chapter 18

>> Uncertainty, Risk, and Private Information

AFTER THE FLOOD

In March 1998 the Flint River overflowed its banks and flooded the town of Albany, Georgia. Nobody was killed, but it was still a catastrophe. The town's residents were especially distraught because many of them had no flood insurance—they had dropped their coverage after a similar flood in 1994 led insurance companies to raise their *premiums.*

The case of Albany's second flood reminds us that uncertainty is an important feature of real-world economies. Up to this point, we have assumed that people make decisions with knowledge of exactly how those decisions will affect their welfare. In reality, people often make economic decisions—such as whether to build a house near a river—without full knowledge of their future consequences. As the residents of Albany learned, making decisions when the future is uncertain carries with it the risk of loss.

Yet it is often possible for individuals to use markets to reduce the risk they face. The Flint River flood made headlines because it's unusual for so many victims of a disaster to be without insurance. In fact, through insurance and other devices, the modern economy offers many ways for individuals to reduce their exposure to *risk*.

Does this mean that a market economy can solve all the problems created by uncertainty? Alas, no. Markets do very well at coping with situations in which nobody knows what will happen. But they run into trouble when some people know things that others do not—a situation known as *private information*. We'll see that private information can cause economic inefficiency by preventing mutually beneficial transactions.

AP/Wide World Photos

By failing to maintain their flood insurance, these Albany, Georgia, homeowners incurred huge losses in 1998.

What you will learn in this chapter:

- That **risk**—uncertainty about future outcomes—is an important feature of the economy, and that most people are **risk-averse:** they would like to avoid risk
- Why diminishing marginal utility makes people risk-averse and determines how much they are willing to pay to reduce risk
- How risk can be traded, with risk-averse people paying others to assume part of their risk
- How some of the economy's risk can be made to disappear through **diversification**
- How special problems are posed by **private information**—situations in which some people know things that other people do not

In this chapter we'll examine the economics of risk and private information. We'll start by looking at why people dislike risk. Then we'll explore how a market economy allows people to reduce risk, at a price. Finally, we'll turn to the special problems created when some people have information that others don't.

The Economics of Risk Aversion

In general, people don't like risk and are willing to pay a price to avoid it. Just ask the U.S. insurance industry, which collects more than $1 trillion in premiums every year. But what exactly is risk? And why don't people like it? To answer these questions, we need to look briefly at the concept of *expected value* and the meaning of uncertainty. Then we can turn to why people dislike *risk.*

Expectations and Uncertainty

The Lee family doesn't know how big its medical bills will be next year. If all goes well, it won't have any medical expenses at all. Let's assume that there's a 50 percent chance of that happening. But if family members require hospitalization or expensive drugs, they may face medical expenses of $10,000. Let's assume that there's also a 50 percent chance that these high medical expenses will materialize.

A **random variable** is a variable with an uncertain future value.

The **expected value** of a random variable is the weighted average of all possible values, where the weights on each possible value correspond to the probability of that value occurring.

In this example—which is designed to illustrate a point, rather than to be realistic—the Lees' medical expenses for the coming year are a **random variable,** a variable that has an uncertain future value. No one can predict which of its possible values, or outcomes, a random variable will take. But that doesn't mean we can say nothing about the Lees' future medical expenses. On the contrary, an actuary (a person trained in evaluating uncertain future events) could calculate the **expected value** of expenses next year—the weighted average of all possible values, where the weights on each possible value correspond to the probability of that value occurring. In this example, the expected value of the Lees' medical expenses is (0.5 × $0) + (0.5 × $10,000) = $5,000.

A **state of the world** is a possible future event.

To derive the general formula for the expected value of a random variable, we imagine that there are a number of different **states of the world**, possible future events. Each state is associated with a different realized value—the value that actually occurs—of the random variable. You don't know which state of the world will actually occur, but you can assign probabilities, one for each state of the world. Let's assume that P_1 is the probability of state 1, P_2 the probability of state 2, and so on. And you know the realized value of the random value in each state of the world: S_1 in state 1, S_2 in state 2, and so on. Let's also assume that there are N possible states. Then the expected value of the random variable is:

(18-1) *Expected value of a random variable*

$$EV = (P_1 \times S_1) + (P_2 \times S_2) + \ldots + (P_N \times S_N)$$

In the case of the Lee family, there are only two possible states of the world, each with a probability of 0.5.

Notice, however, that the Lee family doesn't actually expect to pay $5,000 in medical bills next year, regardless of what occurs. That's because in this example there is no state of the world in which the family pays exactly $5,000. Either the family pays nothing, or it pays $10,000. So the Lees face considerable uncertainty about their future medical expenses.

But what if the Lee family can buy health insurance that will cover its medical expenses, whatever they turn out to be? Suppose, in particular, that the family can pay $5,000 up front in return for full coverage of whatever medical expenses actually arise during the coming year. Then the Lees' future medical expenses are no longer uncertain *for them*: in return for $5,000—an amount equal to the expected value of the medical expenses—the insurance company assumes all responsibility for paying those medical expenses. Would this be a good deal from the Lees' point of view?

Yes, it would—or at least most families would think so. Most people prefer, other things equal, to reduce **risk**—uncertainty about future outcomes. (We'll focus here on **financial risk,** in which the uncertainty is about monetary outcomes, as opposed to uncertainty about outcomes that can't be assigned a monetary value.) In fact, most people are willing to pay a substantial price to reduce their risk; that's why we have an insurance industry. But before we study the market for insurance, we need to understand why people feel that risk is a bad thing. The answer, as we'll now see, is a concept we first encountered in our analysis of consumer demand, back in Chapter 10: *diminishing marginal utility.*

Risk is uncertainty about future outcomes. When the uncertainty is about monetary outcomes, it becomes **financial risk**.

