10. The labor demand curve in the left-hand panel shifts rightward as workers become more productive, just as in panel A of Exhibit 16.9. But now, because the opportunity to earn high wages is temporary rather than permanent, laborers rush to take advantage of this brief opportunity and their labor supply curves

Intertemporal substitution

Working additional hours during temporary periods of high productivity.



also shift rightward. The rightward shift in individual labor supply (in the right-hand panel) implies a rightward shift in the marketwide labor supply (in the left-hand panel).

Because the demand and supply curves both shift rightward in the left-hand panel of Exhibit 16.10, it appears that the equilibrium wage rate can either rise or fall. But a fall in the equilibrium wage rate would be inconsistent with the story we've been telling: The entire reason the supply curve shifts is so that workers can take advantage of temporarily high wages. This can happen only if wages *are* temporarily high; thus, the new (temporary) equilibrium must be higher than the old equilibrium to which the market will eventually return.

The conclusion is that, because of intertemporal substitution, a *temporary* increase in productivity has a much bigger effect on employment than a *permanent* increase in productivity. Therefore, intertemporal substitution might be an important factor in determining the severity of recessions (that is, temporary periods in which average income is low, and, typically, unemployment is high).

16.3 Differences in Wages

We have discussed the determination of “the” market wage. Yet it is a common observation that different people earn different wages. In this section, we will discuss some of the reasons for these differences.

Human Capital

A firm that hires an employee is often hiring not just raw *labor*, but an entire package of productive skills. Some of those skills, like intelligence, may be innate, whereas others, like education and training, are the result of investments by the employee earlier in life. Such skills can productively be viewed as a form of capital, which we will call **human capital**.

We have seen in Section 15.4 that the revenues of the firm are divided among the productive inputs, with each earning its marginal product. A worker who brings both labor and human capital to an enterprise earns both the wage rate for his labor and the market rate of return for his skills. In practice, he usually receives the sum of these returns in a single paycheck, the size of which is described as his *wage*. Of course, workers with different amounts of human capital will earn differing returns.

The use of the word *capital* here is more than just a loose metaphor. As we will see in Chapter 17, capital consists of productive resources that have themselves been produced by forgoing consumption at earlier times. This description fits human capital perfectly. When you attend college, you forgo current consumption, both by making tuition payments that could be used for other things and by allocating time to your studies that could otherwise be spent earning income. The sum of these costs is an investment in human capital.

In the short run, human capital is a fixed factor (its supply curve is vertical). For this reason, payments to human capital are a form of rent. The difference between the earnings of a college graduate and those of an unskilled laborer constitutes the rent on human capital.

In the long run, people can vary their investments in human capital. As more investment takes place, the costs (like college tuition) are driven up and the rents to human capital are driven down. People will continue to invest until the marginal cost and marginal benefit from a unit of human capital are equal.

Human capital

Productive skills.

If all people can make equally productive use of an education, then everyone will be indifferent between becoming educated and not becoming educated. This is because the cost of an education will exactly offset the benefits. (If the benefits of going to college exceed the costs, additional people will enter college until this is no longer the case.) If, on the other hand, people are endowed with varying quantities of other skills (like intelligence or perseverance) that make education more productive, then those who have unusually large endowments of these other skills can benefit from education.

Signaling

In Chapter 9, we discussed the phenomenon of *signaling*, whereby a college education can lead to higher wages even without contributing to productivity, provided that it helps employers identify people with intelligence and perseverance. Education can lead to higher wages either by adding to human capital or by performing a signaling function; in practice, both aspects are surely present.

Education as Consumption

We have used education as an example of an investment in higher wages. We have suggested two ways in which this could happen. Perhaps education is a way to acquire human capital; perhaps it is a signal of certain innate skills; perhaps it is some combination of the two.

In fact, highly educated people do earn higher wages than do less highly educated people. However, there is yet another possible explanation for this. Rather than education causing high wages, perhaps high wages cause education.

