Labor Economics

Fourth Edition

George J. Borjas
Harvard University

McGraw-Hill Irwin

Boston Burr Ridge, IL Dubuque, IA New York San Francisco St. Louis
Bangkok Bogotá Caracas Kuala Lumpur Lisbon London Madrid Mexico City
Milan Montreal New Delhi Santiago Seoul Singapore Sydney Taipei Toronto

Chapter 7

Human Capital

If you think education's expensive, try ignorance!

—Derek Bok

The theory of compensating differentials suggests that wages will vary among workers because jobs are different. Wages also will vary because workers are different. We each bring into the labor market a unique set of abilities and acquired skills, or **human capital**. For instance, some persons train to be research biologists while other persons train to be musicians. This chapter discusses how we choose the particular set of skills that we offer to employers and how our choices affect the evolution of earnings over the working life.

We acquire most of our human capital in school and in formal and informal on-the-job training programs. The skills we acquire in school make up an increasingly important component of our stock of knowledge. In 1940, 75.5 percent of adults in the United States had not graduated from high school, and only 4.6 percent had a college degree. By 2005, only 13.5 percent of adults did not have a high school diploma, and 27.7 percent had at least a college degree.

This chapter analyzes why some workers obtain a lot of schooling and other workers drop out at an early age. Workers who invest in schooling are willing to give up earnings today in return for higher earnings in the future. For example, we earn a relatively low wage while we attend college or participate in a formal apprenticeship program. However, we expect to be rewarded by higher earnings later on as we collect the returns on our investment. The trade-off between lower earnings today and higher earnings later, as well as the financial and institutional constraints that limit access to education, determines the distribution of educational attainment in the population.

We also will discuss whether the money spent on education is a good investment. In particular, how does the rate of return to schooling compare with the rate of return on other investments? Putting aside our own personal interest in knowing whether we are getting a good deal out of our college education, the rate of return to schooling plays an important role in many policy discussions. It is often argued, for instance, that subsidizing investments to education and other learning activities is the surest way of improving the economic well-being of low-income and disadvantaged workers.

We do not typically stop accumulating skills and knowledge the day we finally leave school. Instead, we continue to add to our human capital stock throughout much of our working lives. As a result of the human capital acquired through training and vocational programs, college graduates in their fifties earn twice as much as college graduates in their twenties. This chapter also analyzes how workers choose a particular path for their postschool investments and investigates how these choices influence the evolution of earnings over the life cycle and determine the earnings distribution in the economy.

Our analysis will assume that the worker chooses the level of human capital investments that maximizes the present value of lifetime earnings. This approach to the study of the determinants of the earnings distribution differs fundamentally from alternative approaches that view a worker's wage as determined by luck and other random factors. These random events might include whether the worker happened to meet an aging billionaire on the way to work or whether the worker was having breakfast at a Hollywood diner when an influential agent walked in. The human capital approach does not deny that luck, looks, and being in the right place at the right time influence a worker's earnings. Rather, we stress the idea that our educational and training decisions play an important role in the determination of earnings.

7-1 Education in the Labor Market: Some Stylized Facts

Table 7-1 summarizes the distribution of education in the U.S. population. The table shows clearly that there are only slight differences in educational attainment between men and women, but that there are substantial differences among racial and ethnic groups. By 2005, only about 9 percent of white persons did not have a high school diploma, as opposed to 16 percent of African Americans and 39 percent of Hispanics. In contrast, more than 30 percent of whites had at least a college diploma, as compared to only 18 percent of African Americans and 12 percent of Hispanics.[1]

TABLE 7-1 Educational Attainment of U.S. Population, 2005 (Persons Aged 25 and Over)

Source: U.S. Bureau of Labor Statistics, *Annual Demographic Supplement of the Current Population Surveys*, March 2005.

	Highest Grade Completed (Percentage of Population in Education Category)					
Group	**Less than High School**	**High School Graduates**	**Some College**	**Associate Degree**	**Bachelor's Degree**	**Advanced Degree**
All Persons	13.5%	33.4%	16.8%	8.7%	18.1%	9.6%
Gender:						
Male	13.8	33.4	16.4	7.6	18.3	10.6
Female	13.3	33.5	17.2	9.5	17.9	8.6
Race/ethnicity:						
White	9.1	33.8	17.4	9.3	19.7	10.8
Black	16.4	39.5	18.5	8.0	12.5	5.2
Hispanic	38.8	30.4	13.3	5.6	8.5	3.5

[1] A careful study of the differences in educational attainment among race/ethnic groups is given by Stephen V. Cameron and James J. Heckman, "The Dynamics of Educational Attainment for Black, Hispanic, and White Males," *Journal of Political Economy* 109 (June 2001): 455–99.

TABLE 7-2

Sources: U.S. Burea

All Workers:

Gender:
Men

Women

Race/ethnicity:
White

Black

Hispanic

TABLE 7-2 **Labor Market Characteristics, by Education Group, 2005 (Persons Aged 25–64)**

Sources: U.S. Bureau of Labor Statistics, *Annual Demographic Supplement of the Current Population Surveys*, March 2005.

		Less than High School	High School Graduates	Some College	College Graduates
All Workers:	Labor force participation rate	62.0	75.5	81.2	85.8
	Unemployment rate	9.1	5.6	4.1	2.3
	Annual earnings (in $1,000)	21.1	31.0	37.4	63.2
Gender:					
Men	Labor force participation rate	74.8	83.8	87.6	91.8
	Unemployment rate	7.9	6.1	4.2	2.5
	Annual earnings (in $1,000)	24.6	37.0	46.2	79.0
Women	Labor force participation rate	48.0	67.0	75.8	79.9
	Unemployment rate	11.1	4.9	4.0	2.2
	Annual earnings (in $1,000)	15.5	23.8	29.2	46.1
Race/ethnicity:					
White	Labor force participation rate	56.2	75.9	81.6	86.3
	Unemployment rate	9.2	4.8	3.6	2.2
	Annual earnings (in $1,000)	23.6	32.9	39.1	65.4
Black	Labor force participation rate	55.5	72.6	79.9	87.2
	Unemployment rate	17.0	9.4	7.4	3.5
	Annual earnings (in $1,000)	18.8	25.2	31.1	51.5
Hispanic	Labor force participation rate	69.4	77.7	81.9	84.7
	Unemployment rate	7.4	5.2	3.6	2.4
	Annual earnings (in $1,000)	20.0	27.6	33.8	52.8

The differences in educational attainment among workers are significant because, as Table 7-2 shows, education is strongly correlated with labor force participation rates, unemployment rates, and earnings. The labor force participation rate of persons who lack a high school diploma is only 62 percent, as compared to 86 percent for college graduates. Even though the unemployment rate of high school dropouts is 9 percent, it is only 2 percent for college graduates. Finally, high school dropouts earn just over $21,000 annually, but college graduates earn $63,000, three times as much.

The data also indicate that education has a substantial beneficial impact on the labor market experiences of women and minorities. For example, the unemployment rate of black high school dropouts is 17 percent, as compared to 9.4 percent for black high school graduates and 3.5 percent for black college graduates. Similarly, Hispanic high school dropouts earn only $20,000, as compared to nearly $53,000 for Hispanic college graduates.

Although there are sizable differences in labor market outcomes between men and women and across race and ethnic groups—and these differences will be discussed in detail in Chapter 10—this chapter investigates a different lesson that can be drawn from the data reported in Table 7-2: Education plays a crucial role in improving labor market outcomes for both men and women and for workers in all racial and ethnic groups.

7-2 Present Value

Any study of an investment decision—whether it is an investment in physical or in human capital—must contrast expenditures and receipts incurred at different time periods. In other words, an investor must be able to calculate the returns to the investment by comparing the current cost with the future returns. For reasons that will become obvious momentarily, however, the value of a dollar received today is not the same as the value of a dollar received tomorrow. The notion of **present value** allows us to compare dollar amounts spent and received in different time periods.

Suppose somebody gives you a choice between two monetary offers: You can have either \$100 today or \$100 next year. Which offer would you take? A little reflection should convince you that \$100 today is better than \$100 next year. After all, if you receive \$100 today, you can invest it, and you will then have \$100 × (1 + 0.05) dollars next year (or \$105), assuming that the rate of interest equals 5 percent. Note, moreover, that receiving \$95.24 today (or \$100 ÷ 1.05) would be worth \$100 next year. Hence, the present value (or the current dollar value) of receiving \$100 tomorrow is only \$95.24. In general, the present value of a payment of, say, y dollars next year is given by

$$PV = \frac{y}{1 + r} \qquad \textbf{(7-1)}$$

where r is the rate of interest, which is also called the **rate of discount.** The quantity PV tells us how much needs to be invested today in order to have y dollars next year. In effect, a future payment of y dollars is discounted so as to make it comparable to current dollars.

The discussion clearly suggests that receiving y dollars two years from now is not equivalent to receiving y dollars today or even to receiving y dollars next year. A payment of \$100 today would be worth \$100 × (1 + 0.05) × (1 + 0.05) two years from now. Hence, the present value of receiving y dollars two years from now is

$$PV = \frac{y}{(1 + r)^2} \qquad \textbf{(7-2)}$$

By arguing along similar lines, we can conclude that the present value of y dollars received t years from now equals

$$PV = \frac{y}{(1 + r)^t} \qquad \textbf{(7-3)}$$

These formulas are extremely useful when we study decisions that involve expenditures made or dollars received at different time periods because they allow us to state the value of these expenditures and receipts in terms of today's dollars.

7-3 The Schooling Model

Education is associated with lower unemployment rates and higher earnings. So why don't all workers get doctorates or professional degrees? In other words, what factors motivate some workers to get professional degrees while other workers drop out before they finish high school?

We begin our analysis of this important question by assuming that workers acquire the education level that maximizes the present value of lifetime earnings. Education and other forms of training, therefore, are valued only because they increase earnings. A college education obviously affects a person's utility in many other ways. It teaches the student how to read and appreciate Nietzsche and how to work out complex mathematical equations; it even reduces the cost of entering the "marriage market" by facilitating contact with a large number of potential mates. Important though they may be, we will ignore these side effects of human capital investments and concentrate exclusively on the monetary rewards of an education.[2]

Consider the situation faced by an 18-year-old man who has just received his high school diploma and who is contemplating whether to enter the labor market or attend college and delay labor market entry by an additional four years.[3] Suppose that there is no on-the-job training and that the skills learned in school do not depreciate over time. These assumptions imply that the worker's productivity does not change once he leaves school, so that real earnings (that is, earnings after adjusting for inflation) are constant over the life cycle.

Figure 7-1 illustrates the economic trade-off involved in the worker's decision. The figure shows the age-earnings profile (that is, the wage path over the life cycle) associated with each alternative. Upon entering the labor market, high school graduates earn w_{HS} dollars annually until retirement age, which occurs when the worker turns 65. If the person chooses to attend college, he gives up w_{HS} dollars in labor earnings and incurs "direct" costs of H dollars to cover tuition, books, and fees. After graduation, he earns w_{COL} dollars annually until retirement.

Figure 7-1 indicates that going to college involves two different types of costs. A year spent in college is a year spent out of the labor force (or at least a year spent working in a low-wage part-time job), so that a college education forces the worker to forgo some earnings. This is the opportunity cost of going to school—the cost of not pursuing the best alternative. The opportunity cost is w_{HS} dollars for each year the student goes to college. The student also has out-of-pocket expenses of H dollars for tuition, books, and a variety of other fees.

Because college has no intrinsic value to the student, employers who wish to attract a highly educated (and presumably more productive) worker will have to offer higher wages, so that $w_{COL} > w_{HS}$. In a sense, the high wage paid to workers with more schooling is a compensating differential that compensates workers for their training costs. If college graduates earned less than high school graduates, no one would bother to get a college education because we are assuming that workers do not get any other benefits from attending college.

[2] A more general approach would assume that workers choose to acquire the skill level that maximizes lifetime utility. Most of the key insights of the schooling model, however, are not affected by this generalization; see Robert T. Michael, "Education in Nonmarket Production," *Journal of Political Economy* 81 (March/April 1973): 306–27.

[3] The schooling model was first analyzed by Jacob Mincer, "Investment in Human Capital and Personal Income Distribution," *Journal of Political Economy* 66 (August 1958): 281–302. An even earlier study that anticipated many of the central concepts in the human capital literature is given by Milton Friedman and Simon Kuznets, *Income from Independent Professional Practice,* Princeton, NJ: Princeton University Press, 1954.

We begin our analysis of this important question by assuming that workers acquire the education level that maximizes the present value of lifetime earnings. Education and other forms of training, therefore, are valued only because they increase earnings. A college education obviously affects a person's utility in many other ways. It teaches the student how to read and appreciate Nietzsche and how to work out complex mathematical equations; it even reduces the cost of entering the "marriage market" by facilitating contact with a large number of potential mates. Important though they may be, we will ignore these side effects of human capital investments and concentrate exclusively on the monetary rewards of an education.[2]

Consider the situation faced by an 18-year-old man who has just received his high school diploma and who is contemplating whether to enter the labor market or attend college and delay labor market entry by an additional four years.[3] Suppose that there is no on-the-job training and that the skills learned in school do not depreciate over time. These assumptions imply that the worker's productivity does not change once he leaves school, so that real earnings (that is, earnings after adjusting for inflation) are constant over the life cycle.

Figure 7-1 illustrates the economic trade-off involved in the worker's decision. The figure shows the **age-earnings profile** (that is, the wage path over the life cycle) associated with each alternative. Upon entering the labor market, high school graduates earn w_{HS} dollars annually until retirement age, which occurs when the worker turns 65. If the person chooses to attend college, he gives up w_{HS} dollars in labor earnings and incurs "direct" costs of H dollars to cover tuition, books, and fees. After graduation, he earns w_{COL} dollars annually until retirement.

Figure 7-1 indicates that going to college involves two different types of costs. A year spent in college is a year spent out of the labor force (or at least a year spent working in a low-wage part-time job), so that a college education forces the worker to forgo some earnings. This is the **opportunity cost** of going to school—the cost of not pursuing the best alternative. The opportunity cost is w_{HS} dollars for each year the student goes to college. The student also has out-of-pocket expenses of H dollars for tuition, books, and a variety of other fees.

Because college has no intrinsic value to the student, employers who wish to attract a highly educated (and presumably more productive) worker will have to offer higher wages, so that $w_{COL} > w_{HS}$. In a sense, the high wage paid to workers with more schooling is a compensating differential that compensates workers for their training costs. If college graduates earned less than high school graduates, no one would bother to get a college education because we are assuming that workers do not get any other benefits from attending college.

[2] A more general approach would assume that workers choose to acquire the skill level that maximizes lifetime utility. Most of the key insights of the schooling model, however, are not affected by this generalization; see Robert T. Michael, "Education in Nonmarket Production," *Journal of Political Economy* 81 (March/April 1973): 306–27.

[3] The schooling model was first analyzed by Jacob Mincer, "Investment in Human Capital and Personal Income Distribution," *Journal of Political Economy* 66 (August 1958): 281–302. An even earlier study that anticipated many of the central concepts in the human capital literature is given by Milton Friedman and Simon Kuznets, *Income from Independent Professional Practice,* Princeton, NJ: Princeton University Press, 1954.

FIGURE 7-1 Potential Earnings Streams Faced by a High School Graduate
A person who quits school after getting his high school diploma can earn w_{HS} from age 18 until the age of retirement. If he decides to go to college, he forgoes these earnings and incurs a cost of H dollars for four years and then earns w_{COL} until retirement age.

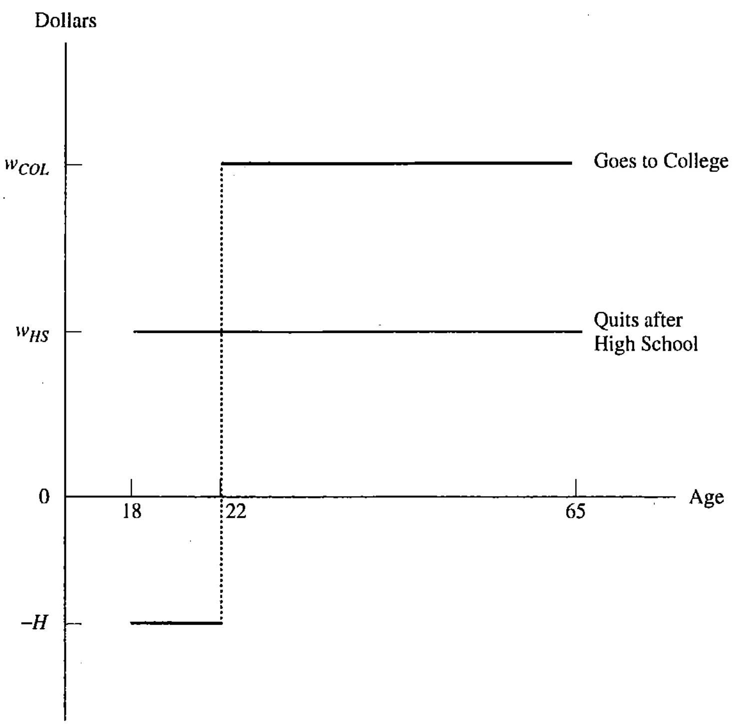

Present Value of Age-Earnings Profiles

The present value of the earnings stream if the worker gets only a high school education is

$$PV_{HS} = w_{HS} + \frac{w_{HS}}{(1 + r)} + \frac{w_{HS}}{(1 + r)^2} + \cdots + \frac{w_{HS}}{(1 + r)^{46}} \qquad \text{(7-4)}$$

where r gives the worker's rate of discount. There are 47 terms in this sum, one term for each year that elapses between the ages of 18 and 64.