The Logic of Risk Aversion

To understand how diminishing marginal utility gives rise to risk aversion, we need to look not only at the Lees' medical costs but also at how those costs affect the income the family has left after medical expenses. Let's assume that the family knows that it will have an income of $30,000 next year. If the family has no medical expenses, it will be left with all of that income. If its medical expenses are $10,000, its income after medical expenses will be only $20,000. Since we have assumed that there is an equal chance of these two outcomes, the Lees' expected income after medical expenses is (0.5 × $30,000) + (0.5 × $20,000) = $25,000.

But as we'll now see, if the family's utility function has the typical shape, its **expected utility**—the expected value of its total utility given uncertainty about future outcomes—is less than it would be if the family didn't face any risk and knew with certainty that its income after medical expenses would be $25,000.

Expected utility is the expected value of an individual's total utility given uncertainty about future outcomes.

To see why, we need to look at how total utility depends on income. Panel (a) of Figure 18-1 on page 434 shows a hypothetical utility function for the Lee family, where total utility depends on income—the amount of money the Lees have available for consumption of goods and services (after they have paid any medical bills). The accompanying table shows how the family's total utility varies over the income range of $20,000 to $30,000. As usual, the utility function is upward sloping, because more income leads to higher utility. Notice as well that the curve gets flatter as we move up and to the right, which relects its diminishing marginal utility.

In Chapter 10 we applied the principle of diminishing marginal utility to individual goods: each successive unit of a good that a consumer purchases adds less to his or her utility. The same principle applies to income used for consumption: each successive dollar of income adds less to utility than the previous dollar. Panel (b) shows how marginal utility varies with income, confirming that marginal utility of income falls as income rises. As we'll see in a moment, diminishing marginal utility is the key to understanding the desire of individuals to reduce risk—their *risk aversion*.

To analyze how a person's utility is affected by risk, economists start from the assumption that individuals facing uncertainty maximize their *expected* utility. We can use the data in Figure 18-1 to calculate the Lee family's expected utility. We'll first do the calculation assuming that the Lees have no insurance, and then we'll recalculate it assuming that they have purchased insurance.

Without insurance, if the Lees are lucky and don't incur any medical expenses, they will have an income of $30,000, generating total utility of 1,080 utils. But if they have no insurance and are unlucky, incurring $10,000 in medical expenses, they will have an income of $20,000 and total utility of only 920 utils. So *without insurance*, the family's expected utility is (0.5 × 1,080) + (0.5 × 920) = 1,000 utils.

Now let's suppose that an insurance company offers to pay whatever medical expenses the family incurs during the next year in return for a **premium**—a payment to the insurance company—of $5,000. Note that the amount of the premium in this

A **premium** is a payment to an insurance company in return for the promise to pay in certain states of the world.

Figure 18-1 The Utility Function of a Risk-Averse Family

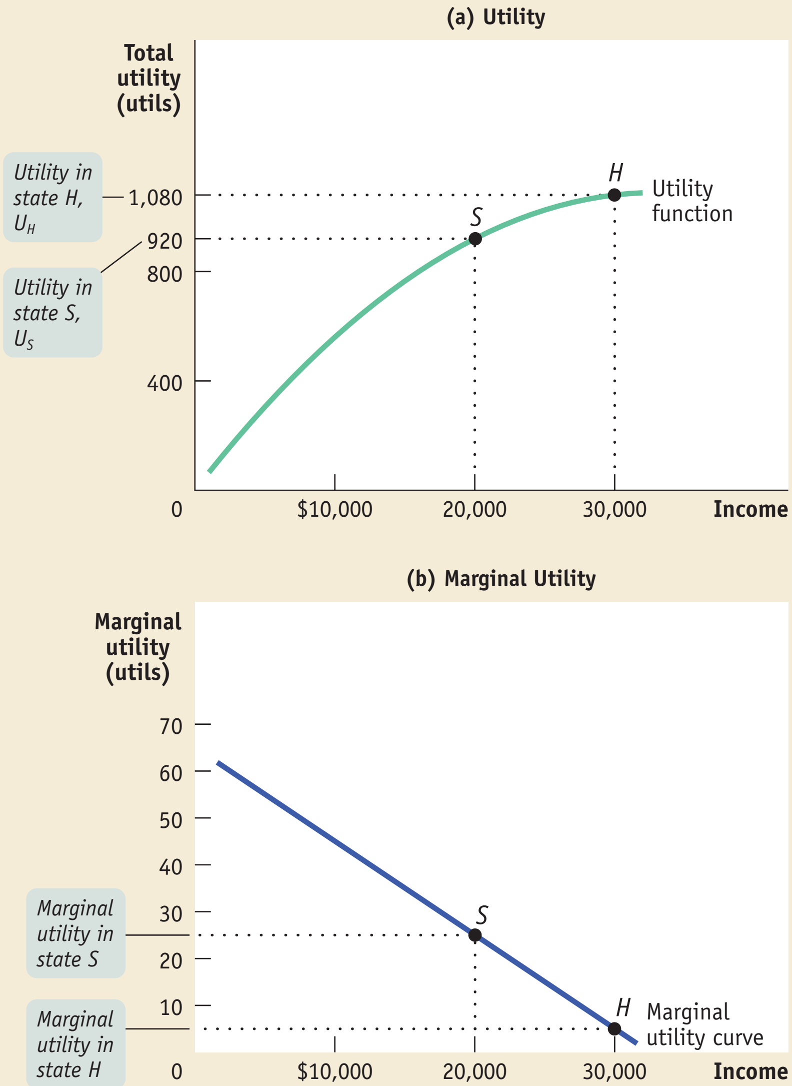

Income	Total utility (utils)
$20,000	920
21,000	945
22,000	968
23,000	989
24,000	1,008
25,000	1,025
26,000	1,040
27,000	1,053
28,000	1,064
29,000	1,073
30,000	1,080

Panel (a) shows how the total utility of the Lee family depends on its income after medical expenses. The curve is upward sloping: more income leads to higher total utility. But it gets flatter as we move up it and to the right, reflecting diminishing marginal utility. Panel (b) reflects the negative relationship between income and marginal utility when there is risk aversion: the marginal utility from an additional dollar is lower the higher your income. So, the marginal utility of income is higher when the family has high medical expenses (point *S*) than when it has low medical expenses (point *H*).

