



# 17

## Markets with Asymmetric Information

For most of this book, we have assumed that consumers and producers have complete information about the economic variables that are relevant for the choices they face. Now we will see what happens when some parties know more than others—i.e., when there is **asymmetric information**.

Asymmetric information is quite common. Frequently, a seller of a product knows more about its quality than the buyer does. Workers usually know their own skills and abilities better than employers. And business managers know more about their firms' costs, competitive positions, and investment opportunities than do the firms' owners.

Asymmetric information also explains many institutional arrangements in our society. It is one reason why automobile companies offer warranties on parts and service for new cars; why firms and employees sign contracts that include incentives and rewards; and why the shareholders of corporations must monitor the behavior of managers.

We begin by examining a situation in which the sellers of a product have better information about its quality than buyers have. We will see how this kind of asymmetric information can lead to market failure. In the second section, we see how sellers can avoid some of the problems associated with asymmetric information by giving potential buyers signals about the quality of their product. Product warranties provide a type of insurance that can be helpful when buyers have less information than sellers. But as the third section shows, the purchase of insurance entails difficulties of its own when buyers have better information than sellers.

In the fourth section, we show that managers may pursue goals other than profit maximization when it is costly for owners of private corporations to monitor their behavior. In other words, managers have better information than owners. We also show how firms can give managers an incentive to maximize profits even when monitoring their behavior is costly. Finally, we show that labor markets may operate inefficiently when employees have better information about their productivity than employers have.

### 17.1 QUALITY UNCERTAINTY AND THE MARKET FOR LEMONS

Suppose you bought a new car for \$20,000, drove it 100 miles, and then decided you really didn't want it. There was nothing wrong with the car—it performed beautifully and met all your expectations. You

#### CHAPTER OUTLINE

- 17.1 Quality Uncertainty and the Market for Lemons 617
- 17.2 Market Signaling 623
- 17.3 Moral Hazard 628
- 17.4 The Principal-Agent Problem 630
- \*17.5 Managerial Incentives in an Integrated Firm 636
- 17.6 Asymmetric Information in Labor Markets: Efficiency Wage Theory 639

#### LIST OF EXAMPLES

- 17.1 Lemons in Major League Baseball 622
- 17.2 Working into the Night 627
- 17.3 Reducing Moral Hazard: Warranties of Animal Health 630
- 17.4 CEO Salaries 632
- 17.5 Managers of Nonprofit Hospitals as Agents 634
- 17.6 Efficiency Wages at Ford Motor Company 641

**• asymmetric information**

Situation in which a buyer and a seller possess different information about a transaction.

simply felt that you could do just as well without it and would be better off saving the money for other things. So you decide to sell the car. How much should you expect to get for it? Probably not more than \$16,000—even though the car is brand new, has been driven only 100 miles, and has a warranty that is transferable to a new owner. And if you were a prospective buyer, you probably wouldn't pay much more than \$16,000 yourself.

Why does the mere fact that the car is second-hand reduce its value so much? To answer this question, think about your own concerns as a prospective buyer. Why, you would wonder, is this car for sale? Did the owner really change his or her mind about the car just like that, or is there something wrong with it? Is this car a “lemon”?

Used cars sell for much less than new cars because *there is asymmetric information about their quality*: The seller of a used car knows much more about the car than the prospective buyer does. The buyer can hire a mechanic to check the car, but the seller has had experience with it and will know more about it. Furthermore, the very fact that the car is for sale indicates that it may be a “lemon”—why sell a reliable car? As a result, the prospective buyer of a used car will always be suspicious of its quality—and with good reason.

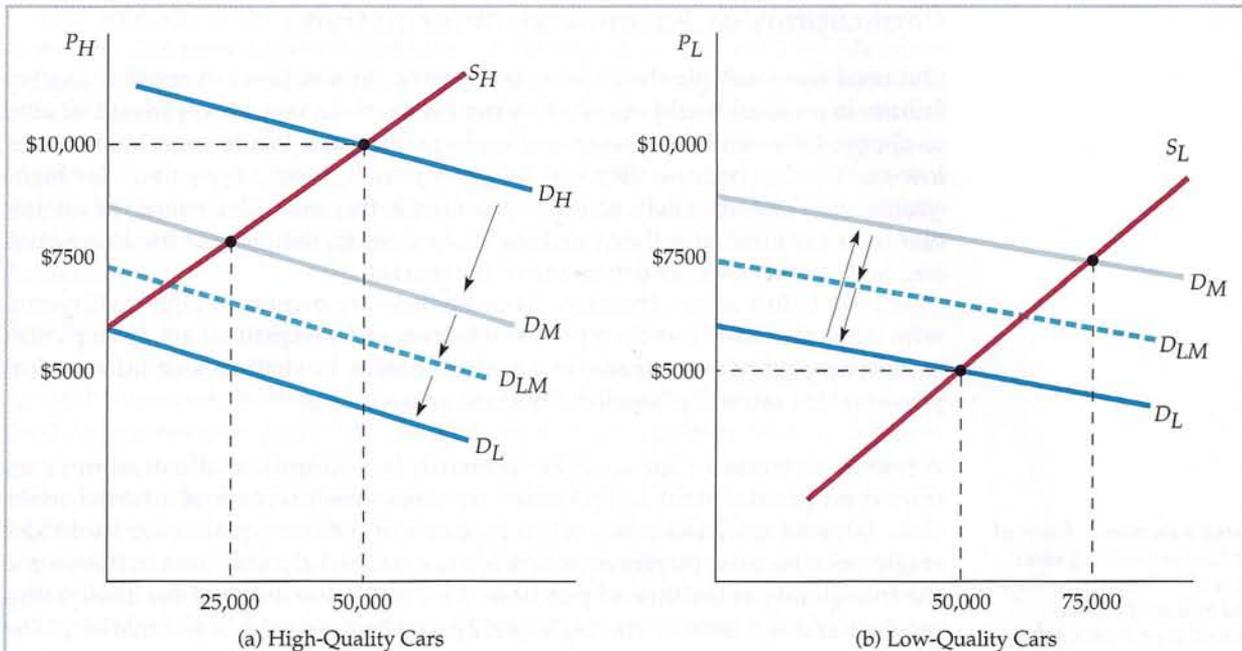
The implications of asymmetric information about product quality were first analyzed by *George Akerlof* and go far beyond the market for used cars.<sup>1</sup> The markets for insurance, financial credit, and even employment are also characterized by asymmetric information about product quality. To understand the implications of asymmetric information, we will start with the market for used cars and then see how the same principles apply to other markets.

## The Market for Used Cars

Suppose two kinds of used cars are available—high-quality cars and low-quality cars. Also suppose that both sellers and buyers can tell which kind of car is which. There will then be two markets, as illustrated in Figure 17.1. In part (a),  $S_H$  is the supply curve for high-quality cars, and  $D_H$  is the demand curve. Similarly,  $S_L$  and  $D_L$  in part (b) are the supply and demand curves for low-quality cars. For any given price,  $S_H$  lies to the left of  $S_L$  because owners of high-quality cars are more reluctant to part with them and must receive a higher price to do so. Similarly,  $D_H$  is higher than  $D_L$  because buyers are willing to pay more to get a high-quality car. As the figure shows, the market price for high-quality cars is \$10,000, for low-quality cars \$5000, and 50,000 cars of each type are sold.

In reality, the seller of a used car knows much more about its quality than a buyer does. (Buyers discover the quality only after they buy a car and drive it for a while.) Consider what happens, then, if sellers know the quality of cars, but buyers do not. Initially, buyers might think that the odds are 50-50 that a car will be high quality. Why? Because when both sellers and buyers know the quality, 50,000 cars of each type are sold. When making a purchase, buyers therefore view all cars as “medium quality,” in the sense that there is an equal chance of getting a high-quality or a low-quality car. (Of course, after buying the car and driving it for a while, they will learn its true quality.) The demand for cars perceived to be medium quality, denoted by  $D_M$  in Figure 17.1, is below  $D_H$  but above  $D_L$ . As the figure shows, these medium-quality cars will sell for about \$7500 each. However, *fewer high-quality cars (25,000) and more low-quality cars (75,000) will now be sold.*

<sup>1</sup>George A. Akerlof, “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics* (August 1970): 488–500.



**FIGURE 17.1** The Market for Used Cars

When sellers of products have better information about product quality than buyers, a “lemons problem” may arise in which low-quality goods drive out high quality goods. In (a) the demand curve for high-quality cars is  $D_H$ . However, as buyers lower their expectations about the average quality of cars on the market, their perceived demand shifts to  $D_M$ . Likewise, in (b) the perceived demand curve for low-quality cars shifts from  $D_L$  to  $D_M$ . As a result, the quantity of high-quality cars sold falls from 50,000 to 25,000, and the quantity of low-quality cars sold increases from 50,000 to 75,000. Eventually, only low quality cars are sold.

As consumers begin to realize that most cars sold (about three-fourths of the total) are low quality, their perceived demand shifts. As Figure 17.1 shows, the new perceived demand curve might be  $D_{LM}$ , which means that, on average, cars are thought to be of low to medium quality. However, the mix of cars then shifts even more heavily to low quality. As a result, the perceived demand curve shifts further to the left, pushing the mix of cars even further toward low quality. *This shifting continues until only low-quality cars are sold.* At that point, the market price would be too low to bring forth any high-quality cars for sale, so consumers correctly assume that any car they buy will be low quality. As a result, the only relevant demand curve will be  $D_L$ .

The situation in Figure 17.1 is extreme. The market may come into equilibrium at a price that brings forth at least some high-quality cars. *But the fraction of high-quality cars will be smaller than it would be if consumers could identify quality before making the purchase.* That is why you should expect to sell your brand new car, which *you know* is in perfect condition, for much less than you paid for it. Because of asymmetric information, low-quality goods drive high-quality goods out of the market. This phenomenon, which is sometimes referred to as the *lemons problem*, is an important source of market failure. It is worth emphasizing:

The lemons problem: With asymmetric information, low-quality goods can drive high-quality goods out of the market.



## Implications of Asymmetric Information

Our used cars example shows how asymmetric information can result in market failure. In an ideal world of fully functioning markets, consumers would be able to choose between low-quality and high-quality cars. While some will choose low-quality cars because they cost less, others will prefer to pay more for high-quality cars. Unfortunately, consumers cannot in fact easily determine the quality of a used car until after they purchase it. As a result, the price of used cars falls, and high-quality cars are driven out of the market.