The present value of the earnings stream if the worker gets a college diploma is

$$PV_{COL} = \underbrace{-H - \frac{H}{(1 + r)} - \frac{H}{(1 + r)^2} - \frac{H}{(1 + r)^3}}_{\text{Direct Costs of Attending College}} + \underbrace{\frac{w_{COL}}{(1 + r)^4} + \frac{w_{COL}}{(1 + r)^5} + \cdots + \frac{w_{COL}}{(1 + r)^{46}}}_{\text{Post-college Earnings Stream}} \qquad \text{(7-5)}$$

The first four terms in this sum give the present value of the direct costs of a college education, whereas the remaining 43 terms give the present value of lifetime earnings in the postcollege period.

A person's schooling decision maximizes the present value of lifetime earnings. Therefore, the worker attends college if the present value of lifetime earnings when he gets a college education exceeds the present value of lifetime earnings when he gets only a high school diploma, or

$$PV_{COL} > PV_{HS} \tag{7-6}$$

Let's illustrate the worker's decision with a simple numerical example. Suppose a worker lives only two periods and chooses from two schooling options. He can choose not to attend school at all, in which case he would earn \$20,000 in each period. The present value of earnings is

$$PV_0 = 20{,}000 + \frac{20{,}000}{1 + r} \tag{7-7}$$

He also can choose to attend school in the first period, incur \$5,000 worth of direct schooling costs, and enter the labor market in the second period, earning \$47,500. The present value of this earnings stream is

$$PV_1 = -5{,}000 + \frac{47{,}500}{1 + r} \tag{7-8}$$

Suppose that the rate of discount is 5 percent. It is easy to calculate that $PV_0 = \$39{,}048$ and that $PV_1 = \$40{,}238$. The worker, therefore, chooses to attend school. Note, however, that if the rate of discount were 15 percent, $PV_0 = \$37{,}391$, $PV_1 = \$36{,}304$, and the worker would not go to school.

As this example shows, the rate of discount *r* plays a crucial role in determining whether a person goes to school or not. The worker goes to school if the rate of discount is 5 percent but does not if the rate of discount is 15 percent. The higher the rate of discount, therefore, the less likely a worker will invest in education. This conclusion should be easy to understand. A worker who has a high discount rate attaches a very low value to future earnings opportunities—in other words, he discounts the receipt of future income "too much." Because the returns to an investment in education are collected in the far-off future, persons with high discount rates acquire less schooling.

It is sometimes assumed that the person's rate of discount equals the market rate of interest, the rate at which funds deposited in financial institutions grow over time. After all, the discounting of future earnings in the present value calculations arises partly because a dollar received this year can be invested and is worth more than a dollar received next year.

The rate of discount, however, also depends on how we feel about giving up some of today's consumption in return for future rewards—or our "time preference." Casual observation (and a large number of psychological experiments) suggests that people differ in how they approach this trade-off. Some of us are "present oriented" and some of us are not. Persons who are present oriented have a high discount rate and are less likely to invest in schooling. Although there is some evidence suggesting that poorer families have a higher rate of discount than wealthier families, we know little about how a person's time preference is formed.[4]

[4] Emily C. Lawrance, "Poverty and the Rate of Time Preference: Evidence from Panel Data," *Journal of Political Economy* 99 (February 1991): 54–77; see also Lawrence L. LeShan, "Time Orientation and Social Class," *Journal of Abnormal and Social Psychology* 47 (July 1952): 589–92. A theoretical analysis of the determinants of the rate of time preference is given by Gary S. Becker and Casey B. Mulligan, "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics* 112 (August 1997): 729–58.

The Wage-Schooling Locus

The rule that a person should choose the level of schooling that maximizes the present value of earnings obviously generalizes to situations when there are more than two schooling options. The person would then calculate the present value associated with each schooling option (for example, one year of schooling, two years of schooling, and so on) and choose the amount of schooling that maximizes the present value of the earnings stream.

There is, however, a different way of formulating this problem that provides an intuitive "stopping rule."[5] This stopping rule tells the individual when it is optimal to quit school and enter the labor market. This alternative approach is useful because it also suggests a way for estimating the rate of return to education.

Figure 7-2 illustrates the **wage-schooling locus,** which gives the salary that employers are willing to pay *a particular worker* for every level of schooling. If the worker gets a

FIGURE 7-2 The Wage-Schooling Locus

The wage-schooling locus gives the salary that a particular worker would earn if he completed a particular level of schooling. If the worker graduates from high school, he earns $20,000 annually. If he goes to college for one year, he earns $23,000.

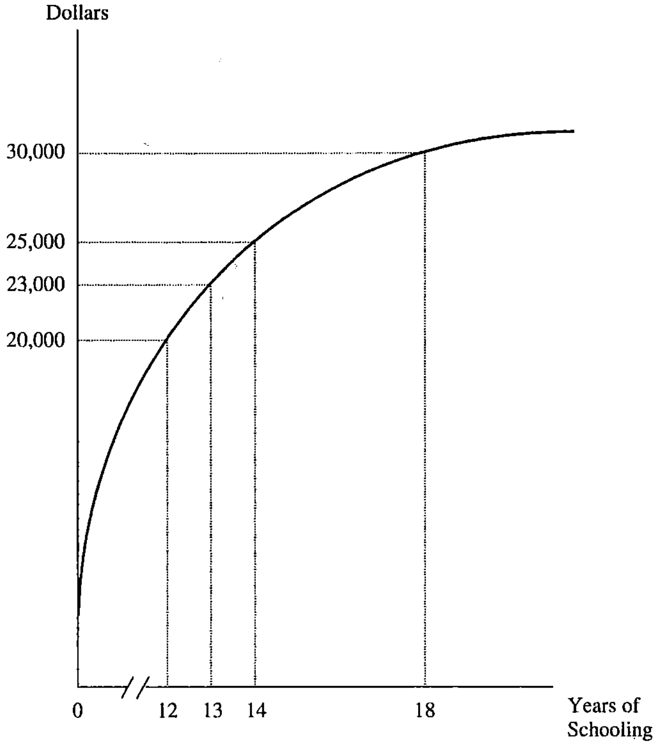

[5] Sherwin Rosen, "Human Capital: A Survey of Empirical Research," *Research in Labor Economics* 1 (1977): 3–39; and David Card, "The Causal Effect of Education on Earnings," in Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, vol. 3A. Amsterdam: Elsevier, 1999, pp. 1801–63.

high school diploma, his annual salary is $20,000; whereas if he gets 18 years of schooling, his annual salary rises to $30,000. The wage-schooling locus is market determined. In other words, the salary for each level of schooling is determined by the intersection of the supply of workers with that particular schooling and the demand for those workers. From the worker's point of view, the salary associated with each level of schooling is a constant.

The wage-schooling locus shown in Figure 7-2 has three important properties:

1. *The wage-schooling locus is upward sloping.* Workers who have more education must earn more as long as educational decisions are motivated only by financial gains. To attract educated workers, employers must compensate those workers for the costs incurred in acquiring an education.
2. *The slope of the wage-schooling locus tells us by how much a worker's earnings would increase if he were to obtain one more year of schooling.* The slope of the wage-schooling locus, therefore, will be closely related to any empirical measure of "the rate of return" to school.
3. *The wage-schooling locus is concave.* The monetary gains from each additional year of schooling decline as more schooling is acquired. In other words, the law of diminishing returns also applies to human capital accumulation.[6] Each extra year of schooling generates less incremental knowledge and lower additional earnings than previous years.

The Marginal Rate of Return to Schooling

The slope of the wage-schooling locus (or $\Delta w/\Delta s$) tells us by how much earnings increase if the person stays in school one more year. In Figure 7-2, for example, the first year of college increases annual earnings in the postschool period by $3,000. The percentage change in earnings from getting this additional year of schooling is 15 percent, (or $^{3,000}/_{20,000} \times 100$). In other words, the worker gets a 15 percent wage increase from staying in school and attending that first year of college. We refer to this *percentage* change in earnings resulting from one more year of school as the marginal **rate of return to schooling.**

The marginal rate of return to schooling gives the percentage increase in earnings per dollar spent in educational investments. To see why, suppose that the only costs incurred in going to college are forgone earnings. The high school graduate who delays his entry into the labor market by one year is then giving up $20,000. This investment outlay increases his future earnings by $3,000 annually, thus yielding an annual 15 percent rate of return for the first year of college.

Because the wage-schooling locus is concave, the marginal rate of return to schooling *must* decline as a person gets more schooling. For example, the marginal rate of return to the second year of college is only 8.7 percent (a $2,000 return on a $23,000 investment). Each additional year of schooling generates a smaller salary increase and it costs more to stay in school. In other words, the wage increase generated by each additional year of college gets smaller at the same time that the cost of staying in school keeps rising. The marginal rate of return schedule, therefore, is a declining function of the level of schooling, as

[6] See George Psacharapoulos, "Returns to Education: A Further International Update and Implications," *Journal of Human Resources* 20 (Fall 1985): 583–604, for evidence supporting the hypothesis that educational production functions exhibit diminishing marginal productivity.

FIGURE 7-3 The Schooling Decision
The *MRR* schedule gives the marginal rate of return to schooling, or the percentage increase in earnings resulting from an additional year of school. A worker maximizes the present value of lifetime earnings by going to school until the marginal rate of return to schooling equals the rate of discount. A worker with discount rate r goes to school for s^* years.

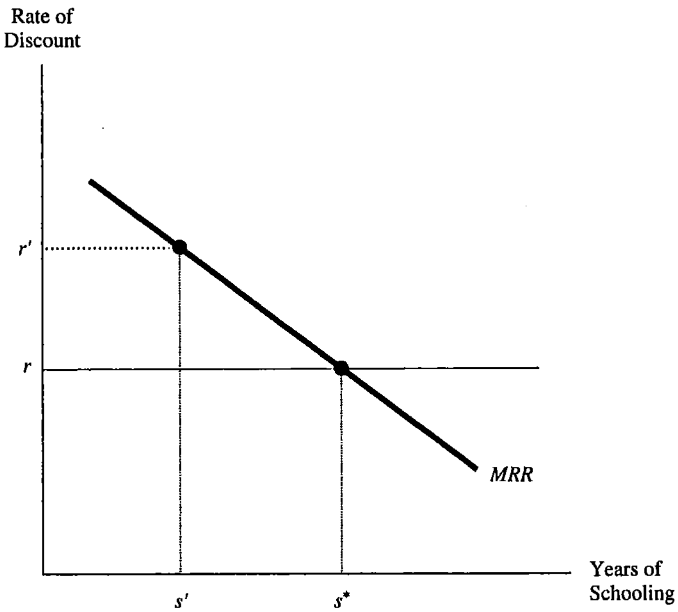

illustrated by the curve labeled *MRR* in Figure 7-3. To summarize, the *MRR* schedule gives the percentage change in annual earnings resulting from each additional year of school.

The Stopping Rule, or When Should I Quit School?

Suppose that the worker has a rate of discount r that is constant; that is, it is independent of how much schooling he gets. The rate of discount schedule, therefore, is perfectly elastic, as illustrated in Figure 7-3.

Which level of schooling should a person choose? It turns out that the intersection of the *MRR* curve and the horizontal rate of discount schedule determines the optimal level of schooling for the worker, or s^* years in the figure. In other words, the stopping rule that tells the worker when he should quit school is

$$\text{Stop schooling when the marginal rate of return to schooling} = r \qquad (7\text{-}9)$$

This stopping rule maximizes the worker's present value of earnings over the life cycle. To see why, suppose that the worker's rate of discount equals the market rate of interest offered by financial institutions. Would it be optimal for the worker to quit school after completing only s' years in Figure 7-3? If the worker were to stay in school for an additional year, he would forgo, say, w' dollars in earnings, and the rate of return to this investment equals r'. His alternative would be to quit school, work, and deposit the w' dollars in a bank that offers a rate of return of only r. Because education yields a higher rate of return, the worker increases the present value of earnings by continuing in school.

Conversely, suppose that the worker gets more than s^* years of school. Figure 7-3 then shows that the marginal rate of return to this "excess" schooling is less than the market rate of interest, so that the extra years of schooling are not profitable.

Equation (7-9)—the stopping rule for schooling investments—describes a general property of optimal investment decisions. The wealth-maximizing student who must decide if he should quit school faces the same economic trade-off as the owner of a forest who must decide when to cut down a tree. The longer the tree is in the ground, the larger it gets and the more lumber and revenue it will eventually generate. But there are forgone profits (as well as maintenance costs) associated with keeping the tree in the ground. The tree should be cut down when the rate of return on investing in the tree equals the rate of return on alternative investments.

It is important to emphasize that the decision of whether to stay in school is influenced by many factors (such as chance encounters with influential teachers or "significant others"), not just the dollar value of the earnings stream. There is also a great deal of uncertainty in the rewards to particular types of education. The assumption that the student knows the shape of the wage-schooling locus—and the marginal rate of return provided by each level of schooling—is clearly false. Economic and social conditions change in unpredictable ways, and it is very difficult to forecast how these shocks shift the rewards to particular types of skills and careers. This uncertainty will surely play a role in our human capital decisions—just like the uncertainty in financial markets affects the type of financial portfolio that maximizes our wealth.[7]

7-4 Education and Earnings

The schooling model summarized by Figure 7-3 tells us how a particular worker decides how much schooling to acquire, and, as a result, also tells us where a worker places in the income distribution in the postschool period. Workers who get more schooling earn more (although they also give up more). The model isolates two key factors that lead different workers to obtain different levels of schooling and, hence, to have different earnings: Workers either have different rates of discount or face different marginal rate of return schedules.

Differences in the Rate of Discount

Consider a labor market with two workers who differ *only* in their discount rates, as illustrated in Figure 7-4*a*. Al's discount rate is r_{AL} and Bo's lower discount rate is r_{BO}. The figure shows that Al (who has a higher discount rate) drops out of high school and gets only 11 years of education; Bo gets a high school diploma. As we saw earlier, workers who discount future earnings heavily do not go to school because they are present oriented.

Figure 7-4*b* shows the implications of these choices for the observed earnings distribution in the postschool period. We have assumed that both workers face the same marginal rate of return schedule. Given our derivation of the marginal rate of return schedule, this assumption is equivalent to saying that both workers face the same wage-schooling locus.

[7] Joseph G. Altonji, "The Demand for and Return to Education When Outcomes Are Uncertain," *Journal of Labor Economics* 11 (January 1993): 48–83.

FIGURE 7-4 Schooling and Earnings When Workers Have Different Rates of Discount
Al has a higher rate of discount (r_{AL}) than Bo (r_{BO}), so that Bo graduates from high school but Al drops out. Al chooses point P_{AL} on the wage-schooling locus and Bo chooses point P_{BO}. The observed data on wages and schooling in the labor market trace out the common wage-schooling locus of the workers.

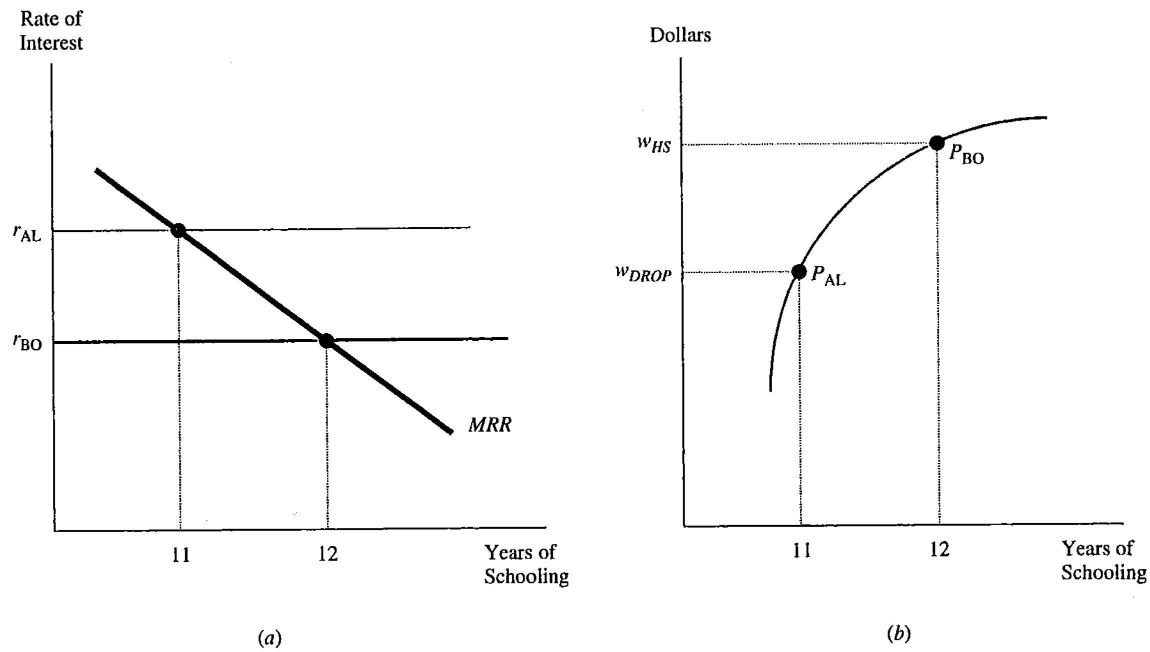

The different schooling decisions of the two workers, therefore, simply place them at different points of the common locus. Al ends up at point P_{AL}, where he goes to school 11 years and earns w_{DROP} dollars; Bo ends up at point P_{BO}, goes to school 12 years, and earns w_{HS} dollars. Note that by connecting points P_{AL} and P_{BO} we can trace out the common wage-schooling locus faced by all workers. Moreover, note also that the wage gap between Al and Bo lets us estimate the rate of return to the 12th grade, the percentage change in earnings that a worker would experience in going from the 11th to the 12th grade.