A **fair insurance policy** is an insurance policy for which the premium is equal to the expected value of the claims.

case is equal to the expected value of the Lees' medical expenses—the expected value of their future claims against the policy. An insurance policy with this feature, for which the premium is equal to the expected value of the claims, has a special name—a **fair insurance policy.**

If the family purchases this fair insurance policy, the expected value of its income available for consumption is the *same* as it would be without insurance: $25,000—that is, $30,000 minus the $5,000 premium. But the family's risk has been eliminated: the family has an income of $25,000 *for sure*, which means that it receives the utility level associated with an income of $25,000. Reading from the table in Figure 18-1, we see that this utility level is 1,025 utils. Or to put it a slightly different way, their expected utility with insurance is 1 × 1,025 = 1,025 utils, because with insurance they will receive a utility of 1,025 utils with a probability of 1. And this is higher than the level of expected utility without insurance—only 1,000 utils. So by eliminating risk through the purchase of a fair insurance policy, the family increases its expected utility even though its expected income hasn't changed.

TABLE 18-1

The Effect of Insurance on the Lee Family's Expected Income and Expected Utility

	Income in different states of the world			
	$0 in medical expenses (0.5 probability)	$10,000 in medical expenses (0.5 probability)	Expected income	Expected utility
Without insurance	$30,000	$20,000	(0.5 × $30,000) + (0.5 × $20,000) = $25,000	(0.5 × 1,080 utils) + (0.5 × 920 utils) = 1,000 utils
With insurance	$25,000	$25,000	(0.5 × $25,000) + (0.5 × $25,000) = $25,000	(0.5 × 1,025 utils) + (0.5 × 1,025 utils) = 1,025 utils

The calculations for this example are summarized in Table 18-1. This example shows that the Lees, like most people in real life, are **risk-averse:** they will choose to reduce the risk they face when the cost of that reduction leaves the expected value of their income or wealth unchanged. So the Lees, like most people, will be willing to buy fair insurance.

Risk-averse individuals will choose to reduce the risk they face when that reduction leaves the expected value of their income or wealth unchanged.

You might think that this result depends on the specific numbers we have chosen. In fact, however, the proposition that purchase of a fair insurance policy increases expected utility depends on only one assumption: diminishing marginal utility. The reason is that *a dollar gained when income is low adds more to utility than a dollar lost when income is high.* That is, having an additional dollar matters more when you are facing hard times than when you are facing good times. And as we will shortly see, a fair insurance policy is desirable because it transfers a dollar from high-income states (where it is valued less) to low-income states (where it is valued more).

But first, let's see how diminishing marginal utility leads to risk aversion by examining expected utility more closely. In the case of the Lee family, there are two states of the world; let's call them *H* and *S,* for healthy and sick. In state *H* the family has no medical expenses; in state *S* it has $10,000 in medical expenses. Let's use the symbols U_H and U_S to represent the Lee family's utility in each state. Then the family's expected utility is

(18-2) Expected utility = (Probability of healthy state × Utility in healthy state) + (Probability of sick state × Utility in sick state)

$$\text{EU} = (0.5 \times U_H) + (0.5 \times U_S)$$

The fair insurance policy *reduces* the family's income in state H by $5,000, but it *increases* the family's income in state S by the same amount. As we've just seen, we can use the utility function to directly calculate the effects of these changes on expected utility. But as we have seen in many other contexts, we often get more insight into individual choice by focusing on *marginal* utility.

To use marginal utility to analyze the effects of fair insurance, let's imagine introducing the insurance a bit at a time, say in 5,000 small steps. At each of these steps, we reduce income in state H by $1 and simultaneously increase income in state S by $1. At each of these steps, utility in state H falls by the marginal utility of income in that state but utility in state S rises by the marginal utility of income in that state.

Now look again at panel (b) of Figure 18-1, which shows how marginal utility varies with income. Point *S* shows marginal utility when the Lee family's income is $20,000; point *H* shows marginal utility when income is $30,000. Clearly, marginal utility is higher when income after medical expenses is low. Because of diminishing marginal utility, an additional dollar of income adds more to utility when the family has low income than when it has high income.

This tells us that the gain in expected utility from increasing income in state S is larger than the loss in expected utility from reducing income in state H. So at each step of the process of reducing risk, expected utility increases: reducing risk increases expected utility. This is the same thing as saying that the family is risk-averse: risk aversion is a result of diminishing marginal utility.