Market failure arises, therefore, because there are owners of high-quality cars who value their cars less than potential buyers of high-quality cars. Both parties could enjoy gains from trade, but, unfortunately, buyers' lack of information prevents this mutually beneficial trade from occurring.

**Adverse Selection** Our used car scenario is a simplified illustration of an important problem that affects many markets—the problem of adverse selection. **Adverse selection** arises when products of different qualities are sold at a single price because buyers or sellers are not sufficiently informed to determine the true quality at the time of purchase. As a result, too much of the low-quality product and too little of the high-quality product are sold in the marketplace. Let's look at some other examples of asymmetric information and adverse selection. In doing so, we will also see how the government or private firms might respond to the problem.

• **adverse selection** Form of market failure resulting when products of different qualities are sold at a single price because of asymmetric information, so that too much of the low-quality product and too little of the high-quality product are sold.

**The Market for Insurance** Why do people over age 65 have difficulty buying medical insurance at almost any price? Older people do have a much higher risk of serious illness, but why doesn't the price of insurance rise to reflect that higher risk? Again, the reason is asymmetric information. People who buy insurance know much more about their general health than any insurance company can hope to know, even if it insists on a medical examination. As a result, adverse selection arises, much as it does in the market for used cars. Because unhealthy people are more likely to want insurance, the proportion of unhealthy people in the pool of insured people increases. This forces the price of insurance to rise, so that more healthy people, aware of their low risks, elect not to be insured. This further increases the proportion of unhealthy people among the insured, thus forcing the price of insurance up more. The process continues until most people who want to buy insurance are unhealthy. At that point, insurance becomes very expensive, or—in the extreme—insurance companies stop selling the insurance.

Adverse selection can make the operation of insurance markets problematic in other ways. Suppose an insurance company wants to offer a policy for a particular event, such as an auto accident that results in property damage. It selects a target population—say, men under age 25—to whom it plans to market this policy, and it estimates that the probability of an accident for people in this group is .01. However, for some of these people, the probability of having an accident is much less than .01; for others, it is much higher than .01. If the insurance company cannot distinguish between high- and low-risk men, it will base the premium on the average accident probability of .01. What will happen? Those people with low probabilities of having an accident will choose not to insure, while those with high probabilities of an accident will purchase the insurance. This in turn raises the accident probability among those who choose to be insured above .01, forcing the insurance company to raise its premium. In the extreme, only those who are likely to be in an accident will choose to insure, making it impractical to sell insurance.



One solution to the problem of adverse selection is to *pool risks*. For health insurance, the government might take on this role, as it does with the Medicare program. By providing insurance for *all* people over age 65, the government eliminates the problem of adverse selection. Likewise, insurance companies will try to avoid or at least reduce the adverse selection problem by offering group health insurance policies at places of employment. By covering all workers in a firm, whether healthy or sick, the insurance company spreads risks and thereby reduces the likelihood that large numbers of high-risk individuals will purchase insurance.<sup>2</sup>

**The Market for Credit** By using a credit card, many of us borrow money without providing any collateral. Most credit cards allow the holder to run a debt of several thousand dollars, and many people hold several credit cards. Credit card companies earn money by charging interest on the debit balance. But how can a credit card company or bank distinguish high-quality borrowers (who pay their debts) from low-quality borrowers (who don't)? Clearly, borrowers have better information—i.e., they know more about whether they will pay than the lender does. Again, the lemons problem arises. Low-quality borrowers are more likely than high-quality borrowers to want credit, which forces the interest rate up, which increases the number of low-quality borrowers, which forces the interest rate up further, and so on.

In fact, credit card companies and banks *can*, to some extent, use computerized credit histories, which they often share with one another, to distinguish low-quality from high-quality borrowers. Many people, however, think that computerized credit histories invade their privacy. Should companies be allowed to keep these credit histories and share them with other lenders? We can't answer this question for you, but we can point out that credit histories perform an important function: They eliminate, or at least greatly reduce, the problem of asymmetric information and adverse selection—a problem that might otherwise prevent credit markets from operating. Without these histories, even the creditworthy would find it extremely costly to borrow money.

## The Importance of Reputation and Standardization

Asymmetric information is also present in many other markets. Here are just a few examples:

- **Retail stores:** Will the store repair or allow you to return a defective product? The store knows more about its policy than you do.
- **Dealers of rare stamps, coins, books, and paintings:** Are the items real or counterfeit? The dealer knows much more about their authenticity than you do.
- **Roofers, plumbers, and electricians:** When a roofer repairs or renovates the roof of your house, do you climb up to check the quality of the work?
- **Restaurants:** How often do you go into the kitchen to check if the chef is using fresh ingredients and obeying health laws?

<sup>2</sup>Some people argue that pooling risks is not the main justification for Medicare, because most people's medical histories are well established by age 65, making it feasible for insurance companies to distinguish among high-risk and low-risk individuals. Another justification for Medicare is a distributional one. After age 65, even relatively healthy people are likely to need more medical care, making insurance expensive even without asymmetric information, and many older people would not have sufficient income to purchase the insurance.



In all these cases, the seller knows much more about the quality of the product than the buyer does. Unless sellers can provide information about quality to buyers, low-quality goods and services will drive out high-quality ones, and there will be market failure. Sellers of high-quality goods and services, therefore, have a big incentive to convince consumers that their quality is indeed high. In the examples cited above, this task is performed largely by *reputation*. You shop at a particular store because it has a reputation for servicing its products; you hire particular roofers or plumbers because they have reputations for doing good work; you go to a particular restaurant because it has a reputation for using fresh ingredients and nobody you know has become sick after eating there.

Sometimes, however, it is impossible for a business to develop a reputation. For example, because most of the customers of highway diners or motels go there only once or infrequently, the businesses have no opportunity to develop reputations. How, then, can they deal with the lemons problem? One way is *standardization*. In your hometown, you may not prefer to eat regularly at McDonald's. But a McDonald's may look more attractive when you are driving along a highway and want to stop for lunch. Why? Because McDonald's provides a standardized product: The same ingredients are used and the same food is served in every McDonald's anywhere in the country. Who knows? Joe's Diner might serve better food, but at least you *know* exactly what you will be buying at McDonald's.

#### EXAMPLE 17.1

#### Lemons in Major League Baseball



How can we test for the presence of a lemons market? One way is to compare the performance of products that are resold with similar products that are seldom put up for resale. In a lemons market, because purchasers of second-hand products will have limited information, resold products should be lower in quality than products that rarely appear on the market. One such

“second-hand” market was created some time ago by a change in the rules governing contracts in major league baseball.<sup>3</sup>

Before 1976, major league baseball teams had the exclusive right to renew a player's contract. After a 1976 ruling declared this system illegal, a new contracting arrangement was created. After six years of major league service, players can now sign new contracts with their original teams or become free agents and sign with new teams. The availability of many free agents creates a second-hand market in baseball players.

Asymmetric information is prominent in the free-agent market. One potential purchaser, the player's original team, has better information about the player's abilities than other teams have. If we were looking at used cars, we could test for the existence of asymmetric information by comparing their repair records. In baseball, we can compare player disability records. If players are working hard and following rigorous conditioning programs, we would expect a low probability of

<sup>3</sup>This example is based on Kenneth Lehn's study of the free-agent market. See “Information Asymmetries in Baseball's Free-Agent Market,” *Economic Inquiry* (1984): 37–44.



injury and a high probability that they will be able to perform if injured. In other words, more motivated players will spend less time on the bench owing to disabilities. If a lemons market exists, we would expect free agents to have higher disability rates than players who are renewed. Players may also have preexisting physical conditions which their original teams know about and which make them less desirable candidates for contract renewal. Because more such players would become free agents, free agents would experience higher disability rates for health reasons.

Table 17.1, which lists the post-contract performance of all players who have signed multiyear contracts, makes two points. First, both free agents and renewed players have increased disability rates after signing new contracts. The disabled days per season increase from an average of 4.73 to an average of 12.55. Second, the postcontract disability rates of renewed and non-renewed players are significantly different. On average, renewed players are disabled for 9.68 days, free agents for 17.23 days.

These two findings suggest that there is a lemons market in free agents that exists because baseball teams know their own players better than the teams with which they compete.

**TABLE 17.1** Player Disability

	Days Spent on Disabled List per Season		
	Precontract	Postcontract	Percentage Change
All players	4.73	12.55	165.4
Renewed players	4.76	9.68	103.4
Free agents	4.67	17.23	268.9

## 17.2 MARKET SIGNALING

We have seen that asymmetric information can sometimes lead to a lemons problem: Because sellers know more about the quality of a good than buyers do, buyers may assume that quality is low, causing price to fall and only low-quality goods to be sold. We also saw how government intervention (in the market for health insurance, for example) or the development of a reputation (in service industries, for example) can alleviate this problem. Now we will examine another important mechanism through which sellers and buyers deal with the problem of asymmetric information: **market signaling**. The concept of market signaling was first developed by Michael Spence, who showed that in some markets, sellers send buyers *signals* that convey information about a product's quality.<sup>4</sup>

To see how market signaling works, let's look at a *labor market*, which is a good example of a market with asymmetric information. Suppose a firm is thinking about hiring some new people. The new workers (the "sellers" of labor) know much more about the quality of the labor they can provide than does the firm (the buyer of labor). For example, they know how hard they tend to work, how

• **market signaling** Process by which sellers send signals to buyers conveying information about product quality.

<sup>4</sup>Michael Spence, *Market Signaling* (Cambridge, MA: Harvard University Press, 1974).



responsible they are, what their skills are, and so forth. The firm will learn these things only after workers have been hired and have been working for some time.

Why don't firms simply hire workers, see how well they work, and then fire those with low productivity? Because this policy is often very costly. In many countries, and in many firms in the United States, it is difficult to fire someone who has been working more than a few months. (The firm may have to show just cause or provide severance pay.) Moreover, in many jobs, workers do not become fully productive for at least six months. Before that time, considerable on-the-job training may be required, for which the firm must invest substantial resources. Thus the firm might not learn how good workers are for six months to a year. Clearly, firms would be much better off if they knew how productive potential employees were *before* they hired them.