Estimates of the returns to schooling play a crucial role in many discussions of public policy. Consider, for example, the impact of a proposed law requiring all students to complete their high school education. By how much would this proposed policy increase the earnings of workers who are now high school dropouts?

In effect, this policy "injects" Al with one more year of schooling. The wage-schooling locus in Figure 7-4 shows that a high school graduate earns w_{HS} dollars. In other words, Al's earnings would increase to w_{HS} if the law went into effect. A compulsory high school diploma would move the worker along the *observed* wage-schooling locus.

As long as workers differ only in their discount rates, therefore, we can calculate the marginal rate of return to schooling from the wage differential between two workers who differ in their educational attainment. We can then correctly predict by how much earnings would increase if we pursued particular policies that injected targeted workers with more education.[8]

Differences in Ability

It is much more difficult to estimate the rate of return to schooling when all workers have the same rate of discount, but each worker faces a different wage-schooling locus—which, in turn, implies that each worker has a different marginal rate of return schedule. It is often assumed that higher ability levels shift the marginal rate of return schedule to the right, so that the earnings gain resulting from an additional year of schooling outweighs the increase in forgone earnings. In other words, more able persons get relatively more from an extra year of schooling. As illustrated in Figure 7-5*a*, Bob's *MRR* schedule lies to the right of Ace's. Because both Bob and Ace have the same rate of discount and because Bob gets more from an additional year of schooling, Bob gets more schooling (12 years versus 11 years).

Figure 7-5*b* illustrates the impact of this ability differential. Bob chooses point P_{BOB} on *his* wage-schooling locus; Bob gets 12 years of schooling and earns w_{HS} dollars. Ace chooses point P_{ACE} on *his* wage-schooling locus; Ace goes to school 11 years and earns w_{DROP} dollars. Note that Bob's wage-schooling locus lies above Ace's because Bob is more able.

The data at our disposal include the education and earnings of the two workers but *do not include* their ability levels. Innate ability, after all, is seldom observed. The observed data, therefore, connect the points P_{ACE} and P_{BOB} in the figure and trace out the line labeled *Z*. It is important to note that this line does *not* coincide with either Ace's or Bob's wage-schooling locus. As a result, the observed data on earnings and schooling do not allow us to estimate the rate of return to schooling.

Suppose that the government proposes a law requiring all persons to complete high school. To determine the economic impact of the proposed legislation, we wish to know by how much Ace's earnings would increase if he were injected with one more year of schooling. The available data tell us that a high school graduate earns w_{HS} and that a high school dropout earns w_{DROP}. Note, however, that the wage differential between Bob and Ace does *not* give the wage gain that Ace would get under the proposed legislation. Line *Z* in Figure 7-5*b* connects points on different wage-schooling curves and provides no information whatsoever about the wage increase that a particular worker would get if he or she were to obtain additional schooling. If the law goes into effect, Ace's earnings would increase only

[8] Many studies examine how credit constraints, student aid, and other financial resources affect the education decision. The relaxation of financial constraints can be interpreted as a decrease in the rate of discount; the additional wealth may make students less present oriented or may allow students to borrow money (to finance their education) at a lower interest rate. The evidence often indicates that the relaxation of financial constraints typically leads to more schooling. See Thomas J. Kane, "College Entry by Blacks since 1970: The Role of College Costs, Family Background, and the Returns to Education," *Journal of Political Economy* 102 (October 1994): 878–911; Pedro Carneiro and James J. Heckman, "The Evidence on Credit Constraints in Post-Secondary Schooling," *Economic Journal* 112 (October 2002): 705–34; and Susan M. Dynarski, "Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion," *American Economic Review* 93 (March 2003): 279–88. Contrary evidence is given by Stephen V. Cameron and Christopher Taber, "Estimation of Educational Borrowing Constraints Using Returns to Schooling," *Journal of Political Economy* 112 (February 2004): 132–82.

FIGURE 7-5 Schooling and Earnings When Workers Have Different Abilities
Ace and Bob have the same discount rate (r), but each worker faces a different wage-schooling locus. Ace drops out of high school and Bob gets a high school diploma. The wage differential between Bob and Ace (or $w_{HS} - w_{DROP}$) arises both because Bob goes to school for one more year and because Bob is more able. As a result, this wage differential does not tell us by how much Ace's earnings would increase if he were to complete high school (or $w_{ACE} - w_{DROP}$).

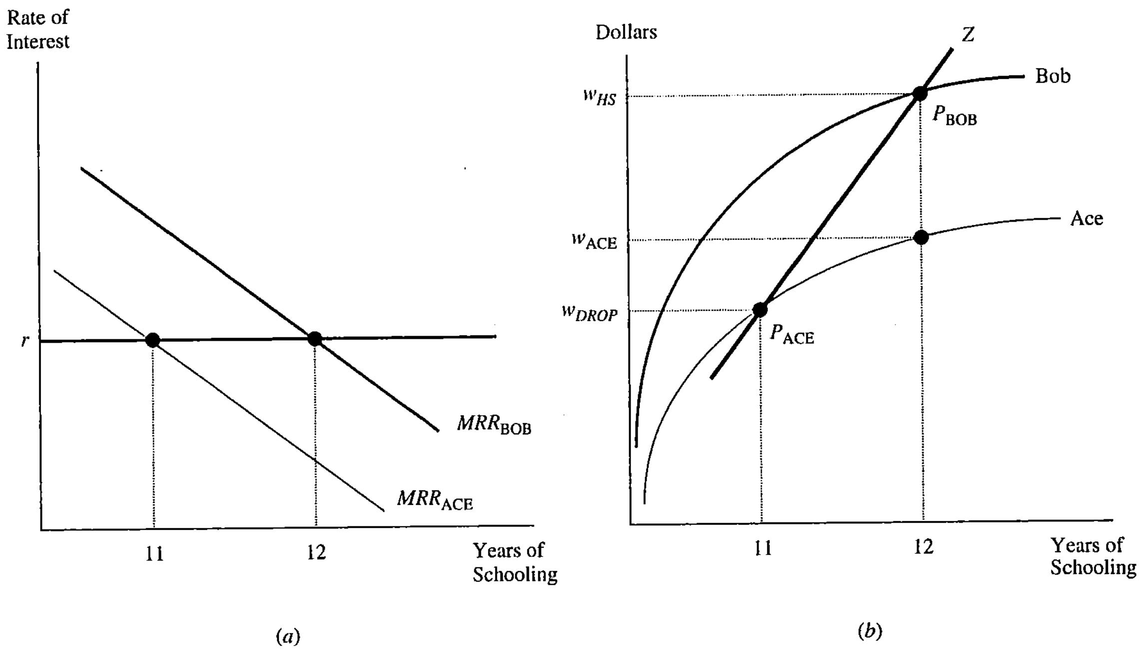

from w_{DROP} to w_{ACE}, which is much less than what a high school graduate like Bob now earns (w_{HS}).

Put differently, the wage gap between Ace and Bob arises for two reasons. Bob has more schooling than Ace and, hence, is getting the returns to additional schooling. Bob, however, also earns more than Ace because Bob is more able (and his wage locus lies above Ace's). The wage differential between these two workers, therefore, incorporates the impact of both education and ability on earnings.

Ability Bias

The model provides an important lesson: If there are unobserved ability differences in the population, earnings differentials across workers do not estimate the returns to schooling. The correlation between schooling and earnings across workers is contaminated by ability differentials, and hence does not provide an answer to the question that initially motivated our analysis: By how much would the earnings of a particular worker increase if he were to obtain more schooling?

Why should one care about this type of ability bias? Suppose that a well-meaning government bureaucrat observes that high school graduates earn $15,000 more per year than high school dropouts. He uses these data to convince policymakers that funding programs that encourage students to complete high school would increase the average wage of high school dropouts by $15,000. In the bureaucrat's calculations, this earnings gain implies that the program "funds itself" (presumably from higher tax revenues, lower expenditures on social assistance programs, and so forth).

We now know that the bureaucrat's argument is fatally flawed. He is assuming that high school graduates and high school dropouts have the same wage-schooling locus and that one can "fix" the earnings disadvantage of dropouts by injecting them with more schooling. It might be the case, however, that high school graduates have a higher wage-schooling locus. Encouraging high school dropouts to complete their high school education would not lead to a $15,000 increase in their earnings upon graduation, and it might be much more difficult to argue that the program pays its way.

7-5 Estimating the Rate of Return to Schooling

As suggested by the discussion in the previous section, the typical method for estimating the rate of return to schooling uses data on the earnings and schooling of different workers and estimates the percentage wage differential associated with one more year of schooling—after adjusting the data for differences in other worker characteristics such as age, sex, and race. The "consensus" estimate of the rate of return to schooling in the United States was probably around 9 percent in the 1990s, so that schooling seems to be a good investment.[9]

The typical study estimates a regression of the form

$$\log w = bs + \text{other variables} \tag{7-10}$$

where w gives the worker's wage and s gives the number of years of schooling acquired by this worker. The coefficient b gives the percent wage differential between two workers who differ by one year of schooling (holding other variables constant) and is typically interpreted as the rate of return to schooling.

Although most of the empirical studies use this regression model to estimate the rate of return to schooling, one must not forget the central point made in the previous section. The percent wage differential between two workers who differ in their educational attainment estimates the rate of return to schooling only if the workers face the same wage-schooling locus so that there is no ability bias. Workers *do* differ in their abilities, however, and a lot of effort has been devoted to ensuring that the other variables included in the regression control for these ability differences. Some studies, for example, include measures of the

[9] Hundreds of studies estimate the rate of return to schooling in this fashion. The classic references include Gary S. Becker and Barry R. Chiswick, "Education and the Distribution of Earnings," *American Economic Review* 56 (May 1966): 358–69; Jacob Mincer, *Schooling, Experience, and Earnings*, New York: Columbia University, 1974; and Giora Hanoch, "An Economic Analysis of Earning and Schooling," *Journal of Human Resources* 2 (Summer 1967): 310–29. The "consensus" estimate of the rate of return to schooling rose between the 1970s and 1990s, from about 7 to 9 percent. Chapter 8 provides a detailed discussion of this rise in the returns to schooling. Recent international estimates of the rate of return to schooling are given by Philip Trostel, Ian Walker, and Paul Woolley, "Estimates of the Economic Return to Schooling for 28 Countries," *Labour Economics* 9 (February 2002): 1–16.

worker's IQ.[10] It is doubtful, however, that these test scores are good measures of a worker's innate productive capacity. After all, there is still an unsettled debate on what IQ measures, even in the context of scholastic achievement.

Using "Natural Experiments" to Compare Workers of the Same Ability

A number of studies have chosen a very clever way out of the fundamental problem raised by unobserved ability differences among workers. Our discussion suggests that the ability bias would disappear if we could compare the earnings of two workers who we know have the same ability but who have different levels of schooling. These two persons would necessarily face the same wage-schooling locus, and the wage gap between the two workers would provide a valid estimate of the rate of return to schooling. The comparison of the earnings of identical twins provides a natural experiment that seems to satisfy these restrictions.

Suppose that we have a sample of identical twins in which each twin reports both earnings and years of schooling. We can calculate the percentage wage differential per year of schooling for each pair of twins and average this number across the twin pairs. The average percentage wage differential is a valid estimate of the rate of return to schooling because ability differences have been completely controlled for.

Although the idea is intuitively appealing, the evidence is mixed. Some early studies reported that the rate of return to schooling in a sample of identical twins is roughly on the order of 3 percent, which is much lower than the rate of return typically estimated in studies that do not adjust for ability bias. These studies conclude that ability differences account for much of the earnings gap between highly educated and less educated workers.[11] More recent studies, however, find that using data on twins raises the rate of return to schooling to about 15 percent, far higher than conventional estimates.[12]

Even if the studies in the literature agreed on the direction of the ability bias, the study of identical twins raises an important question: *Why* do identical twins have different levels of schooling in the first place?

Our theoretical model of the schooling decision isolated two variables that determine how much schooling a person acquires: ability and the rate of discount. Since identical twins that differ in their schooling do not presumably differ in their innate ability, it must be the case that they had different discount rates. The identical twins, in other words, differ in important and unobserved ways. In short, the identical twins do not seem to be completely identical. Unless we can understand how and why identical twins differ, therefore, it is not clear that we should interpret the earnings differential between identical twins as a measure of the "true" rate of return to schooling.

[10] Excellent reviews of the econometric issues involved are given by Zvi Griliches, "Estimating the Returns to Schooling: Some Econometric Problems," *Econometrica* 45 (January 1977): 1–22; and David Card, "The Causal Effect of Education on Earnings."

[11] Paul Taubman, "Earnings, Education, Genetics, and Environment," *Journal of Human Resources* 11 (Fall 1976): 447–61.

[12] Orley C. Ashenfelter and Alan B. Krueger, "Estimates of the Economic Return to Schooling from a New Sample of Twins," *American Economic Review* 84 (December 1994): 1157–73; and Orley Ashenfelter and Cecilia Rouse, "Income, Schooling, and Ability: Evidence from a New Sample of Identical Twins," *Quarterly Journal of Economics* 113 (February 1998): 253–84.

Instrumental Variables: Compulsory Schooling Laws and Month of Birth

Many government policies generate instruments that allow the comparison of earnings among equally able workers. One particularly famous example is the existence of compulsory schooling legislation. Some states, for instance, enact compulsory schooling laws that force workers to remain in school until they reach some predetermined age, such as 16 or 17.

In the United States, children typically are not allowed to enter the first grade unless they are six years old by January 1 of the academic year in which they enter school. That means that persons born early in the year "miss" the deadline and are older when they start school than persons who are born later in the year. A compulsory schooling age of 16 then implies that children born in the early months of the year attain the legal dropout age after having attended school for a shorter time than children born near the end of the year. This variation serves as an instrument that "nudges" some persons along a particular wage-schooling locus and that can be used to estimate the rate of return to schooling.[13]

To easily understand the nature of the empirical exercise, suppose there is a compulsory schooling age of 16 and compare two children: one born on December 31 and the other born a couple of days later, on January 2. The child born on December 31 qualifies to enter the first grade at an earlier chronological age than the child born in early January. In fact, in the 1960 census, children who are born in the first quarter of the year enter the first grade when they are 6.5 years old, as compared to an age of entry of 6.1 years for children born in the last quarter of the year. As a result, even though both children will turn 16 years old at almost the same time, the child born in December will have attended school for a longer period. The relationship between compulsory schooling and month of birth would be a valid instrument—that is, it would nudge persons along the same wage-schooling locus—if the ability of children born on December 31 is the same, on average, as that of children born on January 2.

Put differently, the biological "accident" of a birth just before January 1 means that the child will be required to be in school for a longer period than a comparable child born just after January 1. The wage gap between the two children, therefore, measures the true rate of return to schooling because there should be no ability differences between them. The only reason that earnings could differ is because those born in late December have slightly more schooling, on average, than those born in early January. If one controls for ability bias in this fashion, the estimated rate of return to schooling is on the order of 7.5 percent.[14]

[13] Joshua Angrist and Alan B. Krueger, "Does Compulsory Schooling Affect Schooling and Earnings?" *Quarterly Journal of Economics* 106 (November 1991): 979–1014. A critical appraisal of this study is given by John Bound, David A. Jaeger, and Regina Baker, "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak," *Journal of the American Statistical Association* 90 (June 1995): 443–50. For an application of the basic framework to German data, see Jörn-Steffen Pischke and Till von Wachter, "Zero Returns to Compulsory Schooling in Germany: Evidence and Interpretation," Centre for the Economics of Education, London School of Economics, January 2006.

[14] Another example of an instrument that has been used to estimate the rate of return to schooling is the GI bill that subsidized education expenditures for World War II veterans. See Marcus Stanley, "College Education and the Midcentury GI Bills," *Quarterly Journal of Economics* 118 (May 2003): 671–708; and John Bound and Sarah Turner, "Going to War and Going to College: Did World War II and the G.I. Bill Increase Educational Attainment for Returning Veterans?" *Journal of Labor Economics* 2002 (October 2002): 784–815. A good review of the literature is given by Jeffrey R. Kling, "Interpreting Instrumental Variables Estimates of the Return to Schooling," *Journal of Business and Economic Statistics* 19 (July 2001): 358–64.

7-6 Policy Application: School Construction in Indonesia

Many studies document that the wage gap between high-educated and low-educated workers in developing countries is even higher than the gap in industrialized economies.[15] It is tempting to infer from these findings that developing labor markets offer a high rate of return to schooling, and that these high rates of return justify sizable investments in the education infrastructure. As we have seen, however, these wage gaps need not suggest that increasing schooling opportunities for a wide segment of the population would substantially improve the earnings of those workers.

In Indonesia, children typically go to school between the ages of 7 and 12. In 1973, the Indonesian government launched a major school construction program (INPRES) designed to increase the enrollment of children in disadvantaged areas.[16] By 1978–79, more than 61,000 new primary schools had been built, approximately two schools per 1,000 children. The typical school was designed for three teachers and 120 students. This construction program cost almost $700 million (2002 U.S. dollars), representing 1.5 percent of the Indonesian GDP as of 1973. As a way of grasping the scale of the construction, a similar commitment by the United States (in terms of GDP share) would require an expenditure of around $150 billion.

It has been reported that INPRES was the fastest primary school construction program in world history. The results were immediate: enrollment rates among children aged 7 to 12 rose from 69 percent in 1973 to 83 percent by 1978.