Almost everyone is risk-averse, because almost everyone has diminishing marginal utility. But the degree of risk aversion varies among individuals—some people are more risk-averse than others. To illustrate this point, Figure 18-2 compares two individuals, Danny and Mel. We suppose that each of them is confronted with the possibility of earning either $1,000 more or $1,000 less than his expected income. Panel (a) shows how each individual's utility would be affected by the change in income. Danny would gain very little utility from a rise in income, which moves him from N to H_D, but lose a lot of utility from a fall in income, which moves him from N to L_D. That is, he is highly risk-averse. This is reflected in the lower part of the figure by his steeply declining marginal utility curve. Mel, though, would gain almost as much utility from high-

Figure 18-2

Differences in Risk Aversion

Danny and Mel have different utility functions. Danny is highly risk-averse: a gain of $1,000 in income, which moves him from N to H_D, adds only a little to his utility, but a $1,000 fall in income, which moves him from N to L_D, reduces his utility a lot. By contrast, Mel gains almost as much utility from a $1,000 rise in income (the movement from N to H_M) as he loses from a $1,000 fall in income (the movement from N to L_M). This difference—reflected in the differing slopes of the two men's marginal utility curves—means that Danny would be willing to pay much more than Mel for insurance.

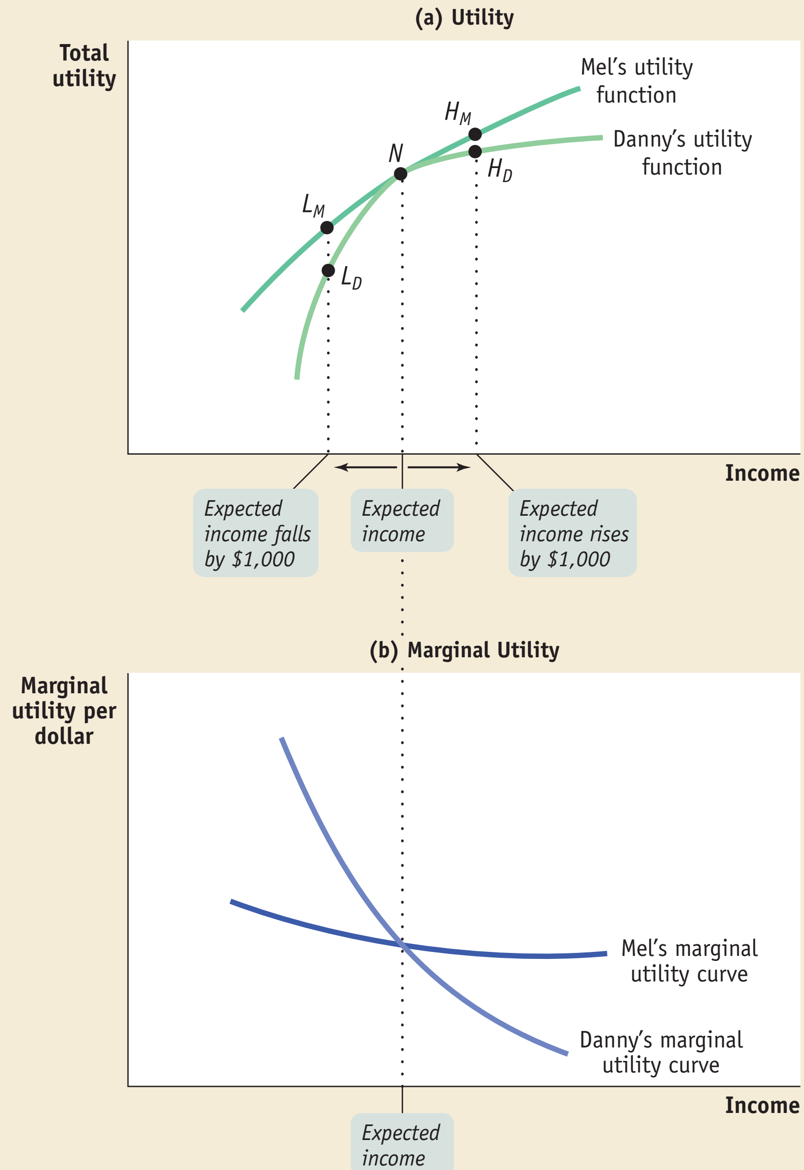

FOR INQUIRING MINDS
THE PARADOX OF GAMBLING

If most people are risk-averse and risk-averse individuals won't take a fair gamble, how come Las Vegas, Atlantic City, Indian reservations, and other places where gambling is legal do so much business?

After all, a casino doesn't even offer gamblers a fair gamble: all the games in any gambling facility are designed so that, on average, the house makes money. So why would anyone play their games?

You might argue that the gambling industry caters to the minority of people who are actually the opposite of risk-averse: risk-loving. But a glance at the customers of Las Vegas hotels quickly refutes that hypothesis: most of them aren't daredevils who also sky-dive and hang-glide. Instead, most of them are ordinary people who have health and life insurance and who wear seat belts. In other words, they are risk-averse like the rest of us.


James Marshall/Corbis

Gambling: Enjoyment or addiction?

So why do people gamble? Presumably because they enjoy the experience.

Also, gambling may be one of those areas where the assumption of rational behavior goes awry. Psychologists have concluded that gambling can be addictive in ways that are not that different from the addictive effects of drugs. Taking dangerous drugs is irrational; so is excessive gambling. Alas, both happen all the same.

er income, which moves him from N to H_M, as he would lose from lower income, which moves him from N to L_M. He is barely risk-averse at all. This is reflected in his marginal utility curve, which is almost flat. So other things equal, Danny will gain a lot more from insurance than Mel will.

Individuals differ in risk aversion for two main reasons: differences in *preferences* and differences in *income* or *wealth*.

- *Differences in preferences*. Other things equal, people simply differ in how much their marginal utility is affected by their level of income. Someone whose marginal utility does not depend very much on income won't be highly risk-averse.
- *Differences in income* or *wealth*. The possible loss of $1,000 makes a big difference to a family living below the poverty line; it makes very little difference to someone who earns $1 million a year. In general, people with high incomes or high wealth will be less risk-averse.

Differences in risk aversion have an important consequence: they affect how much an individual is willing to pay to avoid risk.