What characteristics can a firm examine to obtain information about people's productivity before it hires them? Can potential employees convey information about their productivity? Dressing well for the job interview might convey some information, but even unproductive people can dress well. Dressing well is thus a *weak signal*—it doesn't do much to distinguish high-productivity from low-productivity people. *To be strong, a signal must be easier for high-productivity people to give than for low-productivity people to give, so that high-productivity people are more likely to give it.*

For example, *education* is a strong signal in labor markets. A person's educational level can be measured by several things—the number of years of schooling, degrees obtained, the reputation of the university or college that granted the degrees, the person's grade-point average, and so on. Of course, education can directly and indirectly improve a person's productivity by providing information, skills, and general knowledge that are helpful in work. But even if education did *not* improve productivity, it would still be a useful *signal* of productivity because more productive people find it easier to attain high levels of education. Not surprisingly, productive people tend to be more intelligent, more motivated, more disciplined, and more energetic and hard-working—characteristics that are also helpful in school. More productive people are therefore more likely to attain high levels of education *in order to signal their productivity to firms and thereby obtain better-paying jobs.* Thus, firms are correct in considering education a signal of productivity.

## A Simple Model of Job Market Signaling

To understand how signaling works, we will discuss a simple model.<sup>5</sup> Let's assume that there are only low-productivity workers (Group I), whose average and marginal product is 1, and high-productivity workers (Group II), whose average and marginal product is 2. Workers will be employed by competitive firms whose products sell for \$10,000, and who expect an average of 10 years' work from each employee. We also assume that half the workers in the population are in Group I and the other half in Group II, so that the *average* productivity of all workers is 1.5. Note that the revenue expected to be generated from Group I workers is \$100,000 ( $\$10,000/\text{year} \times 10 \text{ years}$ ) and from Group II workers is \$200,000 ( $\$20,000/\text{year} \times 10 \text{ years}$ ).

If firms could identify people by their productivity, they would offer them a wage equal to their marginal revenue product. Group I people would be paid \$10,000 per year, Group II people \$20,000. On the other hand, if firms could not identify productivity before they hired people, they would pay all workers an

<sup>5</sup>This is essentially the model developed in Spence, *Market Signaling*.



annual wage equal to the average productivity—\$15,000. Group I people would then earn more (\$15,000 instead of \$10,000), at the expense of Group II people (who would earn \$15,000 instead of \$20,000).

Now let's consider what can happen with signaling via education. Suppose all the attributes of an education (degrees earned, grade-point average, etc.) can be summarized by a single index  $y$  that represents years of higher education. All education involves a cost, and the higher the educational level  $y$ , the higher the cost. This cost includes tuition and books, the opportunity cost of foregone wages, and the psychic cost of having to work hard to obtain high grades. What is important is that *the cost of education is greater for the low-productivity group than for the high-productivity group*. We might expect this to be the case for two reasons. First, low-productivity workers may simply be less studious. Second, low-productivity workers may progress more slowly through degree programs. In particular, suppose that for Group I people, the cost of attaining educational level  $y$  is given by

$$C_I(y) = \$40,000y$$

and that for Group II people, it is

$$C_{II}(y) = \$20,000y$$

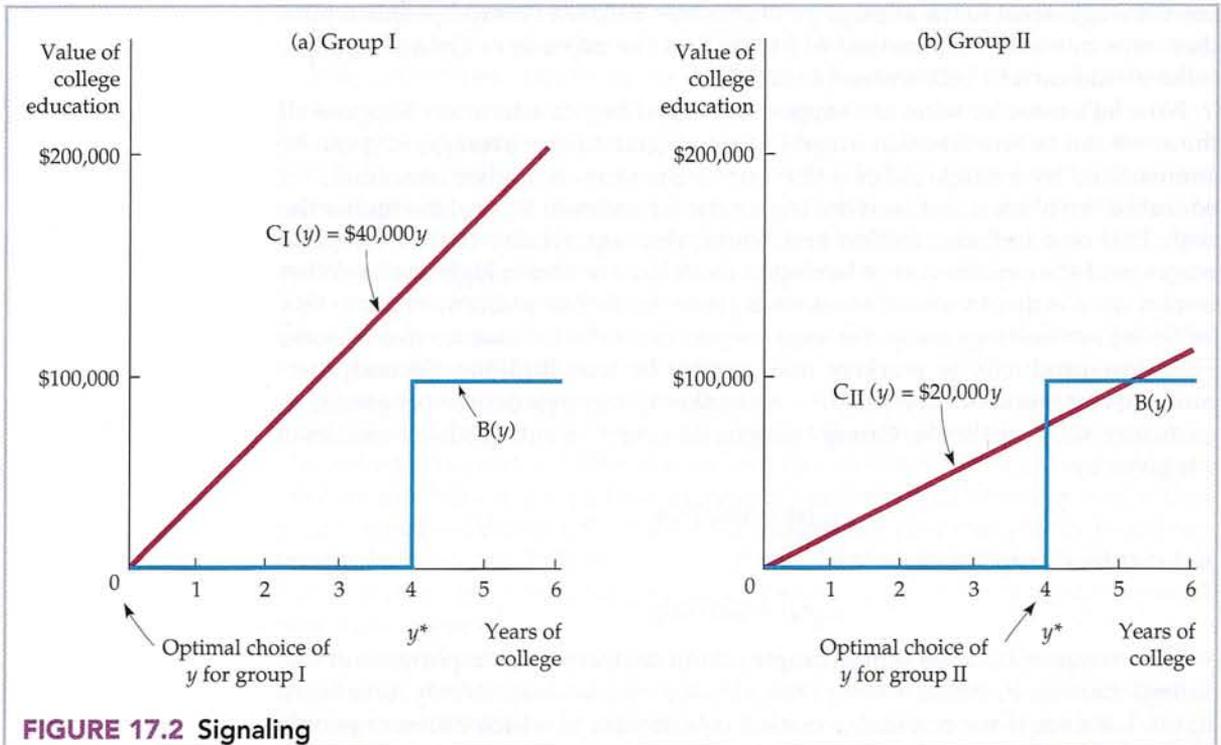
Now suppose (to keep things simple and to dramatize the importance of signaling) that *education does nothing to increase one's productivity; its only value is as a signal*. Let's see if we can find a market equilibrium in which different people obtain different levels of education, and in which firms look at education as a signal of productivity.

**Equilibrium** Consider the following possible equilibrium. Suppose firms use this decision rule: *Anyone with an education level of  $y^*$  or more is a Group II person and is offered a wage of \$20,000, while anyone with an education level below  $y^*$  is a Group I person and is offered a wage of \$10,000*. The particular level  $y^*$  that the firms choose is arbitrary, but for this decision rule to be part of an equilibrium, firms must have identified people correctly. Otherwise, the firms will want to change the rule. Will it work?

To answer this question, we must determine how much education the people in each group will obtain, *given that firms are using this decision rule*. To do this, remember that education allows one to get a better-paying job. The benefit of education  $B(y)$  is the *increase* in the wage associated with each level of education, as shown in Figure 17.2. Observe that  $B(y)$  is 0 initially, which represents the \$100,000 base 10-year earnings that are earned without any college education. For an education level less than  $y^*$ ,  $B(y)$  remains 0, because 10-year earnings remain at the \$100,000 base level. But when the education level reaches  $y^*$  or greater, 10-year earnings increase to \$200,000, increasing  $B(y)$  to \$100,000.

How much education should a person obtain? Clearly the choice is between *no education* (i.e.,  $y = 0$ ) and an education level of  $y^*$ . Why? Any level of education less than  $y^*$  results in the same base earnings of \$100,000. Thus there is no benefit from obtaining an education at a level above 0 but below  $y^*$ . Similarly, there is no benefit from obtaining an educational level above  $y^*$  because  $y^*$  is sufficient to allow one to enjoy the higher total earnings of \$200,000.

**Cost-Benefit Comparison** In deciding how much education to obtain, people compare the benefit of education with the cost. People in each group make the following cost-benefit calculation: *Obtain the education level  $y^*$  if the benefit (i.e., the increase in earnings) is at least as large as the cost of this education*. For both groups,



**FIGURE 17.2** Signaling

Education can be a useful signal of the high productivity of a group of workers if education is easier to obtain for this group than for a low-productivity group. In (a), the low-productivity group will choose an education level of  $y = 0$  because the cost of education is greater than the increased earnings resulting from education. However, in (b), the high-productivity group will choose an education level of  $y^* = 4$  because the gain in earnings is greater than the cost.

the benefit (the increase in earnings) is \$100,000. The costs, however, differ. For Group I, the cost is  $\$40,000y$ , but for Group II it is only  $\$20,000y$ . Therefore, Group I will obtain *no* education as long as

$$\$100,000 < \$40,000y^* \text{ or } y^* > 2.5$$

and Group II will obtain an education level  $y^*$  as long as

$$\$100,000 > \$20,000y^* \text{ or } y^* < 5$$

These results give us an equilibrium *as long as  $y^*$  is between 2.5 and 5*. Suppose, for example, that  $y^*$  is 4.0, as in Figure 17.2. In that case, people in Group I will find that education does not pay and will not obtain any, whereas people in Group II will find that education does pay and will obtain the level  $y = 4.0$ . Now, when a firm interviews job candidates who have no college education, it correctly assumes they have low productivity and offers them a wage of \$10,000. Similarly, when the firm interviews people who have four years of college, it correctly assumes their productivity is high, warranting a wage of \$20,000. We therefore have an equilibrium. High-productivity people will obtain a college education to signal their productivity; firms will read this signal and offer them a high wage.

This is a highly simplified model, but it illustrates a significant point: Education can be an important signal that allows firms to sort workers according to productivity. Some workers (those with high productivity) will want to obtain a college education *even if that education does nothing to increase their*



*productivity*. These workers simply want to identify themselves as highly productive, so they obtain the education needed to send a signal.

In the real world, of course, education *does* provide useful knowledge and *does* increase one's ultimate productivity. (We wouldn't have written this book if we didn't believe that.) But education also serves a signaling function. For example, many firms insist that a prospective manager have an MBA. One reason is that MBAs learn economics, finance, and other useful subjects. But there is a second reason: To complete an MBA program takes intelligence, discipline, and hard work, and people with those qualities tend to be very productive.