A recent study uses data drawn from the Indonesian labor market in 1995 (two decades after the school construction) to determine if the huge investment increased the educational attainment and earnings of the targeted Indonesians, and also to calculate the rate of return to schooling in Indonesia. As noted above, the program attempted to equalize education opportunities across the various regions of Indonesia, building more schools in those parts of Indonesia that had relatively low enrollment rates. Table 7-3 illustrates how education and earnings were affected for persons residing in two different parts of Indonesia—the "high-construction" area, where many new schools were built, and the "low-construction" area, where relatively few schools were built. In rough terms, about one more school per 1,000 children was built in the high-construction area than in the low-construction area.

The table examines the outcomes experienced by two different demographic groups: persons who were 2–6 years old and 12–17 years old as of 1974. The younger of these groups was clearly affected by the construction program. These boys and girls were about to enter school as the construction program began, and they form the treatment group. The older persons—the control group—were past the school-going age, and their educational attainment should not be affected by the presence of more schools.

[15] John Strauss and Duncan Thomas, "Human Resources: Empirical Modeling of Household and Family Decisions," in Jere Behrman and T. N. Srinivasan, editors, *Handbook of Development Economics*. Amsterdam: Elsevier, 1995, pp. 1885–2023.

[16] The discussion in this section is based on the findings reported in Esther Duflo, "Schooling and Labor Market Consequences of School Construction in Indonesia," *American Economic Review* 91 (September 2001): 795–813.

TABLE 7-3 The Impact of School Construction on Education and Wages in Indonesia

Source: Esther Duflo, "Schooling and Labor Market Consequences of School Construction in Indonesia."

	Years of Education			Log Wages		
	Persons Aged 12–17 in 1974	Persons Aged 2–6 in 1974	Difference	Persons Aged 12–17 in 1974	Persons Aged 2–6 in 1974	Difference
Low-construction area	9.40	9.76	0.36	7.02	6.73	−0.29
High-construction area	8.02	8.49	0.47	6.87	6.61	−0.26
Difference-in-differences	—	—	**0.11**	—	—	**0.03**

Table 7-3 uses the difference-in-differences methodology to calculate the impact of the construction on the educational attainment of the targeted population. In the low-construction area, the educational attainment increased by 0.36 year between the older and younger cohorts, while in the high-construction area, the educational attainment rose by 0.47 year. The difference-in-differences approach thus suggests that the additional construction increased educational attainment by 0.11 year. By using a similar approach, the table also shows that the earnings of the younger cohort living in the high-construction area rose by an additional 3 percent.

We can now use the method of instrumental variables to calculate the rate of return to schooling in Indonesia. The instrument is school construction. This variable clearly "nudged" some students along the wage-schooling locus. The instrument is valid if students in the high-construction areas have the same ability as those in the low-construction areas and if the older cohort of students has the same innate ability as the younger cohort. Each additional 0.11 year of schooling increased earnings by 3 percent. This implies that each additional year of school increased earnings by 27 percent (or 0.03 ÷ 0.11). The rate of return to schooling in Indonesia, therefore, seems to be quite high, justifying the sizable expenditure made by the school construction program. In fact, a more thorough analysis of the data, which controls for many of the other factors that also affected trends in educational attainment and wages in Indonesia, suggests that the rate of return to schooling may be as high as 10 percent.

7-7 Policy Application: School Quality and Earnings

Conventional wisdom has it that today's high school graduates are not as good as yesterday's graduates. The media often report that a large fraction of high school graduates are "functionally" illiterate despite the fact that expenditures on primary and secondary education rose dramatically in the past two decades (per-student real expenditures in public schools increased from $4,600 in 1980 to $8,400 in 2006).[17] Does "throwing money" at the

[17] U.S. Bureau of the Census, *Statistical Abstract of the United States, 2006.* Washington, DC: Government Printing Office, 2006, Table 206. A detailed survey of the trends in expenditures and "quality" of the public schools is given by Eric A. Hanushek, "The Economics of Schooling: Production and Efficiency in the Public Schools," *Journal of Economic Literature* 24 (September 1986): 1141–77.

Theory at Work

CAN WE AFFORD TO IMPROVE THE SKILLS OF HIGH SCHOOL DROPOUTS?

Education is clearly a good investment if the rate of return to education is at least 7 percent. Nevertheless, it would be extremely expensive to pursue a policy that injects less-educated workers with additional schooling. There are around 20 million persons over the age of 25 who do not have a high school diploma in the United States. High school dropouts, on average, earn about $10,000 less than the typical high school graduate. Suppose the government injects high school dropouts with enough schooling so as to reduce their wage disadvantage by, say, $5,000 a year. Let's even assume that the rate of return to schooling is high, on the order of 10 percent.

How much would this policy cost the taxpayers? In order to increase annual earnings by $5,000 for a worker, the government has to invest $50,000 on each high school dropout. Because there are 20 million high school dropouts, the policy would cost $1 *trillion.*

It is instructive to compare these projected costs with actual expenditures on programs that subsidize human capital investments. The Clinton administration's highly touted (and ill-fated) 1993 proposal to stimulate the economy by "investing in workers" was budgeted at $16.5 billion annually. Not even an annual appropriation of around $100 billion would come close to narrowing the wage gap between high school dropouts and high school graduates.

Source: Based on James J. Heckman, "Is Job Training Oversold?" *Public Interest* 115 (Spring 1994): 91–115.

public school system raise the rate of return to schooling? Put differently, does school quality, as measured by teacher salaries or pupil/teacher ratios, matter?

Prior to 1992, the consensus was that high levels of school expenditures had little impact on educational or labor market outcomes. As an influential survey concluded, "There appears to be no strong or systematic relationship between school expenditures and student performance."[18] The 1992 publication of an influential study by David Card and Alan Krueger, showing that school quality is indeed positively correlated with the rate of return to schooling, sparked a heated debate over the economic importance of school quality.[19]

Card and Krueger used data on worker earnings from the 1980 census to calculate the rates of return to schooling to cohorts of workers born in a particular state; for example, workers born in Kansas between 1920 and 1929. The two panels of Figure 7-6 summarize the core of the Card-Krueger evidence. There is obviously a great deal of variation in the rate of return to schooling for workers in this age cohort, depending on where they were born. The range in the rate of return is from 3 percent (for workers born in Louisiana) to slightly over 7 percent (for those born in Wyoming). Figure 7-6*a* shows that the rate of return to schooling is negatively correlated with the state's pupil/teacher ratio, while Figure 7-6*b* shows that the rate of return to schooling is positively correlated with the state's average

[18] Hanushek, "The Economics of Schooling: Production and Efficiency in Public Schools."

[19] David Card and Alan B. Krueger, "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy* 100 (February 1992): 1–40; see also David Card and Alan B. Krueger, "School Resources and Student Outcomes: An Overview of the Literature and New Evidence from North and South Carolina," *Journal of Economic Perspectives* 10 (Fall 1996): 31–50.

FIGURE 7-6 School Quality and the Rate of Return to Schooling

Source: David Card and Alan B. Krueger, "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy* 100 (February 1992), Tables 1 and 2. The data in the graphs refer to the rate of return to school and the school quality variables for the cohort of persons born in 1920–1929.

FIGURE 7-6
School Quality and the Rate of Return to Schooling

Source: David Card and Alan B. Krueger, "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy* 100 (February 1992), Tables 1 and 2. The data in the graphs refer to the rate of return to school and the school quality variables for the cohort of persons born in 1920–1929.

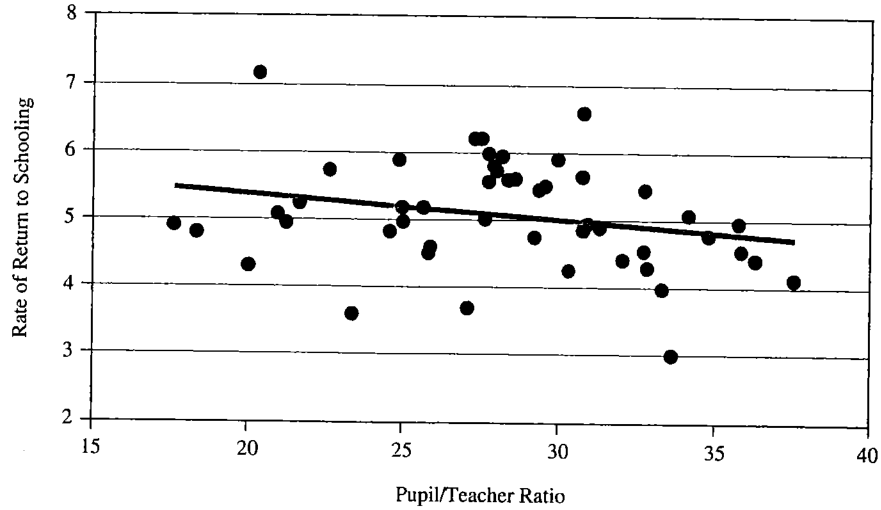

(*a*) Impact of Pupil-Teacher Ratio

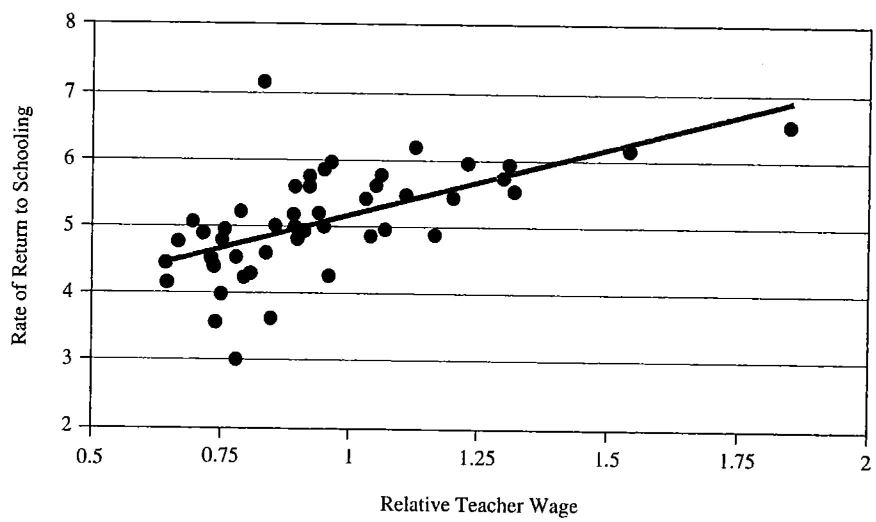

(*b*) Impact of Teacher Salaries

teacher salary (relative to the average wage in the state). After analyzing these data as well as the rates of return to schooling of other birth cohorts, Card and Krueger concluded that children born in states that offered better schools had a substantially higher rate of return to schooling. Decreasing the pupil/teacher ratio by 10 students increased the rate of return by about 1 percentage point, whereas increasing the relative wage of teachers by 30 percent (which presumably attracts better teachers) increased the rate of return to schooling by 0.3 percentage point.

The striking findings in this article motivated a great deal of research in the past decade that attempts to determine the robustness of the correlations. Although many of the subsequent studies report evidence that contradict the Card-Krueger findings, it is difficult to understand why the evidence is so mixed.[20] After all, why do elementary schools incur the extra cost of breaking up 100 third-grade students into four sections with four teachers if the students would be just as well off herded together in one big section? Moreover, there is evidence documenting a strong positive relation between property values and school quality.[21] Why would parents pay more for housing in school districts that offer smaller classes and better teachers if these inputs do not matter?

To resolve some of the confusion, a number of recent studies have analyzed experimental data, observing the outcomes of students who are randomly assigned to classes of different sizes. Beginning in 1985, the Tennessee Student/Teacher Achievement Ratio (STAR) experiment randomly assigned kindergarten students *and* their teachers to small classes (with a pupil/teacher ratio of 13–17) or to larger classes (with a ratio of 22–25 students). After the initial assignment, students remained in the same class type for four years. Between 6,000 and 7,000 students were involved in this experiment each year. A careful evaluation of the data resulting from the STAR experiment indicates that students assigned to the small classes scored higher in achievement tests than students assigned to the larger classes.[22]

Other studies use the method of instrumental variables to estimate the impact of class size on scholastic achievement in nonexperimental settings. The problem, of course, is finding a variable that affects class size but does not affect other outcomes directly. One recent study used an instrument based on the interpretation of the Talmud by the twelfth-century rabbinic scholar Maimonides.[23] According to Maimonides's rule, "Twenty-five children may be put in charge of one teacher. If the number in the class exceed twenty-five but is not more than forty, he should have an assistant to help with the instruction. If there are more than forty, two teachers must be appointed."

[20] See, for example, Julian R. Betts, "Does School Quality Matter? Evidence from the National Longitudinal Survey," *Review of Economics and Statistics* 77 (May 1995): 231–50; James J. Heckman, Anne S. Layne-Farrar, and Petra E. Todd, "Does Measured School Quality Really Matter," in Gary Burtless, editor, *Does Money Matter? The Effect of School Resources on Student Achievement and Adult Success.* Washington, DC: The Brookings Institution, 1996; Jeffrey Grogger, "Does School Quality Explain the Recent Black/White Wage Trend?" *Journal of Labor Economics* 14 (April 1996): 231–53; Carolyn M. Hoxby, "The Effects of Class Size on Student Achievement: New Evidence from Population Variation," *Quarterly Journal of Economics* 115 (November 2000): 1239–85; and Eric A. Hanushek, "Some Simple Analytics of School Quality," National Bureau of Economic Research Working Paper No. 10229, January 2004. A theoretical study of the relation between various measures of school quality and academic achievement is given by Edward P. Lazear, "Educational Production," *Quarterly Journal of Economics* 116 (August 2001): 777–803.

[21] Sandra E. Black, "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics* 114 (May 1999): 577–99.

[22] Alan B. Krueger, "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics* 114 (May 1999): 497–532.

[23] Joshua D. Angrist and Victor Lavy, "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *Quarterly Journal of Economics* 114 (May 1999): 533–75.

The Israeli public school system uses Maimonides's rule to distribute students among various classes. The maximum class size is 40. According to Maimonides's rule, class size increases with enrollment until 40 pupils are enrolled. An extra student, however, implies that class size drops sharply to 20.5. Because there is little reason to suspect that the shift from a class size of 40 to one of 20.5 has anything to do with the underlying ability of the students, Maimonides's rule provides a valid instrument—it shifts class size without affecting any other variables. The analysis of the outcomes experienced by Israeli students consistently suggests a negative relation between class size and academic achievement.[24]

Finally, a number of studies have investigated if attending an "elite" college or university, which presumably offers a higher quality of education, affects earnings. The problem with comparing the earnings of students who attend selective institutions with the earnings of students who do not is that there may be underlying ability differences between the two groups that will affect earnings. Any resulting wage gap may have little to do with the "value added" by the selective institution, and may simply reflect the preexisting ability gap between the two groups of students.

To avoid the problem of ability bias, a recent study compares the earnings of students who attend highly selective schools, as measured by the average SAT score of freshmen, with the earnings of students who were accepted by those institutions but decided to go to a less-selective college.[25] These two groups presumably have the same underlying ability—they were all accepted by the same selective institutions. Interestingly, this comparison reveals that selective schools provide no value-added: Students who graduate from selective schools earn no more than students who were accepted by those schools but decided to go elsewhere. In short, there seems to be little return to attending a selective college.

7-8 Do Workers Maximize Lifetime Earnings?

The schooling model provides the conceptual framework that allows us to estimate the rate of return to schooling. We have seen that—under certain conditions—percent wage differentials among workers who differ in their education can be interpreted as a rate of return to schooling. This calculation of the rate of return to schooling, however, does *not* test the theory. Rather, the calculations use the theory to interpret the earnings differences among workers in a particular way.

Therefore, we want to determine if the schooling model provides a useful "story" of how students actually go about the business of deciding whether to stay in school. The schooling model assumes that persons choose the level of schooling that maximizes the present value

[24] For a study of the relation between school quality and earnings in a very different economic setting (Honduras), see Arjun Singh Bedi and John H. Y. Edwards, "The Impact of School Quality on Earnings and Educational Returns—Evidence from a Low-Income Country," *Journal of Development Economics* 68 (June 2002): 157–85.

[25] Stacy Berg Dale and Alan B. Krueger, "Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables," *Quarterly Journal of Economics* 117 (November 2002): 1491–528. Conflicting evidence is provided by Dominic J. Brewer, Eric R. Eide, and Ronald G. Ehrenberg, "Does It Pay to Attend an Elite Private College? Cross-Cohort Evidence on the Effects of College Type on Earnings," *Journal of Human Resources* 34 (Winter 1999): 104–23.

of lifetime earnings. If we could observe the age-earnings profile of a particular worker both if he were to go to college and if he were to stop after high school, it would be easy to test the key hypothesis of the schooling model. We could use these annual earnings to calculate the present value of each option, compare the two numbers, and check to see if the worker chose the one with the largest present value.

This simple test, however, can *never* be conducted. The reason is both trivial (because it is painfully obvious) and profound (because it raises a number of conceptual questions that have yet to be adequately resolved). *Once a worker makes a particular choice, we can only observe the earnings stream associated with that choice.* Consider the group of workers who go to college. For these college graduates, we can observe only their life cycle earnings after college graduation, and we will never observe what they would have earned had they not attended college. Similarly, consider the group of workers who quit after completing high school. For these high school graduates, we observe the earnings stream subsequent to their high school graduation, and we will never observe what they would have earned had they gone on to college.