Suppose that people actually enjoy going to college and view it as a consumption good. Then we expect people with greater wealth to consume more of this good. Just as richer people buy more Rolls Royces, so richer people buy more education. No one would suggest that because rich people drive Rolls Royces, buying a Rolls Royce will make you rich.

Undoubtedly, education is partly investment and partly consumption. To some extent, people purchase it to raise their incomes, and to some extent they purchase it because they enjoy it. Here is a question to ponder: What observable data would help you determine what percent of educational spending is pure consumption?

Compensating Differentials

Another reason for differences in observed wage rates is that some jobs are more pleasant or less pleasant than others. When there is a large class of equally talented workers available to each of several occupations, these workers must be indifferent as to which occupation they choose.

Exercise 16.3 Why must the workers be indifferent among occupations? If they were not indifferent, what would happen?

There are many reasons why one occupation might be inherently less pleasant than another. In some occupations the work itself is unpleasant, in others the people employed command less respect, and in still others there are greater degrees of risk. In order for workers to remain indifferent, the less pleasant occupations must pay more. We can view the wage in the less pleasant occupation as the sum of the market wage determined elsewhere plus an additional payment to compensate the

Compensating differential

A wage adjustment that comes about in equilibrium to compensate for a particularly pleasant or unpleasant aspect of a job.

worker for the unpleasant aspects of his job. This additional payment is known as a **compensating differential**.

Other occupations are unusually attractive. An employee in such an occupation earns less than one in a more typical job, the compensating differential being negative. For example, many positions offer workers the opportunity to invest in human capital at a cost much lower than the usual market rate. This comes about when an employee, in the course of performing his duties, acquires skills that he will later be able to sell in the marketplace. Such on-the-job training occurs at every level of skill. A postdoctoral instructor in physics at a top university is gaining valuable skills that will increase his marketability in later life, in exchange for which he accepts a wage that might be less than his marginal product. A clerk in a bookstore is observing and learning the business, gaining the skills necessary to be a manager or to open his own shop someday. Again, he pays for this opportunity through a lower wage.

Although on-the-job training is important at every level, it is particularly important at the very bottom of the career ladder, where the skills that are mastered (fundamentals such as knowing the importance of showing up for work on time and how to get along with co-workers) will be useful in any future occupation. In entry-level positions, on-the-job training is often a substantial portion of the employee's total compensation.

Access to Capital

Wages would also differ if workers had access to capital of differing qualities. A secretary in New York City using the latest word processor might be more productive at the margin than a secretary using a manual typewriter in a locality with no electricity.

In making this argument, it is important not to confuse total productivity with marginal productivity. The lone secretary with the manual typewriter in a developing country can certainly be more productive *at the margin* than the 100,000th word-processing New Yorker. In fact, as long as people can move from country to country, wages will tend to become equal everywhere over time because of people leaving the low-wage countries to enter the high-wage countries. This equalization of wages implies an equalization of marginal products. Therefore, the access of different workers to different sorts of capital can explain wage differences in the long run only if there are barriers to the mobility of workers, such as immigration restrictions.

However, even immigration restrictions fail to explain wage differences across countries. If wages are lower in Mexico than in the United States, we at first expect Mexican workers to cross the border until wages are equalized. Then we are reminded that the immigration laws prevent this. But now we should expect U.S. firms to move their capital across the border into Mexico to take advantage of the low wages there. This will raise wages in Mexico and reduce wages in the United States, and the flow of capital across the border should continue until wages are equalized.

We do see some phenomena like this. In recent years, for example, many firms have relocated from the northern to the southern United States to take advantage of lower wage rates. But there has been nothing like the international movement of capital that one would expect on the basis of standard economic theory. Why not?³

³ This riddle was posed by Robert E. Lucas, Jr., in a series of lectures titled "On the Mechanics of Economic Development," *Journal of Monetary Economics* 21 (1988). The answer we will propose is also taken from those lectures, although it is offered there as a clue to the solution of a much deeper riddle, namely: Why do different countries have different levels of economic development and different rates of growth?