## Guarantees and Warranties

We have stressed the role of signaling in labor markets, but it can also play an important role in many other markets in which there is asymmetric information. Consider the markets for such durable goods as televisions, stereos, cameras, and refrigerators. Many firms produce these items, but some brands are more dependable than others. If consumers could not tell which brands tend to be more dependable, the better brands could not be sold for higher prices. Firms that produce a higher-quality, more dependable product must therefore make consumers aware of this difference. But how can they do it in a convincing way? The answer is *guarantees and warranties*.

Guarantees and warranties effectively signal product quality because an extensive warranty is more costly for the producer of a low-quality item than for the producer of a high-quality item. The low-quality item is more likely to require servicing under the warranty, for which the producer will have to pay. In their own self-interest, therefore, producers of low-quality items will not offer extensive warranties. Thus consumers can correctly view extensive warranties as signals of high quality and will pay more for products that offer them.

### EXAMPLE 17.2

### Working into the Night



Job market signaling does not end when one is hired. Even after a few years of employment, a worker will still know more about his abilities than will the employer. This is especially true for workers in knowledge-based fields such as engineering, computer programming, finance, law, management, and consulting. Although an unusually talented computer programmer, for example,

will be more skilled than his co-workers at writing programs that are efficient and bug-free, it may take several years before the firm fully recognizes this talent. Given this asymmetric information, what policy should employers use to determine promotions and salary increases? Can workers who are unusually talented and productive signal this fact and thereby receive earlier promotions and larger salary increases?

Workers can often signal talent and productivity *by working harder and longer hours*. Because more talented and productive workers tend to get more enjoyment and satisfaction from their jobs, it is less costly for them to send this signal than it is for other workers. The signal is therefore strong: It conveys information. As a result, employers can—and do—rely on this signal when making promotion and salary decisions.



This signalling process has affected the way many people work. Rather than an hourly wage, knowledge-based workers are typically paid a fixed salary for a 35- or 40-hour week and do not receive overtime pay if they work additional hours. Yet such workers increasingly work well beyond their weekly schedules. Surveys by the U.S. Labor Department, for example, found that the percentage of all workers who toil 49 hours or more a week rose from 13 percent in 1976 to over 18 percent in 2006.<sup>6</sup> Many young lawyers, accountants, consultants, investment bankers, and computer programmers regularly work into the night and on weekends, putting in 60- or 70-hour weeks. Is it surprising that these people are working so hard? Not at all. They are trying to send signals that can greatly affect their careers.

Employers rely increasingly on the signaling value of long hours as rapid technological change makes it harder for them to find other ways of assessing workers' skills and productivity. A study of software engineers at the Xerox Corporation, for example, found that many people work into the night because they fear that otherwise their bosses will conclude that they are shirkers who choose the easiest assignments. As the bosses make clear, this fear is warranted: "We don't know how to assess the value of a knowledge worker in these new technologies," says one Xerox manager, "so we value those who work into the night."

As corporations become more reluctant to offer lifetime job security, and as competition for promotion intensifies, salaried workers feel more and more pressure to work long hours. If you find yourself working 60- or 70-hour weeks, look at the bright side—the signal you're sending is a strong one.<sup>7</sup>

## 17.3 MORAL HAZARD

When one party is fully insured and cannot be accurately monitored by an insurance company with limited information, the insured party may take an action that increases the likelihood that an accident or an injury will occur. For example, if my home is fully insured against theft, I may be less diligent about locking doors when I leave, and I may choose not to install an alarm system. The possibility that an individual's behavior may change because she has insurance is an example of a problem known as *moral hazard*.

The concept of moral hazard applies not only to problems of insurance, but also to problems of workers who perform below their capabilities when employers cannot monitor their behavior ("job shirking"). In general, **moral hazard** occurs when a party whose actions are unobserved affects the probability or magnitude of a payment. For example, if I have complete medical insurance coverage, I may visit the doctor more often than I would if my coverage were limited. If the insurance provider can monitor its insurees' behavior, it can charge higher fees for those who make more claims. But if the company cannot monitor behavior, it may find its payments to be larger than expected. Under conditions of moral hazard, insurance companies may be forced to increase premiums for everyone or even to refuse to sell insurance at all.

• **moral hazard** When a party whose actions are unobserved can affect the probability or magnitude of a payment associated with an event.

<sup>6</sup>"At the Desk, Off the Clock and Below Statistical Radar," *New York Times*, July 18, 1999. Data on hours worked are available from the Current Population Survey (CPS), Bureau of Labor Statistics (BLS), at <http://www.bls.gov/cps/#charemp>; *Persons at Work in Agriculture and Nonagricultural Industries by Hours of Work*.

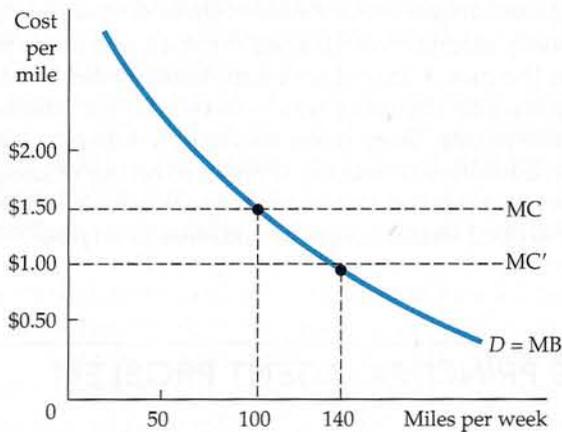
<sup>7</sup>For an interesting study of "time stress," see Daniel Hamermesh and Jungmin Lee, "Stressed Out on Four Continents: Time Crunch or Yuppie Kvetch?" *Review of Econ. and Stat.*, May 2007, 89, 374–383.



Consider, for example, the decisions faced by the owners of a warehouse valued at \$100,000 by their insurance company. Suppose that if they run a \$50 fire-prevention program for their employees, the probability of a fire is .005. Without this program, the probability increases to .01. Knowing this, the insurance company faces a dilemma if it cannot monitor the company's decision to conduct a fire-prevention program. The policy that the insurance company offers cannot include a clause stating that payments will be made only if there is a fire-prevention program. If the program were in place, the company could insure the warehouse for a premium equal to the expected loss from a fire—an expected loss equal to  $.005 \times \$100,000 = \$500$ . Once the insurance policy is purchased, however, the owners no longer have an incentive to run the program. If there is a fire, they will be fully compensated for their financial loss. Thus, if the insurance company sells a policy for \$500, it will incur losses because the expected loss from the fire will be \$1000 ( $.01 \times \$100,000$ ).

Moral hazard is a problem not only for insurance companies. It also alters the ability of markets to allocate resources efficiently. In Figure 17.3, for example,  $D$  gives the demand for automobile driving in miles per week. The demand curve, which measures the marginal benefits of driving, is downward sloping because some people switch to alternative transportation as the cost of driving increases. Suppose that initially, the cost of driving includes the insurance cost and that insurance companies can accurately measure miles driven. In this case, there is no moral hazard and the marginal cost of driving is given by  $MC$ . Drivers know that more driving will increase their insurance premiums and so increase their total cost of driving (the cost per mile is assumed to be constant). For example, if the cost of driving is \$1.50 per mile (50 cents of which is insurance cost), drivers will go 100 miles per week.

A moral hazard problem arises when insurance companies cannot monitor individual driving habits, so that insurance premiums do not depend on miles driven. In that case, drivers assume that any additional accident costs that they incur will be spread over a large group, with only a negligible portion accruing to



**FIGURE 17.3** The Effects of Moral Hazard

Moral hazard alters the ability of markets to allocate resources efficiently.  $D$  gives the demand for automobile driving. With no moral hazard, the marginal cost of transportation  $MC$  is \$1.50 per mile; the driver drives 100 miles, which is the efficient amount. With moral hazard, the driver perceives the cost per mile to be  $MC = \$1.00$  and drives 140 miles.



each of them individually. Because their insurance premiums do not vary with the number of miles that they drive, an additional mile of transportation will cost \$1.00, as shown by the marginal cost curve  $MC'$ , rather than \$1.50. The number of miles driven will increase from 100 to the socially inefficient level of 140.

Moral hazard not only alters behavior; it also creates economic inefficiency. The inefficiency arises because the insured individual perceives either the cost or the benefit of the activity differently from the true social cost or benefit. In the driving example of Figure 17.3, the efficient level of driving is given by the intersection of the marginal benefit (MB) and marginal cost (MC) curves. With moral hazard, however, the individual's perceived marginal cost ( $MC'$ ) is less than actual cost, and the number of miles driven per week (140) is higher than the efficient level at which marginal benefit is equal to marginal cost (100).

### EXAMPLE 17.3

### Reducing Moral Hazard: Warranties of Animal Health



For buyers of livestock, information about the animals' health is very important.<sup>8</sup> Unhealthy animals gain weight more slowly and are less likely to reproduce. Because of asymmetric information in the livestock market (sellers know the health of an animal better than buyers do), most states require warranties on the sale of livestock. Under these laws, sellers not only promise (warrant)

that animals are free from hidden diseases, but are responsible for all costs arising from any diseased animals.

Although warranties solve the problem of the seller having better information than the buyer, they also create a form of moral hazard. Guaranteeing reimbursement to the buyer for all costs associated with diseased animals means that insurance rates are not tied to the level of care that buyers or their agents take to protect their livestock against disease. As a result of these warranties, livestock buyers avoid paying for early diagnoses of diseased livestock, and losses increase.

In response to the moral hazard problem, many states have modified their animal warranty laws by requiring sellers to tell buyers whether livestock are diseased at the time of sale. Some states also require sellers to comply with state and federal animal health regulations, thereby reducing disease. Beyond these measures, however, warranties that animals are free from hidden disease must be in the form of explicit written or oral guarantees to buyers.

## 17.4 THE PRINCIPAL-AGENT PROBLEM

If monitoring the productivity of workers were costless, the owners of a business would ensure that their managers and workers were working effectively. In most firms, however, owners can't monitor everything that employees

<sup>8</sup>This example is based on Terence J. Centner and Michael E. Wetzstein, "Reducing Moral Hazard Associated with Implied Warranties of Animal Health," *American Journal of Agricultural Economics* 69 (1987): 143–50.



do—employees are better informed than owners. This information asymmetry creates a **principal–agent problem**.