It is tempting to work out a simple solution to this problem. Even though we will never observe how much a worker who quits after completing high school would have earned if he had attended college, we do observe the earnings of those workers who did attend college. We could then predict the high school graduate's earnings had he attended college by using the observed data on what college graduates actually make. Similarly, even though we do not observe how much college graduates would have earned had they stopped after high school, we do observe the earnings of high school graduates. We could then predict the college graduate's earnings (had he not attended college) from the salary data for high school graduates.

Our earlier discussion suggests that this exercise is valid only if college graduates and high school graduates lie on the same wage-schooling locus. This calculation is invalid if there are ability differences. The observed wage differential between college graduates and high school graduates reflects not only the returns to college, but also the returns to differences in ability between the two groups. Therefore, using the observed wage differential to determine if workers choose the "right" schooling option yields meaningless results.

A Numerical Example

To illustrate the importance of this problem, let's work through a simple numerical example with two workers, Willie and Wendy. Willie is particularly adept at "blue-collar" work, and this type of work requires no schooling. Wendy is particularly adept at "white-collar" work, and this type of work requires one year of schooling. Suppose also that there are two periods in the life cycle. If a person does not go to school, he works in the blue-collar job in both periods. If the person goes to school, the person would go to school in the first period and work in the white-collar job in the second period.

The wage-schooling locus for each type of worker is summarized by these data:

Worker	Earnings in Blue-Collar Job	Earnings in White-Collar Job
Willie	\$20,000	\$40,000
Wendy	\$15,000	\$41,000

Because Willie is better at doing blue-collar work, he earns more at the blue-collar job (\$20,000) than Wendy would (\$15,000). Similarly, because Wendy is better at white-collar work, she earns more in the white-collar job than Willie would.

Suppose that both Willie and Wendy have a discount rate of 10 percent. Each worker calculates the present value of lifetime earnings for each schooling option and chooses the one that has the highest present value. The present values of Willie's alternative earnings streams are

$$\text{Willie's present value if he does not go to school} = 20{,}000 + \frac{20{,}000}{1.1} = \$38{,}182 \quad (7\text{-}11)$$

$$\text{Willie's present value if he goes to school} = 0 + \frac{40{,}000}{1.1} = \$36{,}364 \quad (7\text{-}12)$$

Willie will decide that he should not go to school and will be a blue-collar worker.

The present values of Wendy's potential earnings streams are

$$\text{Wendy's present value if she does not go to school} = 15{,}000 + \frac{15{,}000}{1.1} = \$28{,}636 \quad (7\text{-}13)$$

$$\text{Wendy's present value if she goes to school} = 0 + \frac{41{,}000}{1.1} = \$37{,}273 \quad (7\text{-}14)$$

Wendy, therefore, goes to school in the first period and works in a white-collar job in the second.

What data do we observe in the labor market? We observe the earnings of persons who do not go to school and work in blue-collar jobs (like Willie). The present value of their earnings is \$38,182. We also observe the earnings of persons who do go to school and work in white-collar jobs (like Wendy). The present value of their earnings stream is \$37,273.

A comparison of the two numbers that can be observed would suggest that Wendy made a terrible mistake. In our numerical example, persons who go to school earn less over their lifetime than persons who do not go to school. Because workers sort themselves into particular occupations, however, this is an irrelevant comparison. Both Willie and Wendy made the right choice. The problem lies in comparing the earnings of the two types of workers. This comparison is akin to comparing apples and oranges and is contaminated by **selection bias**, the fact that workers self-select themselves into jobs for which they are best suited. In our numerical example, the selection bias leads to an incorrect rejection of the validity of the human capital model.

Selection Bias Corrections

In view of the significance of the issues associated with selection bias, it is not surprising that a great deal of study has been devoted to this problem. This research has developed statistical techniques, known as "selection bias corrections," that allow us to test correctly the hypothesis underlying the schooling model.[26] These techniques give a statistically valid

[26] James J. Heckman, "Sample Selection Bias as a Specification Error," *Econometrica* 47 (January 1979): 153–62. See also James J. Heckman, "Varieties of Selection Bias," *American Economic Review* 80 (May 1990): 313–18; and Charles F. Manski, "Anatomy of the Selection Problem," *Journal of Human Resources* 24 (Summer 1989): 343–60.

methodology for predicting what a high school graduate would have earned had he attended college and what a college graduate would have earned had he quit school after getting a high school diploma.

A well-known study uses these selection bias corrections to estimate the life cycle earnings profiles associated with each of two alternatives (going to college or quitting after high school) for a large number of workers.[27] The empirical analysis confirms the basic hypothesis of the theory: On average, workers choose the schooling option that maximizes the present value of lifetime earnings. Moreover, the evidence indicates that when both a high school graduate and a college graduate are placed in the type of job that high school graduates typically fill, the high school graduate would be more productive. Conversely, if both high school graduates and college graduates were placed in jobs typically filled by college graduates, the college graduate would be more productive.

As implied by our numerical example, this empirical result suggests that the notion that there is only one type of ability that inevitably leads to higher earnings does not correctly describe how workers differ in the labor market. There exist various types of abilities, and each of us may be particularly adept at doing some things and quite inept at doing others. Some persons have a knack for doing work that is best learned in college, whereas other persons have a knack for doing blue-collar work. Put differently, some workers have a *comparative advantage* at doing skilled work; other workers have a comparative advantage at doing less-skilled work.

7-9 Schooling as a Signal

The schooling model is based on the idea that education increases a worker's productivity and that this increase in productivity raises wages. An alternative argument is that education need not increase the worker's productivity at all, but that "sheepskin" levels of educational attainment (such as a high school or college diploma) signal a worker's qualifications to potential employers.[28] In this view, education increases earnings not because it increases productivity, but because it certifies that the worker is cut out for "smart" work. Education can play this signaling role only when it is difficult for potential employers to observe the worker's ability directly. If the employer could determine cheaply whether the worker is qualified for the job, the firm would not have to rely on third-party certifications.

To illustrate how workers decide how much schooling to get when education plays only a signaling role, let's work through a simple numerical example. Suppose there are two

[27] Robert J. Willis and Sherwin Rosen, "Education and Self-Selection," *Journal of Political Economy* 87 (October 1979 Supplement): S7–S36. See also Lawrence W. Kenny, Lung-Fei Lee, G. S. Maddala, and R. P. Trost, "Returns to College Education: An Investigation of Self-Selection Bias Based on the Project Talent Data," *International Economic Review* 20 (October 1979): 775–89; and John Garen, "The Returns to Schooling: A Selectivity-Bias Approach with a Continuous Choice Variable," *Econometrica* 52 (September 1984): 1199–218.

[28] A. Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87 (August 1973): 355–74. See also Kenneth J. Arrow, "Higher Education as a Filter," *Journal of Public Economics* 2 (July 1973): 193–216; and Joseph Stiglitz, "The Theory of Screening, Education, and the Distribution of Income," *American Economic Review* 65 (June 1975): 283–300. For evidence on sheepskin effects, see David A. Jaeger and Marianne E. Page, "Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education," *Review of Economics and Statistics* 78 (November 1996): 733–40.

types of workers in the labor market, low-productivity workers and high-productivity workers, and that the distribution of productivity in the population is given by

Type of Worker	Proportion of Population	Present Value of Lifetime Productivity
Low-productivity	q	\$200,000
High-productivity	$1 - q$	300,000

The productivity differences between the two types of workers exist since birth and have *nothing* to do with how much schooling a particular worker gets.

If an employer could determine easily if a job applicant is a high-productivity worker, he would pay the worker \$300,000 over the life cycle. After all, if the employer's wage offer did not match the high-productivity applicant's true value, the job applicant would simply go elsewhere, where his high productivity was better appreciated. Similarly, if the employer could determine easily that the applicant is a low-productivity worker, he would pay the worker only \$200,000.

But life is not quite this easy. Even though a particular worker knows which group he belongs to, it might take some years for the employer to learn that. Moreover, if an employer asks the job applicant if he is a low- or high-productivity worker, the applicant (who wants a high salary) will always reply that he is a high-productivity worker regardless of his true ability. When a job applicant shows up at the firm, therefore, there is a great deal of uncertainty about whether he is a low-productivity or a high-productivity worker.

Pooling Workers

Because low-productivity workers will always lie about their productivity, the firm will disregard what anyone says about their own qualifications. In the absence of any other information, therefore, the employer simply pools all job applicants and treats them identically. The average productivity and salary of the workers hired by the firm is then given by

$$\begin{aligned}\text{Average salary} &= (200{,}000 \times q) + [300{,}000 \times (1 - q)] \\ &= 300{,}000 - 100{,}000q \qquad \text{(7-15)}\end{aligned}$$

The average salary is simply a weighted average of the workers' productivities, where the weights are the proportions in the population that belong to each productivity group.

Because the proportion q is between 0 and 1, the average salary in this "pooled equilibrium" is between \$200,000 and \$300,000. Low-productivity workers prefer a pooled equilibrium because they are being pooled with more productive workers, who push up their salary. Neither employers nor high-productivity workers like the pooled equilibrium. Employers find that they are mismatching workers and jobs. Some high-productivity workers are being assigned to menial jobs, and low-productivity workers are placed in jobs that they are not qualified to perform. This mismatching reduces the firm's efficiency and output. Similarly, the earnings of high-productivity workers are dragged down by the low-productivity workers, and, hence, the high-productivity workers would like to find a way of demonstrating to the employer that they truly are more productive.

A Signal Helps Distinguish the Workers

High-productivity workers have an incentive to provide *and* firms have an incentive to take into account credible information that can be used to allocate the worker into either productivity group. This type of information is called a **signal.** It turns out that an educational diploma or certificate can perform this signaling job and that it can perform the task with absolute precision. *No mismatches occur.*

Suppose a firm chooses the following rule of thumb for allocating workers to the two types of jobs. If a worker has at least $\bar{y}$ years of college, the firm assumes that the worker is a high-productivity worker, allocates him to a job that requires a high level of skills, and pays him a (lifetime) salary of \$300,000. If a worker has fewer than $\bar{y}$ years of college, the firm assumes that the worker is a low-productivity worker, allocates him to an unskilled job, and pays him a salary of \$200,000.

Because employers are willing to pay more to workers who get at least $\bar{y}$ years of college, all workers will want to get the required college credits. Obtaining these credits, however, is expensive. We assume that obtaining credits is more expensive for less-able workers; in particular, a year's worth of college credits costs \$20,000 for a high-productivity worker, but \$25,001 for a low-productivity worker. Obviously, tuition and fees do not differ according to ability, but the *real* cost of a college credit is higher for a low-productivity worker. To attain a particular level of achievement, a low-productivity worker will have to devote more time to studying and may have to pay for tutors, study guides, and special classes. This assumption that low-productivity workers find it more costly to obtain the signal is the fundamental assumption of the signaling model—and, in fact, is what makes the model work.

Given the firm's wage offer, workers must now decide how many years of college to get. A "separating equilibrium" occurs when low-productivity workers choose not to get $\bar{y}$ years of schooling and voluntarily signal their low productivity, and high-productivity workers choose to get at least $\bar{y}$ years of schooling and separate themselves from the pack.

Figure 7-7*a* illustrates the firm's wage offer and the cost function facing a low-productivity worker. The wage offer is such that if the worker has fewer than $\bar{y}$ years of college, he earns \$200,000, and if he has $\bar{y}$ or more years, he earns \$300,000. The cost function is upward sloping and has a slope of \$25,001 because each additional year of college costs \$25,001 for a low-productivity worker.

In our numerical example, a worker will decide either not to go to college at all or to go to college for $\bar{y}$ years. After all, a worker's earnings do not increase if he goes to college for more than $\bar{y}$ years, yet it costs the worker \$25,001 to get an additional year's worth of college credits. Similarly, because the worker's lifetime salary equals \$200,000 for *any* level of education between 0 and $\bar{y}$ years of college, there is no point to getting "just a few" credits.

A separating equilibrium requires that low-productivity workers do not go to college at all. This will occur whenever the net return from getting zero years of college exceeds the net return from getting $\bar{y}$ years. Figure 7-7*a* indicates that when a low-productivity worker does not go to college, he "takes home" \$200,000 (because he does not incur any college attendance costs). If he goes to college $\bar{y}$ years, his net salary is the vertical difference between the \$300,000 wage offer and the cost of going to college for $\bar{y}$ years (which equals \$25,001 $\times$ $\bar{y}$). Therefore, the low-productivity worker will not attend college if

$$\$200{,}000 > \$300{,}000 - (\$25{,}001 \times \bar{y}) \qquad \textbf{(7-16)}$$

FIGURE 7-7
Workers get paid
workers find it e
result, the worke

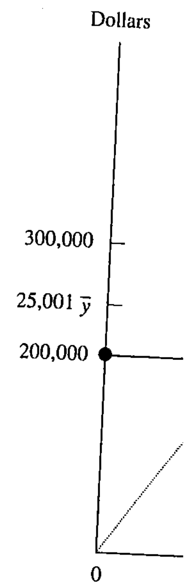

FIGURE 7-7 Education as a Signal

Workers get paid \$200,000 if they get less than $\bar{y}$ years of college, and \$300,000 if they get at least $\bar{y}$ years. Low-productivity workers find it expensive to invest in college, and will not get $\bar{y}$ years. High-productivity workers do obtain $\bar{y}$ years. As a result, the worker's education signals if he is a low-productivity or a high-productivity worker.

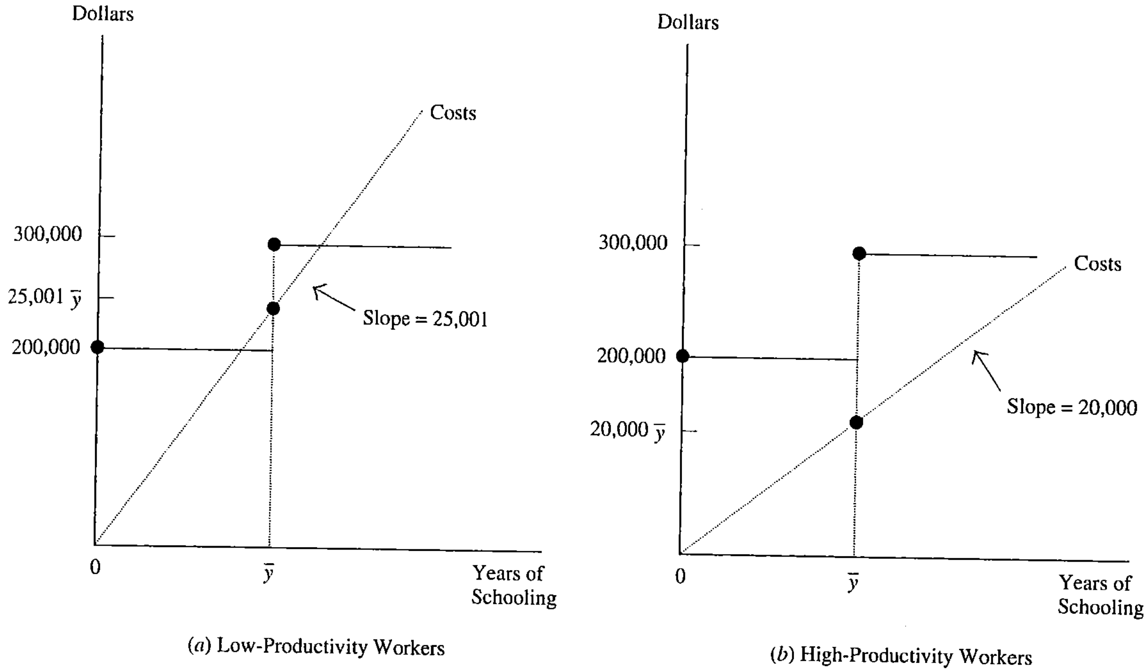

Solving for $\bar{y}$ implies that

$$\bar{y} > 3.999 \tag{7-17}$$

In other words, if the firm chooses a rule of thumb where only workers who get more than 3.999 years of college will be considered high-productivity workers, *no* low-productivity worker will bother going to college—it is just too expensive. By choosing not to attend college, low-productivity workers "voluntarily" signal their low productivity and separate themselves out.

A separating equilibrium also requires that high-productivity workers *do* get $\bar{y}$ years of college. Figure 7-7*b* illustrates their decision. The net salary of a high-productivity worker who does not go to college is \$200,000. The net salary of a high-productivity worker who goes to college for $\bar{y}$ years is the vertical difference between the \$300,000 wage offer and the cost of going to college (which equals \$20,000 $\times$ $\bar{y}$). Therefore, high-productivity workers get $\bar{y}$ years of college whenever

$$\$200,000 < \$300,000 - (\$20,000 \times \bar{y}) \tag{7-18}$$

Solving for $\bar{y}$ yields

$$\bar{y} < 5 \tag{7-19}$$

In other words, as long as the firm does not demand "too many" years of higher education (such as a master's degree or Ph.D.), high-productivity workers go to college and voluntarily signal that they are highly productive.

Putting together both conditions implies that low-productivity workers do not go to college and that high-productivity workers do whenever

$$3.999 < \bar{y} < 5 \qquad \textbf{(7-20)}$$

A firm can choose any hiring standard in this range and generate a separating equilibrium. The firm can say, for instance, that workers who obtain more than 4.5 years of college will be considered high-productivity workers, and the two types of workers will sort themselves out accordingly. There seem to be an infinite number of valid thresholds that the firm can use (4 years of college, 4.5 years, 4.666 years, 4.999 years, and so on). Not all of these solutions, however, can survive the competitive pressures of the marketplace. Suppose, for instance, that some firms require 4.333 years of college to allocate high-productivity workers into skilled jobs, whereas other firms require only 4.000 years. High-productivity workers prefer the firm with the 4.000 hiring threshold because both firms pay the same competitive salary (of \$300,000) and high-productivity workers have nothing to gain from getting more education than the minimum required. The competitive solution, therefore, is the smallest possible threshold, so that using a college diploma (four years of college) to separate out job applicants generates a separating equilibrium.