An answer to this riddle might be found in the external effects of human capital accumulation.⁴ When you invest in training or education, you increase not only your own productivity but that of your fellow workers through a variety of complicated interactions between you and them. Perhaps some of your new knowledge rubs off in conversations around the water cooler. Perhaps you are more likely to make suggestions or to have ideas that other workers can imitate or that will inspire them to formulate related new ideas of their own. These interactions need not be confined to your own workplace. To paraphrase Adam Smith, people of the same trade seldom meet together, even for merriment and diversion, without the conversation ending in a mutually beneficial exchange of ideas and methods or in some contrivance to increase efficiency.

Through such mechanisms, your accumulation of human capital can raise the productivity not only of your co-workers and of other workers in your industry but also of the physical capital with which you interact. In that case, those owners of physical capital who locate themselves in areas with large concentrations of highly trained people will reap a share of these external benefits. They might be willing to pay higher wages, or higher land rents, in exchange for such an opportunity. Consequently, the difference between land rents in, say, Manhattan and a more remote location might be a tolerably good measure of the value of those external benefits.

If human capital investment yields significant positive externalities, then there will be too little of it. People invest in human capital only up to the point where the marginal cost is equal to the marginal increase in their own productivity, without taking account of how their investment affects the productivity of others. This observation constitutes an efficiency-based argument for subsidizing investments in human capital, such as education. If, as we have argued, differences in land rents measure the value of human capital externalities, then the size of such rent differentials could be used in a calculation of the size of the optimal subsidy.

16.4 Discrimination

The average black person earns less than the average white person, and the average woman earns less than the average man. Parts of these differentials are easy to account for. The average African-American is about 6 1/2 years younger than the average white, and younger workers generally earn less than older workers do. A larger percentage of African-Americans live in the South, where wages are lower generally. Women are more likely than men to have studied sociology instead of engineering.

Economists disagree about whether such factors can account for all of the observed wage differentials. The alternative hypothesis is that the differentials are partly due to discrimination. The existence of discrimination is difficult to measure. One must ask not “Do African-Americans earn less than whites do?” but “Do African-Americans earn less than whites *with comparable market characteristics* [education, experience, age, etc.] do?”⁵ The question is empirical but difficult to settle, because of the difficulty of measuring all of the relevant market characteristics.

⁴ An alternative possible answer is that after adjusting for human capital differences, Mexican wages really *aren't* any lower than U.S. wages.

⁵ The real question is “Do African-Americans earn less than do whites with the same marginal product?” In view of the difficulty of measuring marginal product directly, we hope to approximate it with a mix of observable market characteristics.

Theories of Discrimination

If there is discrimination, employers engage in it at a cost. If African-Americans earn lower wages than equally productive whites, any employer who hires whites forgoes an opportunity to hire equally productive African-American labor at a lower wage.

In fact, a relatively small number of nondiscriminating employers could suffice to eliminate all wage differentials, even if the majority of employers discriminate. Suppose that 80% of employers are discriminatory and are unwilling to pay African-Americans more than half their marginal product. Suppose that the remaining 20% of employers are indifferent between hiring whites and hiring African-Americans and that these 20% are enough to employ all of the blacks in the economy. Then as long as African-Americans are paid less than their marginal product, the nondiscriminating firms will hire more of them. This will continue, bidding up the price of black labor, until African-Americans are earning their full marginal product, just as whites are.

It is sometimes alleged that employers discriminate not out of any genuine distaste for a particular group, but as a strategy to employ that group at a lower wage. Such a strategy would require the cooperation of thousands of employers and would be subject to exactly the same pressures that cause cartels to break down. Any individual employer could gain by cheating. In fact, such a strategy is far more implausible than a cartel, because a cartel requires cooperation only by the firms in a single industry, whereas the “fake discrimination” ploy requires the cooperation of all firms that hire labor.