An *agency relationship* exists whenever there is an arrangement in which one person's welfare depends on what another person does. The **agent** is the person who acts, and the **principal** is the party whom the action affects. A *principal–agent problem arises when agents pursue their own goals rather than the goals of the principal*. In our example, the manager and the workers are the agents, and the owners of the firm are the principals. In this case, the principal–agent problem results from the fact that managers may pursue their own goals, even at the cost of lower profits for the owners.

Agency relationships are widespread in our society. For example, doctors serve as agents for hospitals and, as such, may select patients and do procedures which, though consistent with their personal preferences, are not necessarily consistent with the objectives of the hospital. Similarly, managers of housing properties may not maintain the property the way that the owners would like. And sometimes insured parties may be seen as agents and insurance companies as principals.

How does incomplete information and costly monitoring affect the way agents act? And what mechanisms can give managers the incentives to operate in the owner's interest? These questions are central to any principal–agent analysis. In this section, we study the principal–agent problem from several perspectives. First, we look at the owner–manager problem within private and public enterprises. Second, we discuss ways in which owners can use contractual relationships with their employees to deal with principal–agent problems.

## The Principal–Agent Problem in Private Enterprises

Most large firms are controlled by management. Indeed, an individual family or financial institution owns more than 10 percent of the shares of only 16 of the 100 largest U.S. industrial corporations.<sup>9</sup> The fact that most individual stockholders have only a small percentage of a firm's total equity makes it difficult for them to obtain information about how well the firm's managers are performing. One function of owners (or their representatives) is to monitor the behavior of managers. But monitoring is costly, and information is expensive to gather and use, especially for an individual.

Managers of private enterprises can thus pursue their own objectives. But what are these objectives? One view is that managers are more concerned with growth than with profit per se: More rapid growth and larger market share provide more cash flow, which in turn allows managers to enjoy more perks. Another view emphasizes the utility that managers get from their jobs, not only from profit but also from the respect of their peers, the power to control the corporation, the fringe benefits and other perks, and long job tenure.

However, there are limitations to managers' ability to deviate from the objectives of owners. First, stockholders can complain loudly when they feel that managers are behaving improperly. In exceptional cases, they can oust the current management (perhaps with the help of the board of directors, whose job it is to monitor managerial behavior). Second, a vigorous market for corporate control can develop. If a takeover bid becomes more likely when the firm is poorly managed, managers will have a strong incentive to pursue the goal of profit maximization. Third, there can be a highly developed market for managers.

- **principal–agent problem** Problem arising when agents (e.g., a firm's managers) pursue their own goals rather than the goals of principals (e.g., the firm's owners).

- **agent** Individual employed by a principal to achieve the principal's objective.

- **principal** Individual who employs one or more agents to achieve an objective.

<sup>9</sup>See Merritt B. Fox, *Finance and Industrial Performance in a Dynamic Economy* (New York: Columbia University Press, 1987).



If managers who maximize profit are in great demand, they will earn high wages and so give other managers an incentive to pursue the same goal.

Unfortunately, the means by which stockholders control managers' behavior are limited and imperfect. Corporate takeovers may be motivated by personal and economic power, for example, instead of economic efficiency. The managerial labor market may also work imperfectly because top managers are frequently near retirement and have long-term contracts. The problem of limited stockholder control shows up most dramatically in executive compensation, which has grown very rapidly over the past several decades. In 2002, a *Business Week* survey of the 365 largest U.S. companies showed that the average CEO earned \$13.1 million in 2000, and executive pay has continued to increase at a double-digit rate. Even more disturbing is the fact that for the 10 public companies led by the highest-paid CEOs, there was a *negative* correlation between CEO pay and company performance.

It is clear that shareholders have been unable to adequately control managers' behavior. What can be done to address this problem? In theory, the answer is simple: One must find mechanisms that more closely align the interests of managers and shareholders. In practice, however, this is likely to prove difficult. Among those suggestions put into effect recently by the Securities and Exchange Commission, which regulates public companies, are reforms that grant more authority to independent outside directors. Other possible reforms would tie executive pay more closely to the long-term performance of the company. Reward structures that focus on profitability over a 5- to 10-year period are more likely to generate efficient incentives than more shortsighted reward structures. We will consider some additional solutions to this important principal-agent problem in the next section.

#### EXAMPLE 17.4 CEO Salaries

When Jack Welch retired as CEO of General Electric in 2001, his salary was \$16.7 million, and his benefits from stock options and other perks were worth millions more. In addition, his post-retirement benefits included a monthly income of \$2.1 million, the use of a company-owned Manhattan apartment, and unlimited use of the company's Boeing 737 business jet. In 2005, AT&T CEO Edward Whitacre was paid \$17.1 million, bringing his total pay to more than \$85 million over the previous five years—even though shareholder return during the period was negative 40 percent.<sup>10</sup> Other CEOs have also received extremely generous compensation packages, even when their companies were performing poorly.

CEO compensation has increased sharply over time. The average annual salary for production workers in the U.S. went from \$27,632 in 1990 to \$28,315 in 2005. Over the same period, the average annual compensation for CEOs grew from \$2.9 million to \$11 million. In other words, CEO compensation has gone from 107 times the pay of an average production worker to over 411 times as much.<sup>11</sup> Why? Have top managers become more productive, or are CEOs simply becoming more effective at extracting economic rents from their companies? The answer lies in the principal-agent problem, which is at the heart of CEO salary determination.

For years, many economists believed that executive compensation reflected an appropriate reward for talent. Recent evidence, however, suggests that managers

<sup>10</sup>Adam Geller, "Rise in Pay for CEOs Slows but Doesn't Stop," *International Herald Tribune*, April 20, 2006.

<sup>11</sup>Source: Institute for Policy Studies—United for a Fair Economy (2006).



have been able to increase their power over boards of directors and have used that power to extract compensation packages that are out of line with their economic contributions. In essence, managers have steadily increased their ability to extract economic rents. How has this happened?

First, most board of directors do not have the necessary information or independence to negotiate effectively with managers. Directors often cannot properly monitor executives' activities and therefore cannot effectively negotiate compensation packages that are tightly linked to their performance. Furthermore, boards consist of a mix of inside members, who either are or represent top executives, and outside members, who are chosen by and are often on close terms with top executives. Therefore, directors have a strong incentive to support executives in order to be re-nominated to the board or otherwise rewarded. Only if a compensation package is seen as outrageous from the point of view of outsiders do board members appear to bargain aggressively in shareholders' interests.

Second, managers have introduced forms of compensation that camouflage the extraction of rents from shareholders. For example, stock options, post-retirement perks, and pension plans can give executives substantial payoffs that appear costless because companies do not have to count them as expenses on their books.

Why has the amount of rent extraction grown so much over time? One reason is that boards of directors frequently use compensation consultants to advise them on comparable salaries paid to CEOs of other companies. Because a firm usually wants its CEO to be paid at least the average salary of other CEOs, the net result has been a gradual upward trend.

With the wave of corporate scandals that began in late 2001, the picture of rent extraction created by the principal-agent problem changed in 2002 and 2003. The median total executive pay package for 209 companies rose 11.8 percent, well above the rate of inflation, but a much smaller increase than in previous years. However, recent trends suggest that executive compensation may be heading back to its pre-2001 levels.

## The Principal-Agent Problem in Public Enterprises

The principal-agent framework can also help us understand the behavior of the managers of public organizations. These managers may also be interested in power and perks, both of which can be obtained by expanding their organization beyond its "efficient" level. Because it is also costly to monitor the behavior of public managers, there are no guarantees that they will produce the efficient output. Legislative checks on a government agency are not likely to be effective as long as the agency has better information about its costs than the legislature has.

Although the public sector lacks some of the market forces that keep private managers in line, government agencies can still be effectively monitored. First, managers of government agencies care about more than just the size of their agencies. Indeed, many choose lower-paying public jobs because they are concerned about the "public interest." Second, much like private managers, public managers are subject to the rigors of the managerial job market. If public managers are perceived to be pursuing improper objectives, their ability to obtain high salaries in the future might be impaired. Third, legislatures and other government agencies perform an oversight function. For example, the



Government Accounting Office and the Office of Management and Budget spend much of their energy monitoring other agencies.

At the local rather than the federal level, public managers are subject to even more checks. Suppose, for example, that a city transit agency has expanded bus service beyond the efficient level. Citizens can vote the transit managers out of office, or, if all else fails, use alternative transportation (or even move). Competition among agencies can be as effective as competition among private firms in constraining the behavior of managers.

### EXAMPLE 17.5

## Managers of Nonprofit Hospitals as Agents



Do the managers of nonprofit organizations have the same goals as those of for-profit organizations? Are nonprofit organizations more or less efficient than for-profit firms? We can get some insight into these issues by looking at the provision of health care. In a study of 725 hospitals, from 14 major hospital chains, researchers compared the return on investment and average costs of non-

profit and for-profit hospitals to determine if they performed differently.<sup>12</sup>

The study found that for 1977 and 1981, the rate of returns did indeed differ. In 1977, for example, for-profits earned an 11.6-percent return, while nonprofits earned 8.8 percent. In 1981, for-profits earned 12.7 percent and nonprofits only 7.4 percent. A straight comparison of returns and costs is not appropriate, however, because the hospitals perform different functions. For example, 24 percent of the nonprofit hospitals provide medical residency programs, as compared with only 6 percent of the for-profit hospitals. Similar differences can be found in the provision of specialty care, with 10 percent of the nonprofits having open-heart units, as compared to only 5 percent of the for-profits. In addition, while 43 percent of nonprofits have premature infant units, only 29 percent of the for-profits have equivalent units.

Using a statistical regression analysis, which controls for differences in the services performed, one can determine whether differences in services account for the higher costs. The study found that after adjusting for services performed, the average cost of a patient day in nonprofit hospitals was 8 percent higher than in for-profit hospitals. This difference implies that the profit status of the hospital affects its performance in the way principal-agent theory predicts: Without the competitive forces faced by for-profit hospitals, nonprofit hospitals may be less cost-conscious and therefore less likely to serve appropriately as agents for their principals—namely, society at large.

Of course, nonprofit hospitals provide services that society may well wish to subsidize. But the added cost of running a nonprofit hospital should be considered when determining whether it should be granted tax-exempt status.