The signaling model shows that education can play the role of signaling the worker's innate ability without increasing the worker's productivity. It has been extremely difficult, however, to establish empirically if education plays a productivity-enhancing role, a signaling role, or a combination of the two.[29] Regardless of which model is correct, an outsider looking at a particular labor market will observe that more-educated workers earn higher wages. Because both the schooling model and the signaling model predict that more education leads to higher earnings, the positive correlation between earnings and education cannot be used to disentangle which of the two mechanisms is more important in the labor market. As a result, there is no widely accepted calculation that decomposes the wage differential between highly educated and less-educated workers into its productivity and signaling components.

There are reasons to believe, however, that the signaling role of education may be less important than is commonly assumed. It costs over \$150,000 for the typical student to go to college (including both forgone earnings and direct costs). If college provided only an impressive-looking piece of paper, firms that specialize in printing equally impressive pieces of paper at lower prices would appear in the marketplace. The fact that we do not see a large industry of firms that sell credentials certifying a person's innate productivity must

[29] See, for example, Kenneth Wolpin, "Education and Screening," *American Economic Review* 67 (December 1977): 949–58; Kevin Lang and David Kropp, "Human Capital versus Sorting: The Effects of Compulsory Schooling Laws," *Quarterly Journal of Economics* 101 (August 1986): 609–24; Eugene A. Kroch and Kriss Sjoblom, "Schooling as Human Capital or a Signal: Some Evidence," *Journal of Human Resources* 29 (Winter 1994): 156–80; Joseph G. Altonji and Charles R. Pierret, "Employer Learning and the Signaling Value of Education," in Isao Ohashi and Toshiaki Tachibanaki, editors, *Internal Labour Markets, Incentives, and Employment*, Macmillan: New York, 1998, pp. 159–95; and Kelly Bedard, "Human Capital versus Signaling Models: University Access and High School Dropouts," *Journal of Political Economy* (August 2001): 749–75.

Theory at Work
IS THE GED BETTER THAN NOTHING?

There has been a substantial increase in the number of persons who obtain their high school diplomas by passing an equivalency test rather than by going through the normal route of spending 12 years in a classroom and then graduating from high school. In 1968, only 5 percent of high school graduates obtained their diplomas by taking the GED test (which stands for General Equivalency Diploma). By 1987, 14 percent of persons receiving a high school diploma received GED certificates.

A comparison of the earnings of traditional high school graduates with the earnings of workers who get their high school diplomas via the GED can help determine if the schooling process actually matters. In other words, does passing the GED provide the same skills as attending school for 12 years?

A recent study reports that the labor market characteristics of GED graduates and high school *dropouts* are virtually indistinguishable. In particular, the wages of GED graduates are no higher than the wages of high school dropouts. It seems, therefore, that simply certifying someone who passes a standardized test to be a high school graduate is no substitute for the learning that takes place when persons actually go to school. As the authors of the study conclude, "there is no cheap substitute for schooling."

Source: Stephen V. Cameron and James J. Heckman, "The Nonequivalence of High School Equivalents," *Journal of Labor Economics* 11 (January 1993, Part 1): 1–47; and James J. Heckman and Paul A. LaFontaine, "Bias-Corrected Estimates of GED Returns," *Journal of Labor Economics* 24 (July 2006): 661–700. Contrary evidence is provided by John H. Tyler, Richard J. Murnane, and John B. Willett, "Estimating the Labor Market Signaling Value of the GED," *Quarterly Journal of Economics* 115 (May 2000): 431–68.

imply that education does more than just signal a worker's productivity; it also must alter the human capital stock.

Although this is a sensible argument, it would be preferable to document empirically the relative importance of the signaling and productivity-enhancing roles of education. Separating out the two effects is important because the human capital framework and the signaling hypothesis have very different implications for many policy questions. The human capital model, for example, suggests that human capital investments, such as education, provide a way out of low incomes and poverty. Indeed, the rationale behind government programs that subsidize on-the-job training and tuition expenses is that these programs increase the human capital stock of the targeted groups. The signaling model says that education does not really increase a worker's innate productivity. Low-productivity workers remain low-productivity workers regardless of the billions of dollars spent on these government programs.

Private and Social Rates of Return

The different policy recommendations made by the two approaches suggest that the **private rate of return to schooling,** as measured by the increase in a worker's earnings resulting from an additional year of schooling, may differ substantially from the **social rate of return to schooling,** as measured by the increase in national income resulting from the same year of education. Suppose the signaling model is correct and education does not increase productivity. From a worker's point of view, education still has a positive private rate of return. The highly productive worker gains from signaling that he is highly productive. From a social point of view, however, educational expenditures are wasteful. National income is not increased because the worker's productivity is the same both before and after the investment in education. The social rate of return is zero.

These conclusions, however, ignore the fact that—even in the context of the signaling model—education serves the very useful role of sorting workers into the right jobs. The employer can use the education signal to allocate highly productive workers to so-called skilled jobs and to allocate the less-productive workers into other types of jobs. The mismatching of workers and jobs in the labor market—for instance, assigning a low-productivity worker to run a nuclear power plant—would surely have a detrimental effect on national income. As a result, education could have a positive social rate of return even if it does not increase a particular worker's human capital. We know very little about the magnitude of the misallocation costs that would arise if education did not help sort workers among jobs so that many of the questions concerning the "true" magnitude of the social rate of return to education have not been answered.

Some recent studies have argued that the definition of the "social" rate of return to schooling should be expanded to include the impact of education on civic engagement and attitudes in a democracy, or the impact of an educated workforce on the rate of economic growth. In fact, additional schooling raises voter participation rates and support for free speech and leads to a better-informed electorate (as measured by the frequency of newspaper readership).[30] Similarly, the evidence suggests that a more educated workforce may promote faster growth.[31]

7-10 Post-School Human Capital Investments

The evolution of wages over the life cycle is illustrated by the age-earnings profiles presented in Figure 7-8, which report the average weekly earnings of U.S. workers in a particular schooling group at different ages. The figure reveals three important properties of age-earnings profiles:

1. *Highly educated workers earn more than less-educated workers.* We have seen that education increases earnings either because education increases productivity or because education serves as a signal of a worker's innate ability.
2. *Earnings rise over time, but at a decreasing rate.* The wage increase observed over the life cycle suggests that a worker's productivity rises even after leaving school, perhaps as a result of on-the-job or off-the-job training programs.[32] The rate of wage growth, however, slows down as workers get older. Younger workers seem to add more to their human capital than older workers.
3. *The age-earnings profiles of different education groups diverge over time.* Put differently, earnings increase faster for more educated workers. The steeper slope of age-earnings

[30] Thomas S. Dee, "Are There Civic Returns to Education?" *Journal of Public Economics* 88 (August 2004): 1697–720.

[31] Mark Bils and Peter J. Klenow, "Does Schooling Cause Growth?" *American Economic Review* 90 (December 2000): 1160–83; and Eric A. Hanushek and Dennis D. Kimko, "Schooling, Labor-Force Quality, and the Growth of Nations," *American Economic Review* 90 (December 2000): 1184–208.

[32] This interpretation of upward-sloping age-earnings profiles assumes that the worker's productivity rises throughout the life cycle. As we will see in Chapter 12, other models of the labor market imply an upward-sloping age-earnings profile even if the worker's productivity is constant over time.

FIGURE 7-8
Age-Earnings Profiles of Full-Time Workers, 2005

Source: U.S. Bureau of Labor Statistics, *Annual Demographic Supplement of the Current Population Surveys*, 2005.

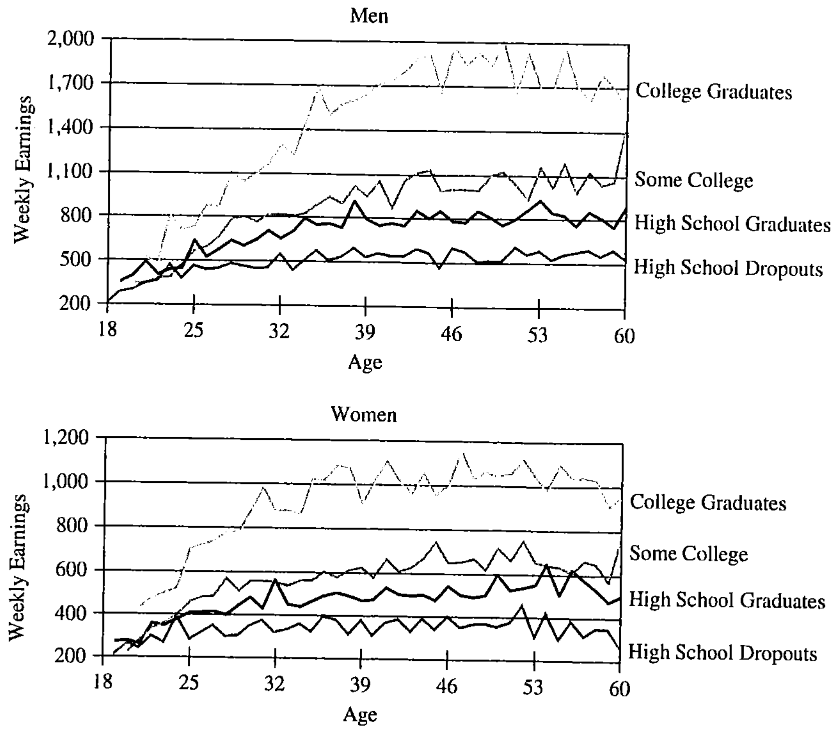

profiles for more-educated workers suggests a complementarity between investments in education and investments in on-the-job training. In other words, workers who are highly educated experience the fastest wage growth because they also invest the most during the postschool period. This complementarity between pre- and postschool investments might arise if some workers have a "knack" for acquiring all types of human capital.[33]

7-11 On-the-Job Training

Until now, we have focused on one particular aspect of human capital investments—the schooling decision. Most workers augment their human capital stock after completing their education, particularly through on-the-job training (OJT) programs. The diversity of OJT investments is striking: Secretaries learn word processing skills, lawyers get courtroom experience, investment bankers concoct new financial instruments, and politicians learn from failed policies. Evidently, OJT is an important component of a worker's human capital stock, making up at least half of a worker's human capital.[34]

[33] A detailed study of the link between schooling and earnings changes over the life cycle is given by Henry S. Farber and Robert Gibbons, "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111 (November 1996): 1007–47.

[34] Jacob Mincer, "On-the-Job Training: Costs, Returns, and Some Implications," *Journal of Political Economy* 70 (October 1962, Part 2): 50–79.

There are two types of OJT: **general training** and **specific training.**[35] General training is the type of training that, once acquired, enhances productivity equally in all firms. These general skills, which include such things as typing, learning how to drive, and learning how to use a calculator, are found frequently in the labor market. Specific training is the type of training that enhances productivity only in the firm where it is acquired and the productivity gains are lost once the worker leaves the firm. Examples of specific training also abound in the labor market: learning how to drive a tank in the army or memorizing the hierarchical nature of a particular organization. In reality, much OJT is a mixture of general and specific training, but the conceptual separation into purely general and purely specific training is extremely useful.

Consider a simple framework where the employment relationship between a competitive firm and the worker lasts two periods. Suppose that in the first period (when the worker is hired), the *total* labor costs equal TC_1 dollars, and in the second period, the costs equal TC_2 dollars. Similarly, the values of marginal product in each of the two periods are VMP_1 and VMP_2, respectively. Finally, let r be the rate of discount. The profit-maximizing condition giving the optimal level of employment for the firm over the two periods is

$$TC_1 + \frac{TC_2}{1+r} = VMP_1 + \frac{VMP_2}{1+r} \qquad \textbf{(7-21)}$$

The left-hand side of the equation gives the present value of the costs associated with hiring a worker over the two-period life cycle. The right-hand side gives the present value of the worker's contribution to the firm. It is easy to see that this equation generalizes the condition that the wage equals the value of marginal product. In a multiperiod framework, the analogous condition is that the present value of employment costs equals the present value of the value of marginal product.

Suppose OJT takes place only in the first period. It costs the firm H dollars to put a worker through the training. These costs include teacher salaries and the purchase of training equipment. The total cost of hiring a worker during the first period can be written as the sum of training costs H and the wage paid to the worker during the training period, or w_1. This implies that $TC_1 = w_1 + H$. Because no training occurs in the second period, the total cost of hiring the worker in the second period simply equals the wage. We can then rewrite equation (7-21) as

$$w_1 + H + \frac{w_2}{1+r} = VMP_1 + \frac{VMP_2}{1+r} \qquad \textbf{(7-22)}$$

Who Pays for General Training?

Consider the case where all training is general. In the posttraining period, the worker's value of marginal product increases to VMP_2 in *all* firms. As a result, many firms are willing to pay the worker a wage equal to VMP_2. The firm that provided the training must

[35] The concepts of general and specific training are due to Gary S. Becker, *Human Capital,* 3rd ed. Chicago: University of Chicago Press, 1993. Becker's framework continues to be the cornerstone of the human capital literature and is an essential component in the toolkit of modern labor economics.

either follow suit and increase the wage to VMP_2 or lose the worker. Therefore, the second period wage, w_2, will equal VMP_2. As a result, equation (7-22) simplifies to

$$w_1 = VMP_1 - H \tag{7-23}$$

Therefore, the first-period wage equals the value of the worker's initial marginal product minus training costs. In other words, workers pay for general training by accepting a lower "trainee wage" during the training period. In the second period, workers get the returns from the training by receiving a wage that equals the value of their posttraining marginal product. *Competitive firms provide general training only if they do not pay any of the costs.*

There are many examples of workers paying for general training through lower wages. It is common for trainees in formal apprenticeship programs to receive low wages during the training period and to receive higher wages after the training is completed. Similarly, medical interns (even though they already have a medical degree) earn low wages and work long hours during their residency, but their investment is well rewarded once they complete their training.

If a firm were to pay for general training, as some firms claim to do when they pay for the tuition of workers who enroll in an MBA program, the firm would surely attract a large number of job applicants. After all, many workers would quickly realize that this firm was offering free general training. Because the firm cannot legally enslave its employees after they receive their degree, the workers would take advantage of the free training opportunities and then run to a firm that offers them a wage commensurate with their newly acquired skills. Therefore, a firm that paid for general training and did not raise the posttraining wage would get an oversupply of trainees and the workers would quit in the posttraining period. This firm faces the worst of all possible outcomes: It pays for the training and gets none of the benefits. A profit-maximizing firm would quickly learn that it can lower the wage because there is an oversupply of trainees, passing on the training costs to the workers.[36]

Who Pays for Specific Training?

The productivity gains resulting from specific training vanish once the worker leaves the firm. As a result, the worker's alternative wage (that is, the wage that other firms are willing to pay) is *independent* of the training and equals his pretraining productivity. Who then pays for specific training and who collects the returns?[37]

Consider what would happen if the firm paid for specific training. The firm could incur the cost and collect the returns by not changing the wage in the posttraining period, even though the worker's value of marginal product in this firm has increased. Because VMP_2

[36] Recent studies have shown that noncompetitive firms may be willing to pay for general training under some circumstances; see Daron Acemoglu and Jörn-Steffen Pischke, "The Structure of Wages and Investment in General Training," *Journal of Political Economy* 107 (June 1999): 539–72. Empirical studies of who pays for training are given by John M. Barron, Mark C. Berger, and Dan A. Black, "Do Workers Pay for On-the-Job Training?" *Journal of Human Resources* 34 (Spring 1999): 235–52; and David H. Autor, "Why Do Temporary Help Firms Provide Free General Skills Training?" *Quarterly Journal of Economics* 116 (November 2001): 1409–48.

[37] A more detailed discussion of these issues is given by Masanori Hashimoto, "Firm-Specific Human Capital as a Shared Investment," *American Economic Review* 71 (June 1981): 475–82. For an alternative approach to modeling specific training, see Edward P. Lazear, "Firm-Specific Human Capital: A Skill-Weights Approach," National Bureau of Economic Research Working Paper No. 9679, May 2003.

would then exceed w_2, there are gains to providing the training. If the worker were to quit in the second period, however, the firm would suffer a capital loss. The firm, therefore, would hesitate paying for specific training unless it has some assurance that the trained worker will not quit.

Suppose instead that the worker pays for the specific training. Workers would then receive a low wage during the training period and higher wages in the posttraining period. The worker, however, does not have an ironclad assurance that the firm will employ him in the second period. If the worker were to get laid off, he would lose his investment because specific training is not portable. The worker, therefore, is not willing to invest in specific training unless he is very confident that he will not be laid off.

Both the firm and the worker, therefore, are reluctant to invest in specific training. The problem arises because there does not exist a legally binding contract that ties together workers and firms "until death do them part." Neither party wishes to take the initiative and pay for the training.

The way out of this dilemma is to note that fine-tuning the posttraining wage can reduce the probabilities of *both* quits and layoffs. Consider a labor contract in which the worker's posttraining wage, w_2, is set such that

$$\overline{w} < w_2 < VMP_2 \qquad \textbf{(7-24)}$$

where $\overline{w}$ is the alternative wage. This contract implies that the worker and the firm share the returns from specific training. The worker's posttraining wage w_2 is higher than his productivity elsewhere, but less than his productivity at the current firm. Note that because the worker is better off in this firm than elsewhere, he has no incentive to quit. Similarly, because the firm is better off by employing the worker than by laying him off (that is, the worker gets paid less than his value of marginal product), the firm does not want to let the worker go. If *both* the firm and the worker share the returns of the specific training, therefore, the possibility of job separation in the posttraining period is eliminated.