One theory of discrimination says that while employers might be indifferent between hiring whites and African-Americans, they nevertheless discriminate because their white employees have a distaste for associating with African-Americans. Whenever a African-American is hired, the employer must increase the white workers' wages or they will leave the firm. Thus, because it is especially costly for employers to hire African-Americans, the demand for African-Americans is lower and they receive lower wages. If this theory is correct, employers should be able to benefit by hiring all-African-American workforces, paying the lower African-American wage without having to worry about the effect on white employees. Employers will adopt this strategy until African-American wages are bid up to the level of white wages. Thus, the theory predicts a heavily segregated workforce, with some all-white firms and some all-African-American firms, but no wage differentials.

Considerable sophistication is needed to find a theory consistent with sustained wage differentials in the face of profit maximization by even some employers. Since most theories predict a tendency toward complete segregation, it is necessary to postulate a force opposing that tendency in order to get realistic results. One possibility is that African-Americans and whites have different skills and that those skills are complementary in production. In this case, it would pay to combine African-American and white workers even if it required paying a premium to the whites. Another possibility is to develop a theory of the costs of changing personnel, so that an employer who would ultimately benefit from an all-African-American workforce will find it optimal to stretch the adjustment out over a long period of time.⁶

⁶ See K. Arrow, “Some Models of Racial Discrimination in the Labor Market” in A. Pascal (ed.), *Racial Discrimination in Economic Life* (RAND Corporation, 1972), for a survey and detailed discussion of such theories. The first serious attempt by an economist to study questions related to discrimination was in G. Becker, *The Economics of Discrimination* (Chicago: University of Chicago Press, 1957).

Wage Differences Due to Worker Preferences

Some apparent discrimination undoubtedly results from the preferences of the workers themselves. Here is an example of how this might come about.

When a worker seeks a job, he or she typically receives several offers at different salaries. Suppose that men and women typically receive the same range of offers, but that men on average are more inclined to accept their highest-paying offer, whereas women apply many other criteria in making their choice. In this case, statistics will show that women earn less than men do, even though men and women both receive exactly the same salary offers on average.

Why might men be more inclined than women to accept their highest-paying offers? One reason is that most married men are trained for more lucrative occupations than their wives are. Thus, if a married couple must live together in the same city, they usually maximize their total family income by moving to the city where the husband has the brightest prospects.

Imagine, for example, a couple in which the husband is a movie director and the wife a professor. The husband is offered a \$100,000 job in California and a \$50,000 job in Massachusetts. The wife is offered a \$10,000 job in California and a \$20,000 job in Massachusetts. In this case the couple maximizes its income by moving to California, where their combined salaries are \$110,000 instead of \$70,000. The wife will earn \$10,000, whereas most male professors (who are not married to movie directors) will live in Massachusetts and earn \$20,000.

Statistics will show that female professors generally earn less than their male counterparts do, while perhaps failing to show the reason why. The point of this example is not its empirical significance, which at any rate is unclear.⁷ The point is that wage differentials can result from supply decisions (by workers) as well as from demand decisions (by employers) through subtle mechanisms that might not be apparent to the researcher. This is why questions about discrimination are so hard to settle.

Human Capital Inheritance

If it is argued that African-Americans earn less than whites only because of inferior human capital, one must still attempt to account for this interracial difference in human capital. A common explanation is that human capital is largely inherited (we learn much from our parents' skills and attitudes) and that African-Americans have inherited less because of past discrimination. Of course, this is scant comfort to an African-American worker who is informed that he earns less than his white colleagues not because he is African-American, but because his parents were. Yet it surely does make a difference whether African-Americans and other groups are suffering only from past discrimination or from present discrimination as well. Although two diseases have the same symptoms, the prescribed medications could differ substantially.

Although past discrimination, via human capital inheritance, might play a role in determining the current incomes of African-Americans, it is at least reasonably certain that this is not true of women. African-American people tend to have mostly African-American ancestors, but women have only the same percentage of female ancestors that their brothers do.