<sup>12</sup>Regina E. Herzlinger and William S. Krasker, "Who Profits from Nonprofits?" *Harvard Business Review* 65 (January–February 1987): 93–106.



## Incentives in the Principal-Agent Framework

We have seen why managers' and owners' objectives are likely to differ within the principal-agent framework. How, therefore, can owners design reward systems so that managers and workers come as close as possible to meeting owners' goals? To answer this question, let's study a specific problem.

A small manufacturer uses labor and machinery to produce watches. The owners want to maximize profit. They must rely on a machine repairperson whose effort will influence the likelihood that machines break down and thus affect the firm's profit level. Revenue also depends on other random factors, such as the quality of parts and the reliability of other labor. As a result of high monitoring costs, the owners can neither measure the effort of the repairperson directly nor be sure that the same effort will always generate the same profit level. Table 17.2 describes these circumstances.

The table shows that the repairperson can work with either a low or high amount of effort. Low effort ( $a = 0$ ) generates either \$10,000 or \$20,000 in revenue (with equal probability), depending on the random factors that we mentioned. We've labeled the lower of the two revenue levels "bad luck" and the higher level "good luck." When the repairperson makes a high effort ( $a = 1$ ), revenue will be either \$20,000 (bad luck) or \$40,000 (good luck). These numbers highlight the problem of incomplete information: When the firm's revenue is \$20,000, the owners cannot know whether the repairperson has made a low or high effort.

Suppose the repairperson's goal is to maximize his wage payment less the cost (in terms of lost leisure and unpleasant work time) of the effort that he makes. To simplify, we'll suppose that the cost of effort is 0 for low effort and \$10,000 for high effort. (Formally,  $c = \$10,000a$ .)

Now we can state the principal-agent problem from the owners' perspective. The owners' goal is to maximize expected profit, given the uncertainty of outcomes and given the fact that the repairperson's behavior cannot be monitored. The owners can contract to pay the repairperson for his work, but the payment scheme must be based entirely on the measurable output of the manufacturing process, not on the repairperson's effort. To signify this link, we describe the payment scheme as  $w(R)$ , stressing that payments can depend only on measured revenue.

What is the best payment scheme? And can that scheme be as effective as one based on effort rather than output? The best payment scheme depends on the nature of production, the degree of uncertainty, and the objectives of both owners and managers. The arrangement will not always be as effective as an ideal scheme directly tied to effort. A lack of information can lower economic efficiency because both the owners' revenue and the repairperson's payment may fall at the same time.

Let's see how to design a payment scheme when the repairperson wishes to maximize his payment received net of the cost of effort made.<sup>13</sup> Suppose first

**TABLE 17.2** Revenue from Making Watches

	Bad Luck	Good Luck
Low effort ( $a = 0$ )	\$10,000	\$20,000
High effort ( $a = 1$ )	\$20,000	\$40,000

<sup>13</sup>We assume that because the repairperson is risk neutral, no efficiency is lost. If, however, the repairperson were risk averse, there would be an efficiency loss.



that the owners offer a fixed wage payment. Any wage will do, but we can see things most clearly if we assume that the wage is 0. (Here, 0 could represent a wage equal to the wage paid in other comparable jobs.) Facing a wage of 0, the repairperson has no incentive to make a high level of effort. The reason is that the repairperson does not share in any of the gains that the owners enjoy from the increased effort. It follows, therefore, that a fixed payment will lead to an inefficient outcome. When  $a = 0$  and  $w = 0$ , the owner will earn an expected revenue of \$15,000 and the repairperson a net wage of 0.

Both the owners and the repairperson will be better off if the repairperson is rewarded for his productive effort. Suppose, for example, that the owners offer the repairperson the following payment scheme:

$$\begin{aligned} \text{If } R = \$10,000 \text{ or } \$20,000, w = 0 \\ \text{If } R = \$40,000, w = \$24,000 \end{aligned} \quad (17.1)$$

Under this bonus arrangement, a low effort generates no payment. A high effort, however, generates an expected payment of \$12,000, and an expected payment less the cost of effort of  $\$12,000 - \$10,000 = \$2,000$ . Under this system, the repairperson will choose to make a high level of effort. This arrangement makes the owners better off than before because they get an expected revenue of \$30,000 and an expected profit of \$18,000.

This is not the only payment scheme that will work for the owners, however. Suppose they contract to have the worker participate in the following revenue-sharing arrangement. When revenues are greater than \$18,000,

$$w = R - \$18,000 \quad (17.2)$$

(Otherwise the wage is zero.) In this case, if the repairperson makes a low effort, he receives an expected payment of \$1,000. But if he makes a high level of effort, his expected payment is \$12,000, and his expected payment less the \$10,000 cost of effort is \$2,000. (The owners' profit is \$18,000, as before.)

Thus, in our example, a revenue-sharing arrangement achieves the same outcome as a bonus-payment system. In more complex situations, the incentive effects of the two types of arrangements will differ. However, the basic idea illustrated here applies to all principal-agent problems: When it is impossible to measure effort directly, an incentive structure that rewards the outcome of high levels of effort can induce agents to aim for the goals that the owners set.

## \*17.5 MANAGERIAL INCENTIVES IN AN INTEGRATED FIRM

We have seen that owners and managers of firms can have asymmetric information about demand, cost, and other variables. We've also seen how owners can design reward structures to encourage managers to make appropriate efforts. Now we focus our attention on firms that are *integrated*—that consist of several divisions, each with its own managers. Some firms are **horizontally integrated**: Several plants produce the same or related products. Others are also **vertically integrated**: Upstream divisions produce materials, parts, and components that downstream divisions use to produce final products. Integration creates organizational problems. We addressed some of these problems in the appendix to Chapter 11, where we discussed *transfer pricing* in

- **horizontal integration**

Organizational form in which several plants produce the same or related products for a firm.

- **vertical integration**

Organizational form in which a firm contains several divisions, with some producing parts and components that others use to produce finished products.



the vertically integrated firm—that is, how the firm sets prices for parts and components that upstream divisions supply to downstream divisions. Here we will examine problems that stem from asymmetric information.

## Asymmetric Information and Incentive Design in the Integrated Firm

In an integrated firm, division managers are likely to have better information about their different operating costs and production potential than central management has. This asymmetric information causes two problems.

1. How can central management elicit accurate information about divisional operating costs and production potential from divisional managers? This information is important because the inputs into some divisions may be the outputs of other divisions, because deliveries must be scheduled to customers, and because prices cannot be set without knowing overall production capacity and costs.
2. What reward or incentive structure should central management use to encourage divisional managers to produce as efficiently as possible? Should they be given bonuses based on how much they produce? If so, how should they be structured?

To understand these problems, consider a firm with several plants that all produce the same product. Each plant's manager has much better information about its production capacity than central management has. In order to avoid bottlenecks and to schedule deliveries reliably, central management wants to learn more about how much each plant can produce. It also wants each plant to produce as much as possible. Let's examine ways in which central management can obtain the information it wants while also encouraging plant managers to run the plants as efficiently as possible.

One way is to give plant managers bonuses based on either the total output of their plant or its operating profit. Although this approach would encourage managers to maximize output, it would penalize managers whose plants have higher costs and lower capacity. Even if these plants produced efficiently, their output and operating profit—and thus their bonuses—would be lower than those of plants with lower costs and higher capacities. Plant managers would also have no incentive to obtain and reveal accurate information about cost and capacity.

A second way is to ask managers about their costs and capacities and *then* base bonuses on how well they do relative to their answers. For example, each manager might be asked how much his or her plant can produce each year. Then at the end of the year, the manager receives a bonus based on how close the plant's output was to this target. For example, if the manager's estimate of the feasible production level is  $Q_f$ , the annual bonus in dollars,  $B$ , might be

$$B = 10,000 - .5(Q_f - Q) \quad (17.3)$$

where  $Q$  is the plant's actual output, 10,000 is the bonus when output is at capacity, and .5 is a factor chosen to reduce the bonus if  $Q$  is below  $Q_f$ .

Under this scheme, however, managers would have an incentive to *underestimate* capacity. By claiming capacities below what they know to be true, they can more easily earn large bonuses even if they do not operate efficiently. For example, if a manager estimates capacity to be 18,000 rather than 20,000, and the plant actually produces only 16,000, her bonus increases from \$8000 to



\$9000. Thus this scheme fails to elicit accurate information about capacity and does not ensure that plants will be run as efficiently as possible.

Now let's modify this scheme. We will still ask managers how much their plants can feasibly produce and tie their bonuses to this estimate. However, we will use a slightly more complicated formula than the one in (17.3) to calculate the bonus:

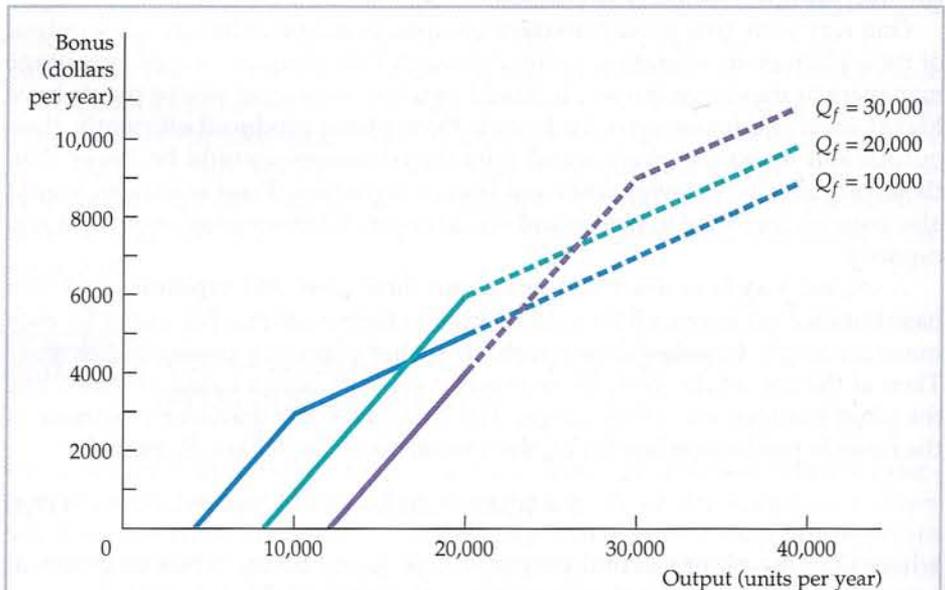
$$\begin{aligned} \text{If } Q > Q_f, \quad B &= .3Q_f + .2(Q - Q_f) \\ \text{If } Q \leq Q_f, \quad B &= .3Q_f - .5(Q_f - Q) \end{aligned} \quad (17.4)$$

The parameters (.3, .2, and .5) have been chosen so that each manager has the incentive to reveal the *true* feasible production level *and* to make  $Q$ , the actual output of the plant, as large as possible.