If firms and workers do share the returns of specific training, they also will have to share the costs. After all, if firms paid all the costs of providing specific training and got only part of the returns, they would attract an oversupply of trainees. Therefore, if firms pay, say, 30 percent of the costs of specific training, they also will get 30 percent of the returns. Otherwise, the firm would attract either too few or too many job applicants.

Some Implications of Specific Training

It is important to note that specific training breaks the link between the worker's wage and the value of marginal product throughout the worker's life cycle. During the training period, workers get paid less than their value of marginal product because they are paying part of the training costs. In the posttraining period, workers get paid less than their value of marginal product in the firm that provided the training, but get paid more than their marginal product in other firms (that is, they get paid more than the alternative wage).

As a result of this contract, workers who have specific training are effectively granted a type of tenure or lifetime contract in the firm. Neither workers nor firms that have invested in specific training want to terminate the employment contract. It might seem surprising to argue that lifetime contracts might be common in labor markets where workers and firms are evidently very mobile, such as in the United States. Nevertheless, the evidence indicates that jobs lasting more than 20 years are the rule rather than the exception even in the United States.[38]

The concept of specific training has many other implications for labor markets. It provides a simple explanation of the "last hired, first fired" rule that typically determines who gets laid off during an economic downturn. Workers who have been with a firm for many years probably have more specific training than newly hired workers. When the demand for the firm's output falls, the price of the output and the value of the worker's marginal product decline. Workers with seniority have a buffer between their value of marginal product and their wage, so that the drop in the value of the worker's contribution to the firm protects these senior workers from layoffs. Put differently, because specifically trained workers produce more than they get paid, the firm need not lay off many of these workers when it experiences a sudden drop in the demand for its product. Profit-maximizing employers who want to cut the size of the workforce, therefore, will lay off newly hired workers.

Moreover, if a specifically trained worker does get laid off, he will have little incentive to find alternative employment. After all, these workers will suffer a capital loss if they change employers. Specifically trained workers, therefore, will prefer to "wait out" the unemployment spell until they are recalled by their former employers. There is, in fact, a very high incidence of **temporary layoffs** in many labor markets. At least 60 percent of the layoffs in the United States end when their former employers recall laid-off workers.[39]

Because specific training "marries" firms and workers, the probability of job separation for a given worker (either through a quit or a layoff) declines with job seniority. Newly hired workers will have high turnover rates, whereas more senior workers will have low turnover rates. This negative correlation between job turnover propensities and job seniority would not arise if all training were general. General training is portable and can be carried to any firm at any time. As a result, there would be no reason to expect the worker's economic opportunities in the current firm (relative to other firms) to improve over time. The important relationship between specific training and job turnover is discussed in more detail in Chapter 9.[40]

7-12 On-the-Job Training and the Age-Earnings Profile

The shape of the age-earnings profile depends on the timing of human capital investments over the working life.[41] At every age, we will want to invest in human capital up to the point where the marginal revenue of the investment equals the marginal cost of the investment. To describe

[38] Robert E. Hall, "The Importance of Lifetime Jobs in the U.S. Economy," *American Economic Review* 72 (September 1982): 716–24; see also Manuelita Ureta, "The Importance of Lifetime Jobs in the U.S. Economy, Revisited," *American Economic Review* 82 (March 1992): 322–35.

[39] Martin S. Feldstein, "Temporary Layoffs in the Theory of Unemployment," *Journal of Political Economy* 84 (October 1976): 937–57.

[40] Although the discussion assumes that the specific capital embodied in the worker is specific to the firm, it may well be the case that some of the specific capital is specific to the industry where the worker is employed; see Daniel Parent, "Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics," *Journal of Labor Economics* 18 (April 2000): 306–23.

[41] Yoram Ben-Porath, "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy* 75 (August 1967): 352–65; and James J. Heckman, "A Life-Cycle Model of Earnings, Learning, and Consumption," *Journal of Political Economy* 84 (August 1976 Supplement): S11–S46.

the timing of human capital acquisitions, therefore, we must describe what happens to the marginal revenue and the marginal cost of human capital investments as workers get older.

For convenience, let's measure the human capital stock in **efficiency units.** Efficiency units are standardized units of human capital. The total human capital stock of the worker equals the total number of efficiency units embodied in him or her. If David has 100 efficiency units and Mac has only 50 units, David is equivalent to two Macs—at least in terms of his labor market productivity.

An efficiency unit of human capital can be rented out in the labor market, and the rental rate per efficiency unit is R dollars. The market for efficiency units is competitive, so the per-unit rental price is R dollars regardless of how many efficiency units a worker has. Finally, to keep things simple, let's assume that all training is general and that there is no depreciation of the human capital stock over time. Therefore, an efficiency unit of human capital generates R dollars per year from the date when it is acquired until retirement, which occurs at age 65.

Suppose that the worker enters the labor market at age 20 and plans to retire at age 65. The marginal revenue of acquiring one efficiency unit of human capital at age 20 is

$$MR_{20} = R + \frac{R}{1+r} + \frac{R}{(1+r)^2} + \frac{R}{(1+r)^3} + \cdots + \frac{R}{(1+r)^{45}} \qquad \textbf{(7-25)}$$

where r is the discount rate. The intuition behind equation (7-25) is easy to understand. If a worker acquires one efficiency unit at age 20, this investment yields a return of R dollars during that first year in the labor market. In the second year, the present value of the return to that same efficiency unit is $R/(1 + r)$ dollars; in the third year, the return equals $R/(1 + r)^2$ dollars; and so on. Equation (7-25) simply adds the discounted returns to the efficiency unit over the entire working life.

The curve MR_{20} in Figure 7-9 illustrates the relationship between the marginal revenue of an efficiency unit acquired at age 20 and the number of efficiency units that the worker acquires. Because we assumed that the rental rate R is the same regardless of how much human capital the worker acquires, the marginal revenue curve MR_{20} is horizontal.

Suppose the worker looks into the future and wants to know how many efficiency units of human capital he would have acquired if he were 30 years old. The marginal revenue of an efficiency unit acquired at age 30 is given by

$$MR_{30} = R + \frac{R}{1+r} + \frac{R}{(1+r)^2} + \frac{R}{(1+r)^3} + \cdots + \frac{R}{(1+r)^{35}} \qquad \textbf{(7-26)}$$

Equation (7-26) indicates that the marginal revenue of acquiring an efficiency unit at age 30 is the discounted sum of the returns collected at age 30, at age 31, and so on. Note that the worker is now 10 years closer to retirement, so the sum in equation (7-26) has 10 fewer terms than the sum in equation (7-25).

By comparing the marginal revenue of acquiring an efficiency unit at ages 20 and 30, we can see that the marginal revenue of investing at age 20 exceeds the marginal revenue of investing at age 30. This fact is illustrated in Figure 7-9, which shows that the MR_{30} curve lies below the MR_{20} curve. The marginal revenue of human capital investment falls as the worker ages for a simple reason: We do not live forever. Human capital acquired when young can be rented out for a long period of time, whereas investments undertaken at

FIGURE 7-9
The marginal re
revenue of a uni
marginal cost, s

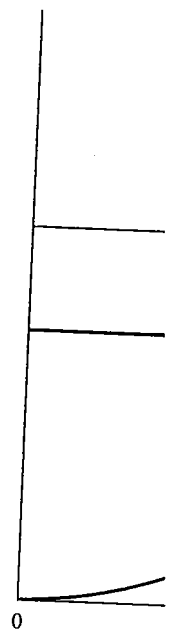

FIGURE 7-9 The Acquisition of Human Capital over the Life Cycle
The marginal revenue of an efficiency unit of human capital declines as the worker ages (so that MR_{20}, the marginal revenue of a unit acquired at age 20, lies above MR_{30}). At each age, the worker equates the marginal revenue with the marginal cost, so that more units are acquired when the worker is younger.

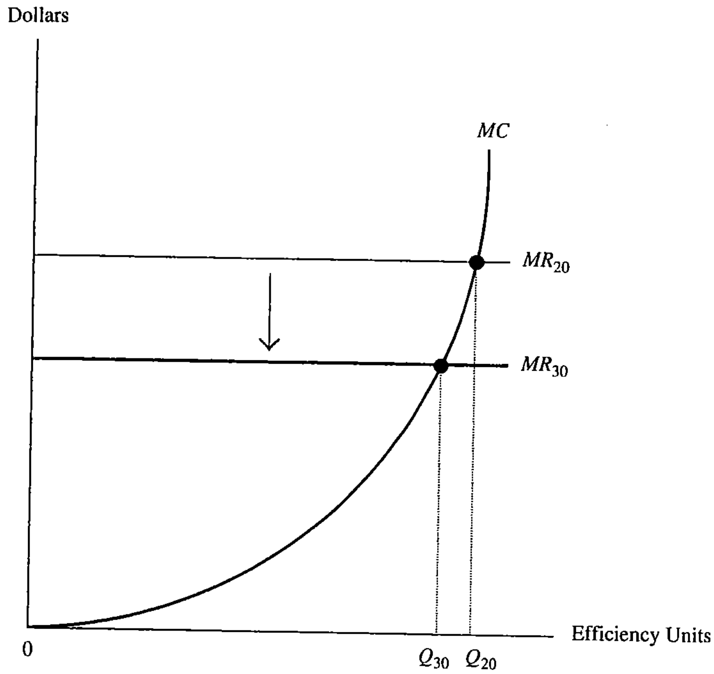

older ages can be rented out only for shorter periods. As a result, human capital investments are more profitable the earlier they are undertaken.

As noted earlier, the actual number of efficiency units acquired at any age is determined by equating the marginal revenue with the marginal cost of human capital investments. The marginal cost curve (*MC*), also illustrated in Figure 7-9, has the usual shape: Marginal costs rise as more efficiency units are acquired. The shape of the marginal cost curve is determined by the underlying production function for human capital. The assumption of diminishing returns in the production of efficiency units guarantees that marginal costs increase at an increasing rate as the worker attempts to acquire more and more human capital.

The intersection of MR_{20} and the marginal cost curve in Figure 7-9 implies that the worker will acquire Q_{20} efficiency units at age 20. Because the marginal revenue of human capital investments declines over time, the optimal investment level at age 30 falls to Q_{30}. In other words, the worker acquires fewer efficiency units as he gets older. This result helps us understand why workers typically go to school when young, why this period of complete specialization in human capital investments is followed by a period of considerable on-the-job training, and why on-the-job training activities taper off as the

FIGURE 7-10 The Age-Earnings Profile Implied by Human Capital Theory
The age-earnings profile is upward-sloping and concave. Older workers earn more because they invest less in human capital and because they are collecting the returns from earlier investments. The rate of growth of earnings slows down over time because workers accumulate less human capital as they get older.

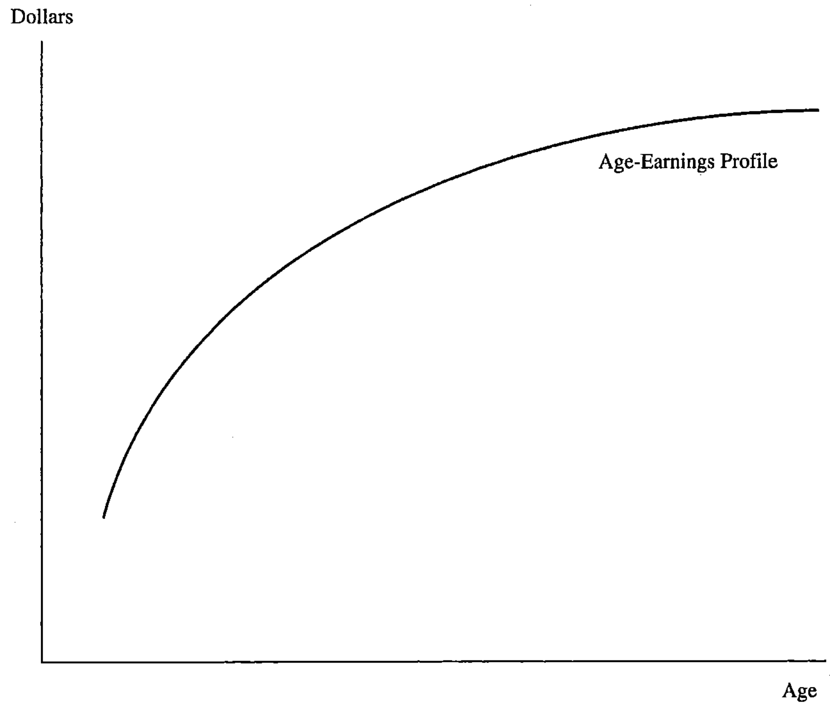

worker ages. This timing of investments over the life cycle maximizes the present value of lifetime earnings.[42]

Because the worker acquires more human capital when he is young, the worker's age-earnings profile is upward sloping, as illustrated in Figure 7-10. As we have seen, workers pay for on-the-job training through reduced wages. Older workers, therefore, earn more than younger workers because older workers acquire fewer efficiency units of human capital and, hence, have lower forgone earnings. Older workers also earn more because they are collecting the returns made on prior investments.

The optimal timing of investments over the working life also implies that the age-earnings profile is concave so that earnings increase over time but at a decreasing rate. Year-to-year

[42] Our discussion assumes that the marginal cost curve is constant over time (that is, it does not shift as the worker ages). It may be that older workers are more efficient at producing human capital, and, hence, the marginal cost curve would shift down. The forgone earnings incurred in producing human capital, however, are higher for older workers so that the marginal cost curve shifts up with age. It is sometimes assumed that these two opposing effects exactly outweigh each other (this assumption is called the "neutrality assumption"). As a result, the marginal cost curve does not shift over time. For a discussion of these issues, see Yoram Ben-Porath, "The Production of Human Capital over Time," in W. Lee Hansen, editor, *Education, Income, and Human Capital*, New York: Columbia University Press, 1970.

Theory at Work
EARNINGS AND SUBSTANCE ABUSE

Most of the examples that dominate discussions of the economic impact of human capital deal with investments that have a beneficial impact on the worker's productivity, such as schooling and on-the-job training. Many workers, however, also undertake activities that presumably have an adverse impact on the value of their human capital stock, such as alcoholism and drug use.

Alcoholism is a major social and economic problem in many countries. In the United States, this disorder afflicts about 5 percent of the population at any point in time, and nearly 10 percent of the population at some point in their lives. There is strong evidence that alcoholics pay a heavy price not only in terms of their health and the well-being of their families, but also in the labor market. Among workers aged 30 to 59, alcoholics are 15 percentage points less likely to work and earn 17 percent less than nonalcoholics, even if we look at alcoholics whose health has not yet been impaired.

Drug use is an equally important problem. By the time workers reach the age of 30, nearly 30 percent have used cocaine and about 3 percent have used it in the past month. Surprisingly, the evidence does not suggest that cocaine users have systematically lower wages or employment rates.

It is important to stress that these correlations between substance abuse and labor market characteristics do not necessarily prove that alcoholism "causes" lower wages or that cocaine use "does not reduce" productivity. The population of substance abusers is self-selected. Perhaps alcoholism does not reduce earnings, but those workers who are less successful in the labor market have a greater chance of becoming alcoholics. Similarly, it may be that only those workers who can handle the adverse consequences of cocaine use—or who can afford to buy cocaine—become habitual users.

The range of activities that may have adverse labor market consequences will likely increase as employers increasingly use Internet searches to gain information on the personal life of job applicants, often by searching through personal Web pages published by the applicants themselves. There are widespread reports that some employers now engage in this type of systematic search as a routine part of the hiring process, and that various forms of personal behavior are seen as "deal breakers" in the process.

Sources: Thomas S. Dee and William N. Evans, "Teen Drinking and Educational Attainment: Evidence from Two-Sample Instrumental Variables Estimates," *Journal of Labor Economics* 21 (January 2003): 178–209; John Muhally and Jody L. Sindelar, "Alcoholism, Work, and Income," *Journal of Labor Economics* 11 (July 1993): 494–520; and Robert Kaestner, "New Estimates of the Effect of Marijuana and Cocaine Use on Wages," *Industrial and Labor Relations Review* 47 (April 1994): 454–70.

wage growth depends partly on how many additional efficiency units the worker acquires. Because fewer units are acquired as the worker ages, the rate of wage growth declines over time.

The Mincer Earnings Function

The implications of the human capital model for the age-earnings profile have been the subject of extensive empirical analysis. This line of research culminated in the development of Jacob Mincer's human capital earnings function.[43] In particular, Mincer showed that the human capital model generates an age-earnings profile of the form

$$\log w = as + bt - ct^2 + \text{other variables} \qquad \textbf{(7-27)}$$

[43] Jacob Mincer, *Schooling, Experience, and Earnings*. New York: Columbia University Press, 1974. This literature is surveyed by Robert J. Willis, "Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions," in Orley C. Ashenfelter and Richard Layard, editors, *Handbook of Labor Economics*, vol. 1. Amsterdam: Elsevier, 1986, pp. 525–602; and James J. Heckman, Lance J. Lochner, and Petra E. Todd, "Fifty Years of Mincer Earnings Functions," National Bureau of Economic Research Working Paper No. 9732, May 2003.

where w is the worker's wage rate, s is the number of years of schooling, t gives the number of years of labor market experience, and t^2 is a quadratic on experience that captures the concavity of the age-earnings profile.