⁷ For some evidence, see R. Frank, "Why Women Earn Less: The Theory and Estimation of Differential Overqualification," *American Economic Review* 68 (1978): 360–373.

Summary

Individuals supply labor to firms, which produce outputs that individuals demand. Labor supply, like output demand, depends on the tastes of individuals.

Thus, we need to study the individual's indifference curves between consumption and labor. We can begin by drawing his indifference curves between consumption and leisure, which are both goods, and then reversing the leisure axis.

The budget line is determined by nonlabor income (which gives the intercept) and the wage rate (which gives the slope). Once we have the indifference curves and the budget line, we can determine how much labor is supplied.

An increase in nonlabor income corresponds to a parallel shift of the budget line. We always assume that consumption and leisure are both normal goods, so that after a rise in nonlabor income, consumption increases and less labor is supplied.

A rise in an individual worker's wage rate has both an income effect and a substitution effect. The substitution effect, which is a movement to a steeper part of the original indifference curve, results in more labor supplied. The income effect, which is a movement to a higher indifference curve, results in less labor supplied. Either effect could dominate. When wages are low, however, income effects are small, so at least at low wages the substitution effect dominates. Thus, at low wages the individual's labor supply curve slopes upward, whereas at high wages it could either continue to slope upward or it could bend backward.

Combining the worker's labor supply curve with the firm's labor demand curve (and remembering that the labor demand curve is the MRP_L curve), we can find the market equilibrium and study how it changes in response to changing market conditions. A rise in nonlabor income leads to a leftward shift in labor supply. A rise in marginal productivity leads to a rightward shift in labor demand. If the rise in marginal productivity increases the nonlabor income of workers (by increasing the value of the capital that they own) then it leads to a leftward shift in labor supply as well.

When wage changes are perceived to be temporary, intertemporal substitution takes place. That is, the labor supply curve shifts to reflect workers' response to their perception that the situation is temporary. If wages are perceived to be temporarily high, workers will reschedule their current vacation plans for later; if wages are perceived to be temporarily low, workers will reschedule their future vacation plans for today. Thus, it is possible that even small wage changes, if perceived to be temporary, could yield very large changes in employment. This is consistent with what we know of the history of recessions.

Different workers receive different wages for different reasons. Often, a portion of the worker's paycheck is not really a wage at all, but a return on human capital. Workers can benefit by having access to capital of differing qualities, including their colleagues' human capital, from which they receive external benefits. Some workers receive positive compensating differentials for work that is especially pleasant or negative ones for work that has special advantages.

There are substantial wage differences between African-Americans and whites and between men and women. Many factors, including discrimination, might be

part of the explanation. Most of these factors, including differences in human capital, are very difficult to measure, making it hard to determine the significance of discrimination. Some wage differences result from the choices of workers themselves, as when married women choose to live in the cities where their husbands can earn the highest wage, rather than in the cities where they themselves can earn the highest wage. Economists do not know how important a role such phenomena play in determining wage differences.

Review Questions

- R1.** Explain the income and substitution effects of a rise in an individual's wage. Which causes him to work less, and why?
- R2.** Under what circumstances can we be sure that the substitution effect will outweigh the income effect? What implications does this have for the shape of the individual's labor supply curve?
- R3.** What are the possible shapes for an individual's labor supply curve? Interpret them in terms of income and substitution effects.
- R4.** Draw a diagram with two panels depicting the supply and demand for labor both in the market as a whole and for an individual worker.
- R5.** In the preceding question, how is the supply curve affected by a change in nonlabor income? How is the demand curve affected by a change in productivity?
- R6.** Will employment fall more in response to a permanent fall in wages or in response to a temporary fall? Why?
- R7.** List some reasons why different people earn different wages.
- R8.** List some theories that might explain wage differences between African-Americans and whites. How might you go about testing some of these theories? What problems might you run into?