To see that this scheme does the job, look at Figure 17.4. Assume that the true production limit is  $Q^* = 20,000$  units per year. The bonus that the manager will receive if she states feasible capacity to be the true production limit is given by the line labeled  $Q_f = 20,000$ . This line is continued for outputs beyond 20,000 to illustrate the bonus scheme but dashed to signify the infeasibility of such production. Note that the manager's bonus is maximized when the firm produces at its limits of 20,000 units; the bonus is then \$6000.

Suppose, however, that the manager reports a feasible capacity of only 10,000. Then the bonus is given by the line labeled  $Q_f = 10,000$ . The maximum bonus is now \$5000, which is obtained by producing an output of 20,000. But note that this is less than the bonus that the manager would receive if she correctly stated the feasible capacity to be 20,000.

The same line of argument applies when the manager exaggerates available capacity. If the manager states feasible capacity to be 30,000 units per year, the



**FIGURE 17.4** Incentive Design in an Integrated Firm

A bonus scheme can be designed that gives a manager the incentive to estimate accurately the size of the plant. If the manager reports a feasible capacity of 20,000 units per year, equal to the actual capacity, then the bonus will be maximized (at \$6000).



bonus is given by the line  $Q_f = 30,000$ . The maximum bonus of \$4000, which is achieved at an output of 20,000, is less than the bonus that she could have received by reporting feasible capacity correctly.<sup>14</sup>

## Applications

Because the problem of asymmetric information and incentive design comes up often in managerial settings, incentive schemes like the one described above arise in many contexts. How, for example, can managers encourage salespeople to set and reveal realistic sales targets and then work as hard as possible to meet them?

Most salespeople cover specific territories. A salesperson assigned to a densely populated urban territory can usually sell more product than a salesperson assigned to a sparsely populated area. The company, however, wants to reward all salespeople equitably. It also wants to give them the incentive to work as hard as possible and to report realistic sales targets, so that it can plan production and delivery schedules. Companies have always used bonuses and commissions to reward salespeople, but incentive schemes have often been poorly designed. Typically, salespeople's commissions were proportional to their sales. This approach elicited neither accurate information about feasible sales targets nor maximum performance.

Today, companies are learning that bonus schemes like the one given by equation (17.4) provide better results. The salesperson can be given an array of numbers showing the bonus as a function of both the sales target (chosen by the salesperson) and the actual level of sales. (The numbers would be calculated from equation (17.4) or some similar formula.) Salespeople will quickly figure out that they do best by reporting feasible sales targets and then working as hard as possible to meet them.<sup>15</sup>

## 17.6 ASYMMETRIC INFORMATION IN LABOR MARKETS: EFFICIENCY WAGE THEORY

When the labor market is competitive, all who wish to work will find jobs for wages equal to their marginal products. Yet most countries have substantial unemployment even though many people are aggressively seeking work. Many of the unemployed would presumably work for an even lower wage rate than that being received by employed people. Why don't we see firms cutting wage rates, increasing employment levels, and thereby increasing profit? Can our models of competitive equilibrium explain persistent unemployment?

In this section, we show how the **efficiency wage theory** can explain the presence of unemployment and wage discrimination.<sup>16</sup> We have thus far determined

Recall from §14.1 that in a perfectly competitive labor market, firms hire labor to the point at which the real wage (the wage divided by the price of the product) is equal to the marginal product of labor.

• **efficiency wage theory**  
Explanation for the presence of unemployment and wage discrimination which recognizes that labor productivity may be affected by the wage rate.

<sup>14</sup>Any bonus of the form  $B = \beta Q_f + \alpha(Q = Q_f)$  for  $Q > Q_f$ , and  $B = \beta Q_f - \gamma(Q_f - Q)$  for  $Q \leq Q_f$ , with  $\gamma > \beta > \alpha > 0$  will work. See Martin L. Weitzman, "The New Soviet Incentive Model," *Bell Journal of Economics* 7 (Spring 1976): 251–6. There is a dynamic problem with this scheme that we have ignored: Managers must weigh a large bonus for good performance this year against being assigned more ambitious targets in the future. This is discussed in Martin Weitzman, "The 'Ratchet Principle' and Performance Incentives," *Bell Journal of Economics* 11 (Spring 1980): 302–8.

<sup>15</sup>See Jacob Gonik, "Tie Salesmen's Bonuses to Their Forecasts," *Harvard Business Review* (May–June 1978): 116–23.

<sup>16</sup>See Janet L. Yellen, "Efficiency Wage Models of Unemployment," *American Economic Review* 74 (May 1984): 200–5. The analysis relies on Joseph E. Stiglitz, "The Causes and Consequences of the Dependence of Quality on Price," *Journal of Economic Literature* 25 (March 1987): 1–48.



• **shirking model** Principle that workers still have an incentive to shirk if a firm pays them a market-clearing wage, because fired workers can be hired somewhere else for the same wage.

• **efficiency wage** Wage that a firm will pay to an employee as an incentive not to shirk.

In §14.2, we explain that the equilibrium wage is given by the intersection of the demand for labor curve and the supply of labor curve.

labor productivity according to workers' abilities and firms' investment in capital. Efficiency wage models recognize that labor productivity also depends on the wage rate. There are various explanations for this relationship. Economists have suggested that the productivity of workers in developing countries depends on the wage rate for nutritional reasons: Better-paid workers can afford to buy more and better food and are therefore healthier and can work more productively.

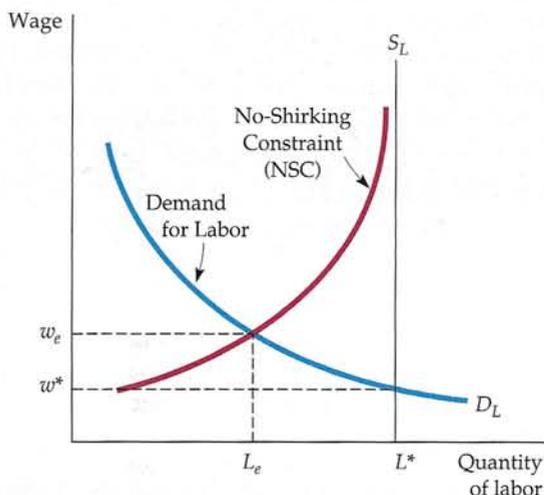
A better explanation for the United States is found in the **shirking model**. Because monitoring workers is costly or impossible, firms have imperfect information about worker productivity, and there is a principal-agent problem. In its simplest form, the shirking model assumes perfectly competitive markets in which all workers are equally productive and earn the same wage. Once hired, workers can either work productively or slack off (shirk). But because information about their performance is limited, workers may not get fired for shirking.

The model works as follows. If a firm pays its workers the market-clearing wage  $w^*$ , they have an incentive to shirk. Even if they get caught and are fired (and they might not be), they can immediately get hired somewhere else for the same wage. Because the threat of being fired does not impose a cost on workers, they have no incentive to be productive. As an incentive not to shirk, a firm must offer workers a higher wage. At this higher wage, workers who are fired for shirking will face a decrease in wages when hired by another firm at  $w^*$ . If the difference in wages is large enough, workers will be induced to be productive, and the employer will not have a problem with shirking. The wage at which no shirking occurs is the **efficiency wage**.

Up to this point, we have looked at only one firm. But all firms face the problem of shirking. All firms, therefore, will offer wages greater than the market-clearing wage  $w^*$ —say,  $w_e$  (efficiency wage). Does this remove the incentive for workers not to shirk because they will be hired at the higher wage by other firms if they get fired? No. Because all firms are offering wages greater than  $w^*$ , the demand for labor is less than the market-clearing quantity, and there is unemployment. Consequently, workers fired for shirking will face spells of unemployment before earning  $w_e$  at another firm.

Figure 17.5 shows shirking in the labor market. The demand for labor  $D_L$  is downward-sloping for the traditional reasons. If there were no shirking, the intersection of  $D_L$  with the supply of labor ( $S_L$ ) would set the market wage at  $w^*$ , and full employment would result ( $L^*$ ). With shirking, however, individual firms are unwilling to pay  $w^*$ . Rather, for every level of unemployment in the labor market, firms must pay some wage greater than  $w^*$  to induce workers to be productive. This wage is shown as the *no-shirking constraint (NSC) curve*. This curve shows the minimum wage, for each level of unemployment, that workers must earn in order not to shirk. Note that the greater the level of unemployment, the smaller the difference between the efficiency wage and  $w^*$ . Why is this so? Because with high levels of unemployment, people who shirk risk long periods of unemployment and therefore don't need much inducement to be productive.

In Figure 17.5, the equilibrium wage will be at the intersection of the NSC curve and  $D_L$  curves, with  $L_e$  workers earning  $w_e$ . This equilibrium occurs because the NSC curve gives the lowest wage that firms can pay and still discourage shirking. Firms need not pay more than this wage to get the number of workers they need, and they will not pay less because a lower wage will encourage shirking. Note that the NSC curve never crosses the labor supply curve. This means that there will always be some unemployment in equilibrium.

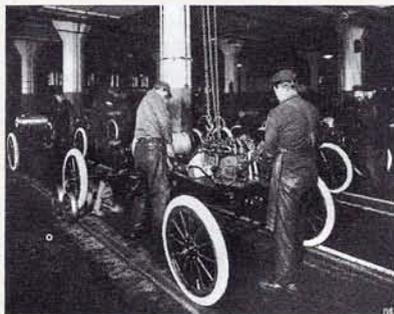


**FIGURE 17.5** Unemployment in a Shirking Model

Unemployment can arise in otherwise competitive labor markets when employers cannot accurately monitor workers. Here, the “no shirking constraint” (NSC) gives the wage necessary to keep workers from shirking. The firm hires  $L_e$  workers (at a higher than competitive efficiency wage  $w_e$ ), creating  $L^* - L_e$  of unemployment.