In the **Mincer earnings function,** as this widely used equation has come to be known, the coefficient on schooling a estimates the percent increase in earnings resulting from one additional year of schooling and is typically interpreted as the rate of return to schooling. We have seen that this interpretation is correct only when workers do not differ in their unobserved ability. The coefficients on experience and experience squared estimate the rate of growth in earnings resulting from one additional year of labor market experience and are typically interpreted as measuring the impact of on-the-job training on earnings. If the worker did not invest in OJT, the coefficients of the experience variables would be zero because there would be no reason for real earnings to increase with labor market experience.

Hundreds of studies have found that the Mincer earnings function provides a reasonably accurate description of age-earnings profiles not only in the United States, but also in the labor markets of many other countries (even in countries with very different labor market institutions). As we saw earlier in this chapter, actual age-earnings profiles in the United States are concave and are higher for more-educated workers. The evidence also suggests that differences in education and labor market experience among workers account for about a third of the variation in wage rates in the population. The human capital model, therefore, goes a long way toward providing a useful story of how the earnings distribution arises.

7-13 Policy Application: Evaluating Government Training Programs

Perhaps the most important policy implication of the human capital model is that the provision of training to low-skill workers may substantially improve their economic well-being. Since the declaration of the War on Poverty in the mid-1960s, a large number of government programs have indeed attempted to provide training to disadvantaged workers. These programs include the Manpower Development and Training Act of 1962 (MDTA), the Comprehensive Employment and Training Act of 1973 (CETA), and the Job Training Partnership Act of 1982 (JTPA).

Each of these programs spent a lot of money trying to "expose" minority and other low-income groups to formal training programs. Federal expenditures on job training programs now exceed over $4 billion per year. In view of the large cost of setting up, maintaining, and operating these programs, it is not surprising that a large number of studies attempt to determine if these programs do what they are supposed to do—namely, increase the human capital and earnings of the trainees.[44]

[44] This literature is surveyed in James J. Heckman, Robert J. LaLonde, and Jeffrey A. Smith, "The Economics and Econometrics of Active Labor Market Programs," in Orley C. Ashenfelter and David Card, editors, *Handbook of Labor Economics*, vol. 3A, Amsterdam: Elsevier, 1999, pp. 1865–2097.

The program evaluations have raised a number of still-unresolved conceptual issues. It would seem that by comparing the earnings of trainees before and after the "treatment," we could measure the effectiveness of the program (at least in terms of the program's impact on earnings capacity). A number of studies have made this type of before-and-after comparison and have found that there are some earnings gains associated with the training programs. Typically, trainees earn about $300 to $1,500 more per year after the program than before the program.[45]

Unfortunately, this calculation may not be very meaningful. As in many other contexts in labor economics, the problem of self-selection mars the analysis. In particular, only those workers who have the most to gain from the program and are most committed to "self-improvement" are likely to enroll and subject themselves to the treatment. The earnings gain achieved by this nonrandom sample of workers, therefore, tells us something about how the training programs affect motivated workers but may say nothing about how the program would affect a randomly chosen person in the disadvantaged population. From a policy point of view, therefore, the before-and-after calculation is useless because it cannot be used to predict how the earnings of targeted workers (such as persons currently receiving public assistance) would respond to the treatment.

Social Experiments

To avoid these pitfalls, there has been a revolutionary shift in the methodology used in program evaluation in recent years. The newer evaluations use randomized experiments, akin to the experimental methods used in the physical sciences, to estimate the impact of the program on trainee earnings. In these experiments, potential trainees are randomly assigned to participate in the training program. Every other applicant, for instance, is allocated to the "treatment" group (that is, they are exposed to the training program), whereas the remaining applicants form the control group and are administered a placebo (that is, they are not put through the training program).

The National Supported Work Demonstration (NSW) provides a good example of such a randomization scheme.[46] The key objective of the NSW was to ease the transition of disadvantaged workers into the labor market by exposing them to a work environment where experience and counseling could be provided. In this experiment, eligible applicants were assigned randomly to one of two tracks. The lucky workers who were treated by the program

[45] See, for example, Orley C. Ashenfelter, "Estimating the Effect of Training Programs on Earnings," *Review of Economics and Statistics* 60 (February 1978): 47–57; Orley C. Ashenfelter and David Card, "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs," *Review of Economics and Statistics* 67 (November 1985): 648–60; and Burt Barnow, "The Impact of CETA Programs on Earnings: A Review of the Literature," *Journal of Human Resources* 22 (Spring 1987): 157–93.

[46] Robert J. LaLonde, "Evaluating the Econometric Evaluations of Training Programs with Experimental Data," *American Economic Review* 76 (September 1986): 604–20. Other studies of experimental data include Stephen H. Bell and Larry L. Orr, "Is Subsidized Employment Cost Effective for Welfare Recipients? Experimental Evidence from Seven State Demonstrations," *Journal of Human Resources* 29 (Winter 1994): 42–61; and Alberto Abadie, Joshua D. Angrist, and Guido W. Imbens, "Instrumental Variables Estimates of the Effects of Training on the Quantiles of Trainee Earnings," *Econometrica* 70 (January 2002): 91–117.

TABLE 7-4 The Impact of the NSW Program on the Earnings of Trainees (in 1998 dollars)

Source: Robert J. LaLonde, "Evaluating the Econometric Evaluations of Training Programs with Experimental Data," *American Economic Review* 76 (September 1986): 604–20, Table 2.

Group	Pretraining Annual Earnings (1975)	Posttraining Annual Earnings (1979)	Difference
Treatment group	1,512	7,888	6,376
Control group	1,481	6,450	4,969
Difference-in-differences	—	—	**1,407**

received all the benefits provided by the NSW, whereas those assigned to the control group received none of the benefits and were left on their own. The NSW guaranteed persons in the treatment group a job for 9 to 18 months, at which time they had to find regular employment. The program cost about $12,500 per participant (in 1998 dollars).

It is easy to estimate the impact of the program on the worker's earnings capacity in the context of this experimental scheme. Table 7-4 summarizes the evidence from an influential evaluation of the NSW program. The typical worker who was treated by the program earned $1,512 annually in the pretraining period and $7,888 after the training. The typical trainee, therefore, experienced a wage gain of almost $6,400.

This wage gain, however, does *not* estimate the impact of the training program because the earnings of trainees could have changed between 1975 and 1979 for other reasons, such as aging and changes in aggregate economic conditions. To isolate the true impact of the NSW program, therefore, we must net out the impact of these extraneous events on earnings. It turns out that workers in the control group—those who were *not* exposed to the training activities provided by the program—earned $1,481 annually before the training and $6,450 after the training, so they experienced an earnings gain of almost $5,000. Since earnings would have increased by $5,000 regardless of whether the worker was injected with the training, the true impact of the training program is the difference-in-differences, or about $1,400.

As noted above, the NSW program cost about $12,500 per participant. Therefore, it would take longer than a decade (if future earnings gains are discounted) for the program to reach the "breakeven" point, the point at which the per-worker training costs equal the present value of the benefits accrued by the worker. Nevertheless, the rate of return to the investment is on the order of 10 percent.

Although the experimental approach is rapidly becoming the standard way of evaluating the impact of worker training programs, the methodology has its detractors.[47] These

[47] See James J. Heckman and V. Joseph Hotz, "Choosing among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American Statistical Association* 84 (December 1989): 862–74; and James J. Heckman and Jeffrey A. Smith, "Assessing the Case for Social Experiments," *Journal of Economic Perspectives* 9 (Spring 1995): 85–110.

Summary

detractors argue that it is incorrect to assume that the $1,400 increase in earnings is the net gain that would be observed if the program were made available to the entire disadvantaged population, and a randomly chosen person in that population were admitted into the program. The criticism is valid because the treatment and control groups do not truly represent a natural experiment. Only persons who are interested in receiving the training in the first place bother to go to the training center and fill out an application. As a result, there is already self-selection in the sample of persons who end up in the treatment group. Moreover, some persons allocated to the treatment group may not show up for the training, whereas persons allocated to the control group may find a way of qualifying for some type of training program (perhaps by trying out at other training sites). Experimental methods, therefore, may not entirely get rid of the selection bias that is at the heart of the evaluation problem.

Summary

- A dollar received today does not have the same value as a dollar received tomorrow. The present value of a future income receipt gives the value of that amount in terms of today's dollars.
- The wage-schooling locus gives the salary that a worker earns if he or she completes a particular level of schooling.
- Workers choose the point on the wage-schooling locus that maximizes the present value of lifetime earnings. In particular, workers quit school when the marginal rate of return to schooling equals the rate of discount.
- When workers differ only in their discount rates, the rate of return to schooling can be estimated by comparing the earnings of different workers. When workers differ in their innate abilities, the wage differential among workers does not measure the rate of return to schooling because the wage gap also depends on the unobserved ability differential.
- Workers sort themselves into those occupations for which they are best suited. This self-selection implies that we cannot test the hypothesis that workers choose the schooling level that maximizes the present value of lifetime earnings by comparing the earnings of different workers.
- In the United States, the rate of return to schooling was around 9 percent in the 1990s.
- Schooling can play a signaling role in the labor market, indicating to employers that the worker carrying the certificate or diploma is a highly productive worker. The signaling value of education can help firms differentiate highly productive workers from less productive workers.
- If education plays only a signaling role, workers with more schooling earn more not because education increases productivity, but because education signals a worker's innate ability.

- The observed age-earnings profile is upward sloping and concave. Earnings increase over the life cycle, but at a decreasing rate.
- General training is valuable in all firms. Specific training is valuable only in the firm that provides the training. Workers pay for and collect the returns from general training. Workers and firms share both the costs and the returns of specific training.
- The optimal timing of human capital investments over the life cycle implies that the age-earnings profile is upward-sloping and concave.

Key Concepts

age-earnings profile, *247*
efficiency units, *280*
general training, *276*
human capital, *243*
Mincer earnings function, *284*
opportunity cost, *247*
present value, *246*
private rate of return to schooling, *273*
rate of discount, *246*
rate of return to schooling, *251*
selection bias, *267*
signal, *270*
social rate of return to schooling, *273*
specific training, *276*
temporary layoffs, *279*
wage-schooling locus, *250*

Review Questions

1. Discuss how the present value of a future income payment is calculated.
2. Discuss how the wage-schooling locus is determined in the labor market, and why it is upward sloping and concave.
3. Derive the stopping rule for investments in education.
4. Why does the percentage gain in earnings observed when a worker gets one more year of schooling measure the marginal rate of return to education?
5. Discuss how differences in discount rates or in ability across workers lead to differences in earnings and schooling. Under what conditions can the rate of return to schooling be estimated?
6. Discuss the relationship between ability bias in the estimation of the rate of return to schooling and selection bias in tests of the hypothesis that workers choose the level of schooling that maximizes the present value of earnings.
7. Discuss how empirical studies estimate the rate of return to schooling and the methods used to avoid the problem of ability bias.
8. Show how education can signal the worker's innate ability in the labor market. What is a pooled equilibrium? What is a perfectly separating signaling equilibrium?
9. How can we differentiate between the hypothesis that education increases productivity and the hypothesis that education is a signal for the worker's innate ability?
10. Discuss the difference between general training and specific training. Who pays for and collects the returns from each type of training?

Problems

11. Discuss the implications of general and specific training for the worker's age-earnings profile.
12. Why are experimental methods now commonly used to evaluate the impact of training programs? Discuss how and under what conditions we can use the results of an experiment to estimate the rate of return to the program.

Problems

7-1. Debbie is about to choose a career path. She has narrowed her options to two alternatives. She can become either a marine biologist or a concert pianist. Debbie lives two periods. In the first, she gets an education. In the second, she works in the labor market. If Debbie becomes a marine biologist, she will spend $15,000 on education in the first period and earn $472,000 in the second period. If she becomes a concert pianist, she will spend $40,000 on education in the first period and then earn $500,000 in the second period.

a. Suppose Debbie can lend and borrow money at a 5 percent annual rate. Which career will she pursue? What if she can lend and borrow money at a 15 percent rate of interest? Will she choose a different option? Why?

b. Suppose musical conservatories raise their tuition so that it now costs Debbie $60,000 to become a concert pianist. What career will Debbie pursue if the discount rate is 5 percent?

7-2. Peter lives for three periods. He is currently considering three alternative education-work options. He can start working immediately, earning $100,000 in period 1, $110,000 in period 2 (as his work experience leads to higher productivity), and $90,000 in period 3 (as his skills become obsolete and his physical abilities deteriorate). Alternatively, he can spend $50,000 to attend college in period 1 and then earn $180,000 in periods 2 and 3. Finally, he can receive a doctorate degree in period 2 after completing his college education in period 1. This last option will cost him nothing when he is attending graduate school in the second period as his expenses on tuition and books will be covered by a research assistantship. After receiving his doctorate, he will become a professor in a business school and earn $400,000 in period 3. Peter's discount rate is 20 percent per period. What education path maximizes Peter's net present value of his lifetime earnings?

7-3. Jane has three years of college, Pam has two, and Mary has one. Jane earns $21 per hour, Pam earns $19, and Mary earns $16. The difference in educational attainment is due completely to different discount rates. How much can the available information reveal about each woman's discount rate?

7-4. Suppose the skills acquired in school depreciate over time, perhaps because technological change makes the things learned in school obsolete. What happens to a worker's optimal amount of schooling if the rate of depreciation increases?

7-5. Suppose workers differ in their ability, but have the same discount rate. Is it possible for the more able workers to choose less schooling?

7-6. Suppose Carl's wage-schooling locus is given by

Years of Schooling	Earnings
9	\$18,500
10	\$20,350
11	\$22,000
12	\$23,100
13	\$23,900
14	\$24,000

Derive the marginal rate of return schedule. When will Carl quit school if his discount rate is 4 percent? What if the discount rate is 12 percent?

7-7. Table 217 of the 2006 *U.S. Statistical Abstract* shows that, among all 25–34 year olds, the average annual earnings of a high school graduate with no further education was \$26,073 while the average annual earnings of a college graduate with no further education was \$43,794 in 2003.

a. Assuming college requires five years, show that the annual return to each year of college education averages 10.9 percent.

b. It is typically thought that this type of calculation of the returns to schooling is biased, because it doesn't take into account innate ability (i.e., ability in the workplace not due to college) or innate motivation. If this criticism is true, is the actual return to each year of a college education more than or less than 10.9 percent?

7-8. Suppose there are two types of persons: high-ability and low-ability. A particular diploma costs a high-ability person \$8,000 and costs a low-ability person \$20,000. Firms wish to use education as a screening device where they intend to pay \$25,000 to workers without a diploma and \$*K* to those with a diploma. In what range must *K* be to make this an effective screening device?

7-9. It has been argued that the minimum wage prevents workers from investing in on-the-job training and discourages employers from providing specific training to low-income workers. Why would the minimum wage have an adverse effect on human capital accumulation for low-income workers?

7-10. Jill is planning the timing of her on-the-job training investments over the life cycle. What happens to Jill's OJT investments at every age if

a. The market-determined rental rate to an efficiency unit falls?

b. Jill's discount rate increases?

c. The government passes legislation delaying the retirement age until age 70?

d. Technological progress is such that much of the OJT acquired at any given age becomes obsolete within the next 10 years?

Selected Readings

7-11. Suppose 3 million high school graduates start college each year. Those who earn a college degree will do so in four years. However, some students will drop out along the way. The first-year attrition rate is 20 percent, while the second- and third-year attrition rates are 10 and 2.5 percent respectively.

a. What is the distribution of college students by year in college? How many students graduate from college each year?

b. Believing that education is the key to the future, a presidential candidate proposes that the federal government pay the first $3,000 of college expenses each year for everyone attending a four-year college. It is expected that this proposal will encourage 1 million more high school graduates to enroll in a four-year college each year. Of these 1 million new college students, the first-, second-, and third-year attrition rates are 40, 20, and 5 percent. Why is it likely that attrition rates will be higher among these groups of students?

c. What is the yearly projected cost of the program in part (b)? What is the average cost of each new four-year college graduate?

7-12. In 1970, men aged 18 to 25 were subject to the military draft to serve in the Vietnam War. A man could qualify for a student deferment, however, if he was enrolled in college and made satisfactory progress on obtaining a degree. By 1975, the draft was no longer in existence. The draft did not pertain to women. Using the data in Table 263 of the 2006 edition of the *U.S. Statistical Abstract*, use women as the control group to estimate (using the difference-in-differences methodology) the effect abolishing the draft had on male college enrollment.

Selected Readings

Joshua Angrist and Alan B. Krueger, "Does Compulsory Schooling Affect Schooling and Earnings?" *Quarterly Journal of Economics* 106 (November 1991): 979–1014.

Orley C. Ashenfelter and Alan B. Krueger, "Estimates of the Economic Return to Schooling from a New Sample of Twins," *American Economic Review* 84 (December 1994): 1157–73.

Yoram Ben-Porath, "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy* 75 (August 1967): 352–65.

David Card, "The Causal Effect of Education on Earnings," in Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, vol. 3A. Amsterdam: Elsevier, 1999, pp. 1801–63.

David Card and Alan B. Krueger, "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy* 100 (February 1992): 1–40.

James J. Heckman, Lance J. Lochner, and Petra E. Todd, "Fifty Years of Mincer Earnings Functions," National Bureau of Economic Research Working Paper No. 9732, May 2003.

A. Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87 (August 1973): 355–74.

Robert J. Willis and Sherwin Rosen, "Education and Self-Selection," *Journal of Political Economy* 87 (October 1979 Supplement): S7–S36.

Web Links

The U.S. Department of Education's Web site provides an essential introduction to many of the education programs in the United States: www.ed.gov/index.jsp.

The American Council on Education gives useful information to prospective takers of the General Education Development (GED) exam: www.acenet.edu/calec/ged.

Chap