Paying to Avoid Risk

The risk-averse Lee family is clearly better off taking out a fair insurance policy—a policy that leaves their expected income unchanged but eliminates their risk. Unfortunately, real insurance policies are rarely fair: because insurance companies have to cover other costs, such as salaries for salespeople and actuaries, they charge more than they expect to pay in claims. Will the Lee family still want to purchase an "unfair" insurance policy—one for which the premium is larger than the expected claims?

It depends on the size of the premium. Look again at Table 18-1. We know that without insurance expected utility is 1,000 utils and

PITFALLS
BEFORE THE FACT VERSUS AFTER THE FACT

Why is an insurance policy different from a doughnut?

No, it's not a riddle. Although the supply and demand for insurance behave like the supply and demand for any good, the payoff is very different. When you buy a doughnut, you know what you're going to get; when you buy insurance, by definition you *don't* know what you're going to get. If you bought car insurance and then didn't have an accident, you got nothing from the policy, except peace of mind, and might wish that you hadn't bothered. But if you did have an accident, you probably wished that you had bought insurance that covered more of the cost.

This means we have to be careful in assessing the rationality of insurance purchases (or, for that matter, any decision made in the face of uncertainty). *After the fact*—after the uncertainty has been resolved—such decisions are almost always subject to second-guessing. But that doesn't mean that the decision was wrong *before the fact*, given the information available at the time.

One highly successful Wall Street investor told us that he never looks back—that as long as he believes he made the right decision given what he knew when he made it, he never reproaches himself if things turn out badly. That's the right attitude, and it probably contributes to his success.

that insurance costing $5,000 raises expected utility to 1,025 utils. If the premium were $6,000, the Lees would be left with an income of $24,000, which, as you can see from Figure 18-1, would give them a utility level of 1,008 utils—which is still higher than their expected utility if they had no insurance at all. So the Lees would be willing to buy insurance with a $6,000 premium. But they wouldn't be willing to pay $7,000, which would reduce their income to $23,000 and their utility to 989 utils.

This example shows that risk-averse individuals are willing to make deals that reduce their expected income but also reduce their risk: they are willing to pay a premium that exceeds their expected claims. The more risk-averse they are, the higher the premium they are willing to pay. That willingness to pay is what makes the insurance industry possible.

economics in action

Warranties

Many expensive consumer goods—stereos, TVs, cars—come with some form of *warranty*. Typically, the manufacturer guarantees that it will repair or replace the item if something goes wrong with it during some specified period after purchase—usually six months or one year.

Why do manufacturers offer warranties? Part of the answer is that they *signal* to consumers that the goods are of high quality (see the discussion of private information later in this chapter). But mainly warranties are a form of consumer insurance.

For many people, the cost of repairing or replacing an expensive item like a stereo—or, worse yet, a car—would be a serious burden. If they were obliged to come up with the cash, their consumption of other goods would be restricted; as a result, their marginal utility of income would be higher than normal.

So a warranty that covers the cost of repair or replacement increases the consumer's expected utility, even if the cost of the warranty is greater than the expected future claims paid by the manufacturer. ■

>>QUICK REVIEW

- Uncertainty about future outcomes entails *risk*. When faced with uncertainty, a consumer chooses the option that yields the highest level of *expected utility*.
- Most people are *risk-averse*: they would be willing to purchase a fair insurance policy for which the premium is equal to the expected value of the claims.
- Risk aversion arises from diminishing marginal utility. Differences in preferences and in income or wealth lead to differences in risk aversion.
- Depending on the size of the premium, a risk-averse person may be willing to purchase an "unfair" insurance policy—a policy with a premium larger than the expected claims. The greater your risk aversion, the greater the premium you are willing to pay.

>>CHECK YOUR UNDERSTANDING 18-1

1. Explain which of the following events will make you more likely to buy car insurance.
   a. You must work to pay your living expenses, and you need a car to get to work.
   b. Your parents are wealthy and can easily buy you another car if you need it.
2. Karma's income next year is uncertain: there is a 60% probability she will make $22,000 and a 40% probability she will make $35,000. The accompanying table shows some income and utility levels for Karma:

Income	Total utility (utils)
$22,000	850
25,000	1,014
26,000	1,056
35,000	1,260

   a. What is Karma's expected income? her expected utility?
   b. What certain income level leaves her as well off as her uncertain income? What does this imply about Karma's attitudes toward risk? Explain.
   c. Would Karma be willing to pay some amount of money greater than zero for an insurance policy that guarantees her an income of $26,000? Explain.

Solutions appear at back of book.

Buying, Selling, and Reducing Risk

Lloyd's of London is the oldest existing commercial insurance company, and it is an institution with an illustrious past. Originally formed in the eighteenth century as a commercial venture to help merchants cope with the risks of commerce, it grew in the heyday of the British Empire into a mainstay of imperial trade.

The basic idea of Lloyd's was simple. In the eighteenth century, shipping goods via sailing vessels was risky: the chance that a ship would sink in a storm or be captured by pirates was fairly high. The merchant who owned the ship and its cargo could easily be financially ruined by such an event. Lloyd's matched shipowners seeking insurance with wealthy investors who promised to compensate a merchant if his ship were lost. In return, the merchant paid the investor a fee in advance; if his ship *didn't* sink, the investor still kept the fee. In effect, the merchant paid a price to relieve himself of risk. By matching people who wanted to purchase insurance with people who wanted to provide it, Lloyd's performed the functions of a market. The fact that British merchants could use Lloyd's to reduce their risk made many more people in Britain willing to undertake merchant trade.