Problem Set

- 1.** **True or False:** If an individual suddenly found that he needed less sleep per night than previously, his consumption would go up.
- 2.** Jack can work up to 8 hours a day at a wage rate of W and as much more as he wants at the higher overtime rate of W' . He chooses to work 10 hours. Jill can work as many hours as she wants at a wage of W'' . Jack and Jill have the same tastes, the same assets, and are equally happy. What can you conclude about the size of W'' compared with W and W' ? What can you conclude about the number of hours Jill works?
- 3.** Suppose that all people have identical tastes and identical talents, but that those who attend college become more productive and hence earn higher wages. On the other hand, college students have to pay tuition.

- a. Explain why college graduates and nongraduates must be equally happy. (*Hint:* What would happen to tuition if they weren't?) Use this observation and an indifference curve diagram to illustrate the equilibrium tuition cost.
 - b. **True or False:** Because college graduates earn higher wages, they might choose to work fewer hours than nongraduates.
4. Dick recently received a substantial inheritance from his aunt and immediately started working more hours at his job. If Dick's wage rate increases, can you predict what will happen to the number of hours that he works? Justify your answer.
5. Jane recently received a substantial inheritance from her aunt and immediately started working fewer hours at her job. If Jane's wage rate increases, can you predict what will happen to the number of hours that she works? Justify your answer.
6. Leisure is an inferior good for Horace.
 - a. Use indifference curves to show the income and substitution effects of an increase in Horace's wage rate.
 - b. Could Horace's labor supply curve be backward-bending? How do you know?
7. Hortense earns a wage of \$10 per hour and chooses to work 35 hours per week. One day, her employer tells her that while he will continue to pay her \$10 an hour for her first 35 hours each week, he will now pay her \$15 per hour for any additional hours beyond the first 35.
 - a. Illustrate Hortense's situation with indifference curves.
 - b. **True or False:** Hortense might choose to continue working exactly 35 hours per week.
8. Car wash attendants currently earn \$5 per hour and choose to work 50 hours per week. A law has just been passed requiring car washes to pay double wages for any hours in excess of 40 per week. The law does not, however, apply to any other occupations.
 - a. Explain why car wash attendants must remain on the same indifference curve. What must happen to their basic wage rate?
 - b. **True or False:** Car wash attendants will certainly now work more hours than they did previously.
9. **True or False:** A man who earns his entire income in wages will respond more sharply to a rise in the wage than will a man whose income is mostly from property.
10. **True or False:** Workers who like their jobs will be more productive at the margin than those who don't.
11. Suppose that an unexpected blight wipes out a large portion of this year's agricultural harvest. What happens to the wage rate, the amount of labor supplied to the marketplace, and the amount of labor supplied by any given individual?
12. Suppose that a tornado destroys a large number of major factories.
 - a. What is the effect on the demand for labor?
 - b. If the factories are owned by workers (say, through stock ownership), what is the effect on the supply of labor?
 - c. What is the effect on the wage rate, the amount of labor supplied to the marketplace, and the amount of labor supplied by any individual?

13. Suppose that an epidemic kills half the workers in an industry that produces goods for export. What is the effect on the wage rate, the amount of labor supplied to the marketplace, and the amount of labor supplied by any individual surviving worker?
14. In the preceding problem, suppose that instead of being produced for export, the good being manufactured is sold to the very workers who produce it. How does your answer change?
15. **True or False:** If the capital stock is fixed and if the level of output is fixed, then a rise in the marginal productivity of labor benefits the owners of capital.
16. How would the wage rate and the level of employment be affected by the invention of a costless pill that made it unnecessary for anyone to sleep?
17. Contrast the effects on employment, output, and wages of (a) a year of bad weather resulting in low agricultural productivity and (b) nuclear contamination that lowers agricultural productivity permanently.
18. Contrast the effects on employment, output, and wages of (a) an income tax that is expected to be in effect for 1 year and (b) an income tax that is expected to be permanent.
19. The current federal tax law allows deductions for the depreciation of physical capital. **True or False:** One effect of this deduction is to reduce the average level of education.