### EXAMPLE 17.6

### Efficiency Wages at Ford Motor Company



One of the early examples of the payment of efficiency wages can be found in the history of Ford Motor Company. Before 1913, automobile production depended heavily on skilled workers. But the introduction of the assembly line drastically changed the workplace. Now jobs demanded much less skill, and production depended on maintaining assembly-line equipment. But as automobile plants changed, workers became increasingly disenchanted. In 1913, turnover at

Ford was 380 percent. The following year, it rose to 1000 percent, and profit margins fell sharply.

Ford needed to maintain a stable workforce, and Henry Ford (and his business partner James Couzens) provided it. In 1914, when the going wage for a day’s work in industry averaged between \$2 and \$3, Ford introduced a pay policy of \$5 a day. The policy was prompted by improved labor efficiency, not generosity. The goal was to attract better workers who would stay with their jobs—and eventually to increase profits.

Although Henry Ford was attacked for it, his policy succeeded. His workforce did become more stable, and the publicity helped Ford’s sales. In addition, because Ford had his pick of workers, he could hire a group that was on average more productive. Ford stated that the wage increase did in fact



increase the loyalty and personal efficiency of his workers, and quantitative estimates support his statements. According to calculations by Ford's chief of labor relations, productivity increased by 51 percent. Another study found that absenteeism had been cut in half and discharges for cause had declined sharply. Thus the productivity increase more than offset the increase in wages. As a result, Ford's profitability rose from \$30 million in 1914 to \$60 million in 1916.

## SUMMARY

1. The seller of a product often has better information about its quality than the buyer. Asymmetric information of this type creates a market failure in which bad products tend to drive good products out of the market. Market failure can be eliminated if sellers offer standardized products, provide guarantees or warranties, or find other ways to maintain good reputations for their products.
2. Insurance markets frequently involve asymmetric information because the party buying insurance has better information about the risk involved than the insurance company. This can lead to adverse selection, in which poor risks choose to insure and good risks do not. Another problem for insurance markets is moral hazard, in which the insured takes less care to avoid losses after being insured.
3. Sellers can deal with the problem of asymmetric information by sending buyers signals about the quality of their products. For example, workers can signal high productivity by obtaining high levels of education.
4. Asymmetric information may make it costly for the owners of firms (principals) to monitor accurately the behavior of their managers (agents). Managers may seek higher fringe benefits for themselves or a goal of sales maximization, even though shareholders would prefer to maximize profit.
5. Owners can avoid some principal-agent problems by designing contracts that give their agents the incentive to perform productively.
6. Asymmetric information can explain why labor markets have unemployment even though some workers are actively seeking work. According to efficiency wage theory, a wage higher than the competitive wage (the efficiency wage) increases worker productivity by discouraging workers from shirking on the job.

## QUESTIONS FOR REVIEW

1. Why can asymmetric information between buyers and sellers lead to market failure when a market is otherwise perfectly competitive?
2. If the used car market is a "lemons" market, how would you expect the repair record of used cars that are sold to compare with the repair record of those not sold?
3. Explain the difference between adverse selection and moral hazard in insurance markets. Can one exist without the other?
4. Describe several ways in which sellers can convince buyers that their products are of high quality. Which methods apply to the following products: Maytag washing machines, Burger King hamburgers, large diamonds?
5. Why might a seller find it advantageous to signal the quality of a product? How are guarantees and warranties a form of market signaling?
6. Joe earned a high grade-point average during his four years of college. Is this achievement a strong signal to Joe's future employer that he will be a highly productive worker? Why or why not?
7. Why might managers be able to achieve objectives other than profit maximization, which is the goal of the firm's shareholders?
8. How can the principal-agent model be used to explain why public enterprises, such as post offices, might pursue goals other than profit maximization?
9. Why are bonus and profit-sharing payment schemes likely to resolve principal-agent problems, whereas a fixed-wage payment will not?
10. What is an efficiency wage? Why is it profitable for the firm to pay it when workers have better information about their productivity than firms do?



## EXERCISES

1. Many consumers view a well-known brand name as a signal of quality and will pay more for a brand-name product (e.g., Bayer aspirin instead of generic aspirin, or Birds Eye frozen vegetables instead of the super-market's own brand). Can a brand name provide a useful signal of quality? Why or why not?
2. Gary is a recent college graduate. After six months at his new job, he has finally saved enough to buy his first car.
  - a. Gary knows very little about the difference between makes and models. How could he use market signals, reputation, or standardization to make comparisons?
  - b. You are a loan officer in a bank. After selecting a car, Gary comes to you seeking a loan. Because he has only recently graduated, he does not have a long credit history. Nonetheless, the bank has a long history of financing cars for recent college graduates. Is this information useful in Gary's case? If so, how?
3. A major university bans the assignment of D or F grades. It defends its action by claiming that students tend to perform above average when they are free from the pressures of flunking out. The university states that it wants all its students to get As and Bs. If the goal is to raise overall grades to the B level or above, is this a good policy? Discuss this policy with respect to the problem of moral hazard.
4. Professor Jones has just been hired by the economics department at a major university. The president of the board of regents has stated that the university is committed to providing top-quality education for undergraduates. Two months into the semester, Jones fails to show up for his classes. It seems he is devoting all his time to research rather than to teaching. Jones argues that his research will bring prestige to the department and the university. Should he be allowed to continue exclusively with research? Discuss with reference to the principal-agent problem.
5. Faced with a reputation for producing automobiles with poor repair records, a number of American companies have offered extensive guarantees to car purchasers (e.g., a seven-year warranty on all parts and labor associated with mechanical problems).
  - a. In light of your knowledge of the lemons market, why is this a reasonable policy?
  - b. Is the policy likely to create a moral hazard problem? Explain.
6. To promote competition and consumer welfare, the Federal Trade Commission requires firms to advertise truthfully. How does truth in advertising promote competition? Why would a market be less competitive if firms advertised deceptively?
7. An insurance company is considering issuing three types of fire insurance policies: (i) complete insurance coverage, (ii) complete coverage above and beyond a \$10,000 deductible, and (iii) 90 percent coverage of all losses. Which policy is more likely to create moral hazard problems?
8. You have seen how asymmetric information can reduce the average quality of products sold in a market, as low-quality products drive out high-quality products. For those markets in which asymmetric information is prevalent, would you agree or disagree with each of the following? Explain briefly:
  - a. The government should subsidize *Consumer Reports*.
  - b. The government should impose quality standards—e.g., firms should not be allowed to sell low-quality items.
  - c. The producer of a high-quality good will probably want to offer an extensive warranty.
  - d. The government should require *all* firms to offer extensive warranties.
9. Two used car dealerships compete side by side on a main road. The first, Harry's Cars, always sells high-quality cars that it carefully inspects and, if necessary, services. On average, it costs Harry's \$8000 to buy and service each car that it sells. The second dealership, Lew's Motors, always sells lower-quality cars. On average, it costs Lew's only \$5000 for each car that it sells. If consumers knew the quality of the used cars they were buying, they would pay \$10,000 on average for Harry's cars and only \$7000 on average for Lew's cars.
 

Without more information, consumers do not know the quality of each dealership's cars. In this case, they would figure that they have a 50–50 chance of ending up with a high-quality car and are thus willing to pay \$8500 for a car.

Harry has an idea: He will offer a bumper-to-bumper warranty for all cars that he sells. He knows that a warranty lasting  $Y$  years will cost  $\$500Y$  on average, and he also knows that if Lew tries to offer the same warranty, it will cost Lew  $\$1000Y$  on average.

  - a. Suppose Harry offers a one-year warranty on all of the cars he sells.
    - i. What is Lew's profit if he *does not* offer a one-year warranty? If he *does* offer a one-year warranty?
    - ii. What is Harry's profit if Lew *does not* offer a one-year warranty? If he *does* offer a one-year warranty?
    - iii. Will Lew's match Harry's one-year warranty?
    - iv. Is it a good idea for Harry to offer a one-year warranty?
  - b. What if Harry offers a two-year warranty? Will this offer generate a credible signal of quality? What about a three-year warranty?
  - c. If you were advising Harry, how long a warranty would you urge him to offer? Explain why.



- \*10. As chairman of the board of ASP Industries, you estimate that your annual profit is given by the table below. Profit ( $\Pi$ ) is conditional upon market demand and the effort of your new CEO. The probabilities of each demand condition occurring are also shown in the table.

Market Demand	Low Demand	Medium Demand	High Demand
Market Probabilities	.30	.40	.30
Low Effort	$\Pi = \$5$ million	$\Pi = \$10$ million	$\Pi = \$15$ million
High Effort	$\Pi = \$10$ million	$\Pi = \$15$ million	$\Pi = \$17$ million

You must design a compensation package for the CEO that will maximize the firm's expected profit. While the firm is risk neutral, the CEO is risk averse. The CEO's utility function is

$$\text{Utility} = W^5 \text{ when making low effort}$$

$$\text{Utility} = W^5 - 100 \text{ when making high effort}$$

where  $W$  is the CEO's income. (The  $-100$  is the "utility cost" to the CEO of making a high effort.) You know the CEO's utility function, and both you and the CEO know all of the information in the preceding table. You

do *not* know the level of the CEO's effort at time of compensation or the exact state of demand. You do see the firm's profit, however.

Of the three alternative compensation packages below, which do you as chairman of ASP Industries prefer? Why?

Package 1: Pay the CEO a flat salary of \$575,000 per year

Package 2: Pay the CEO a fixed 6 percent of yearly firm profits

Package 3: Pay the CEO a flat salary of \$500,000 per year and then 50 percent of any firm profits *above* \$15 million

11. A firm's short-run revenue is given by  $R = 10e - e^2$ , where  $e$  is the level of effort by a typical worker (all workers are assumed to be identical). A worker chooses his level of effort to maximize wage less effort  $w - e$  (the per-unit cost of effort is assumed to be 1). Determine the level of effort and the level of profit (revenue less wage paid) for each of the following wage arrangements. Explain why these different principal-agent relationships generate different outcomes.
- $w = 2$  for  $e \geq 1$ ; otherwise  $w = 0$ .
  - $w = R/2$ .
  - $w = R - 12.5$ .