Insurance companies have changed quite a lot from the early days of Lloyd's. They no longer consist of wealthy individuals deciding on insurance deals over port and boiled mutton. But asking why Lloyd's worked to the mutual benefit of merchants and investors is a good way to understand how the market economy as a whole "trades" and thereby transforms risk.

The insurance industry rests on two principles. The first is that trade in risk, like trade in any good, can produce mutual gains from trade; in this case, the gains come when people who are less willing to bear risk transfer it to people who are more willing to bear it. The second is that some risk can be made to disappear through *diversification*. Let's consider each principle in turn.

Trading Risk

It may seem a bit strange to talk about "trading" risk. After all, risk is a bad thing—and aren't we supposed to be trading *goods and services?*

But people often trade away things they don't like to other people who dislike them less. Suppose you have just bought a house for $100,000, the average price for a house in your community. But you have now learned, to your horror, that the building next door is being turned into an all-night disco. You want to sell the house immediately and are willing to accept $95,000 for it. But who will now be willing to buy it? The answer: a person who doesn't really mind late-night noise. Such a person might be willing to pay up to $100,000. So there is an opportunity here for a mutually beneficial deal—you are willing to sell for as little as $95,000, and the other person is willing to pay as much as $100,000, so any price in between will benefit both of you.

The key point is that the two parties have different sensitivities to noise, which enables those who most dislike noise, in effect, to pay other people to make their lives quieter. Trading risk works exactly the same way: people who want to reduce the risk they face can pay other people to take some of their risk away.

As we saw in the previous section, individual personality accounts for some of the variations in people's attitudes toward risk, but differences in wealth are probably the principal reason behind different risk sensitivities. Lloyd's made money by matching wealthy investors who were more risk-tolerant with less wealthy and therefore more risk-averse shipowners who wanted to purchase insurance.

Suppose, staying with our Lloyd's of London story, that a merchant whose ship went down would lose £1,000, and that there was 1 chance in 10 of such a disaster. The expected loss in this case, then, would be 1/10 of £1,000, or £100. But the merchant, whose whole livelihood was at stake, might have been willing to pay £150 to be compensated in the amount of £1,000 if the ship sank. Meanwhile, a wealthy investor for whom the loss of £1,000 was no big deal would have been willing to take this risk for a return only slightly better than the expected loss—say, £110. Clearly, there is room for a mutually beneficial deal here: the merchant pays something less than £150 and more than £110—say, £130—in return for compensation if the ship goes down. In effect, he has paid a less risk-averse individual to bear the burden of his risk. Everyone has been made better off by this transaction.

The funds that an insurer places at risk when agreeing to provide insurance is called his or her **capital at risk**.

The funds that an insurer places at risk when agreeing to provide insurance are called his or her **capital at risk.** In our example, the wealthy Lloyd's investor places capital of £1,000 at risk in return for a premium of £130. In general, the amount of capital that potential insurers are willing to place at risk depends, other things equal, on the premium offered. If ships are worth £1,000 and have a 1 in 10 chance of going down, nobody would offer insurance for less than a £100 premium, equal to the expected claim. In fact, only an investor who isn't risk-averse at all and would insure a ship for a premium of £100 would be willing to offer a policy at that price, because accepting a £100 premium would leave the insurer's expected income unchanged while increasing his or her risk. Suppose there is one investor who isn't risk-averse; but the next most willing investor is slightly risk-averse and insists on a £105 premium. The next investor, being somewhat more risk-averse, demands a premium of £110, and so on. By varying the premium and asking how many insurers would be willing to provide insurance, we can trace out a supply curve for insurance, as shown in Figure 18-3. As the premium increases as we move up the supply curve, more risk-averse investors are induced to provide coverage.

Meanwhile, potential buyers will consider their willingness to pay a given premium, defining the demand curve for insurance. In Figure 18-4, the highest premium that any shipowner is willing to pay is £200. Who's willing to pay this? The most risk-averse shipowner, of course. The next most risk-averse shipowner might be willing to pay £190, the next after that £180, and so on.

Now imagine a market in which there are thousands of shipowners and potential insurers, so that the supply and demand curves for insurance are smooth lines. In this market, as in markets for ordinary goods and services, there will be an equilibrium price and quantity. In Figure 18-5 the equilibrium premium is £130, with a total quantity of 5,000 policies bought and sold, representing a total capital at risk of £5,000,000.

An **efficient allocation of risk** is an allocation of risk in which those who are most willing to bear risk are those who end up bearing it.

Notice that in this market risk is transferred from the people who most want to get rid of it (the most risk-averse shipowners) to the people least bothered by risk (the least risk-averse investors). So just as markets for goods and services typically produce an efficient allocation of resources, markets for risk also typically lead to an **efficient allocation of risk**—an allocation of risk in which those who are most willing to bear

Figure 18-3

The Supply of Insurance

This is the supply of insurance policies to provide £1,000 in coverage to a merchant ship that has a 1/10 probability of being lost. Each investor has £1,000 of capital at risk. The lowest possible premium at which a policy is offered is £100, equal to the expected loss, and only a non-risk-averse investor is willing to supply this policy. As the premium increases, investors who are more risk-averse are induced to supply policies to the market, increasing the quantity of policies offered.

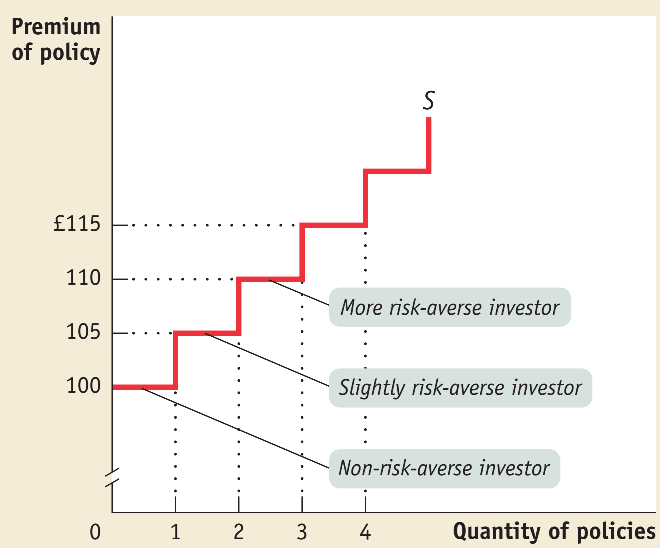

Figure 18-4

The Demand for Insurance

This is the demand for insurance policies for £1,000 in coverage of a merchant ship that has a 1/10 probability of being lost. In this example, the highest premium at which anyone demands a policy is £200, which only the most risk-averse shipowner will desire. As the premium falls, shipowners who are less risk-averse are induced to demand policies, increasing the quantity of policies demanded.

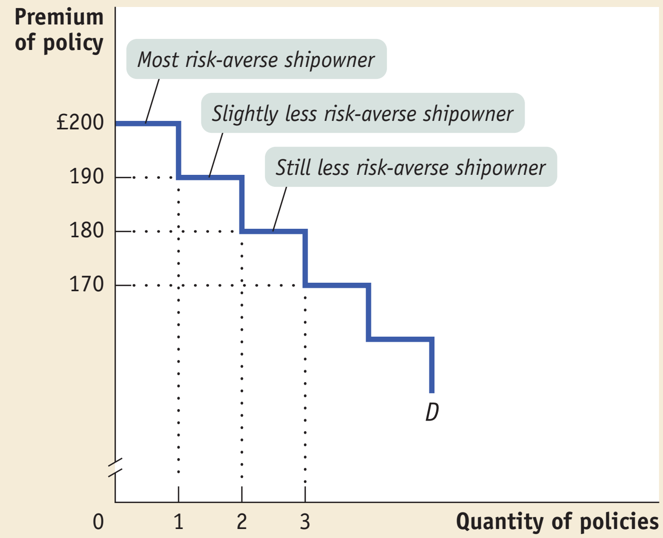

risk are those who end up bearing it. But as in the case of the markets for goods and services, there is an important qualification to this result: there is a well-defined case in which the market for risk fails to achieve efficiency. This arises from the presence of private information, an important topic that we will cover in the next section.

The trading of risk between individuals who differ in their degree of risk aversion plays an extremely important role in the economy, but it is not the only way that markets can help people cope with risk. Under some circumstances, markets can perform a sort of magic trick: they can make some (though rarely all) of the risk that individuals face simply disappear.

Figure 18-5

The Insurance Market

Here we represent the hypothetical market for insuring a merchant ship, where each ship requires £1,000 in coverage. The demand curve is made up of shipowners who wish to buy insurance, and the supply curve is made up of wealthy investors who wish to supply insurance. In this example, at a premium of £200, only the most risk-averse shipowners will purchase insurance; at a premium of £100, only the most non-risk-averse investors are willing to supply insurance. The equilibrium is at a premium of £130 with 5,000 policies bought and sold. In the absence of private information, the market generates an efficient allocation of risk.

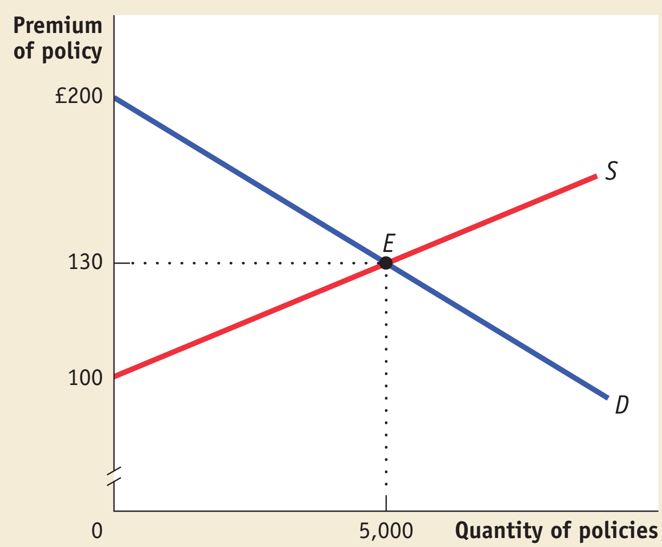

Making Risk Disappear: The Power of Diversification

In the early days of Lloyd's, British merchant ships traversed the world, trading spices and silk from Asia, tobacco and rum from the New World, and textiles and wool from Britain, among many other goods. Each of the many routes that British ships took had its own unique risks—pirates in the Caribbean, gales in the North Atlantic, typhoons in the Indian Ocean.

In the face of all these risks, how were merchants able to survive? One important way was by reducing their risks by not putting all their eggs in one basket: by sending different ships to different destinations, they could reduce the probability that all their ships would be lost. A strategy of investing in such a way as to reduce the probability of severe losses is known as *diversification*. As we'll now see, diversification can often make some of the economy's risk disappear.

Let's stay with our shipping example. It was all too likely that a pirate might seize a merchant ship in the West Indies or that a typhoon might sink another ship in the East Indies. But the key point here is that the various threats to shipping didn't have much to do with each other. So it was considerably less likely that a merchant who had one ship in the West Indies and another ship in the East Indies would lose them both, one to a pirate and the other to a typhoon. After all, there was no connection: the actions of cutthroats in the West had no influence on weather in the East, or vice versa.

Two possible events are **independent events** if each of them is neither more nor less likely to happen if the other one happens.

Statisticians refer to such events—events that have no connection, so that one is no more likely to happen if the other does than if it does not—as **independent events.**

Many unpredictable events are independent of each other. If you toss a coin twice, the probability that it will come up heads on the second toss is the same whether it came up heads or tails on the first toss. If your house burns down today, it does not affect the probability that my house will burn down the same day (unless we live next door to each other or share the services of the same incompetent electrician.)

There is a simple rule for calculating the probability that two independent events will both happen: multiply the probability that one event would happen on its own by the probability that the other event would happen on its own. If you toss a coin once, the probability that it will come up heads is 0.5; if you toss the coin twice, the probability that it will come up heads *both* times is $0.5 \times 0.5 = 0.25$.

But what did it matter to shipowners or Lloyd's investors that ship losses in the Caribbean and ship losses in the Indian Ocean were independent events? The answer is that by spreading their investments across different parts of the world, shipowners or Lloyd's investors could make some of the riskiness of the shipping business simply disappear.

Let's suppose that Joseph Moneypenny, Esq., is wealthy enough to outfit two ships—and let's ignore for a moment the possibility of insuring his cargoes. Should Mr. Moneypenny equip two ships for the West Indian trade and send them off together? Or should he send one ship to Barbados and one to Calcutta?

Assume that both voyages will be equally profitable if successful, yielding £1,000 if the voyage is completed. Also assume that there is 1 chance in 10 both that a ship sent to Barbados will run into a pirate and that a ship sent to Calcutta will be sunk by a typhoon. If Mr. Moneypenny were to send both his ships to either destination, he would face a probability of 0.1 of losing all his investment.

But if Mr. Moneypenny were instead to send one ship east and one west, the probability that he would lose both of them would be only $0.1 \times 0.1 = 0.01$. His expected payoff would be the same—but the chance of losing it all would be much less. So by engaging in **diversification**—investing in several different things, where the possible losses are independent events—he could make some of his risk disappear.

An individual can engage in **diversification** by investing in several different things, so that the possible losses are independent events.

Table 18-2 summarizes Mr. Moneypenny's options and their possible consequences. If he sends both ships to the same destination, he runs a 10% chance of losing them both. If he sends them to different destinations, there are three possible

TABLE 18-2

How Diversification Reduces Risk

(a) If both ships sent to the same destination

Outcome	Probability	Payoff
Both ships arrive	0.9 = 90%	£2,000
Both ships lost	0.1 = 10%	0

(b) If one ship sent east, one west

Outcome	Probability	Payoff
Both ships arrive	0.9 × 0.9 = 81%	£2,000
Both ships lost	0.1 × 0.1 = 1%	0
One ship arrives	9% + 9% = 18%	1,000

outcomes. Both ships could arrive safely: because there is a 0.9 probability of either one making it, the probability that both will make it is $0.9 \times 0.9 = 81\%$. Both could be lost—but the probability of that happening is only $0.1 \times 0.1 = 1\%$. Finally, there are two ways that only one ship can arrive. The probability that the first ship arrives and the second ship is lost is $0.9 \times 0.1 = 9\%$. The probability that the first ship is lost but the second ship arrives is $0.1 \times 0.9 = 9\%$. So the probability that only one ship makes it is $9\% + 9\% = 18\%$.

You might think that diversification is a strategy available only to those with a lot of money to begin with. How could Mr. Moneypenny have diversified if he were able to finance only one ship? But there are ways for even small investors to diversify. Even if Mr. Moneypenny were wealthy enough to equip only one ship, he could enter a partnership with another merchant. They could jointly outfit two ships, agreeing to share the profits equally, and then send those ships to different destinations. That way each would face less risk than if he equipped one ship on his own.

In the modern economy, diversification is made much easier for investors by the fact that they can easily buy shares in many companies by using the *stock market.* The owner of a **share** in a company is the owner of part of that company—typically a very small part, one-millionth or less. An individual who put all of his or her wealth in shares of a single company would lose all of that wealth if the company went bankrupt. But most investors hold shares in many companies, which makes the chance of losing it all very small.

A **share** in a company is a partial ownership of that company.

In fact, Lloyd's of London wasn't just a way to trade risks; it was also a way for investors to diversify. To see how this worked, let's introduce Lady Penelope Smedley-Smythe, a wealthy aristocrat, who decides to increase her income by placing £1,000 of her capital at risk via Lloyd's. She could use that capital to insure just one ship. But more typically she would enter a "syndicate," a group of investors, who would jointly insure a number of ships going to different destinations, agreeing to share the cost if any one of those ships went down. Because it would be much less likely for all the ships insured by the syndicate to sink than for any one of them to go down, Lady Smedley-Smythe would be at much less risk of losing her entire capital.

In some cases, an individual can make risk almost entirely disappear by taking a small share in many

"Your mother called to remind you to diversify."

Pooling is a strong form of diversification in which an individual takes a small share in many independent events. This produces a payoff with very little uncertainty.

independent events. This strategy is known as **pooling.** Consider the case of a health insurance company, which might have millions of policyholders, with thousands of its clients requiring expensive treatment each year. The insurance company can't know whether any given individual will, say, require a heart bypass operation. But heart problems for two different individuals are pretty much independent events. And when there are many possible independent events, it is possible, using statistical analysis, to predict with great accuracy *how many* of those events will happen. For example, if you toss a coin 1,000 times, it will come up heads about 500 times—and it is very unlikely to be more than a percent or two off that figure. So a company offering fire insurance can predict very accurately how many of its clients' homes will burn down in a given year; a company offering health insurance can predict very accurately how many of its clients will need heart surgery in a given year; a life insurance company can predict how many of its clients will . . . Well, you get the idea.

When an insurance company is able to take advantage of the predictability that comes from looking at large numbers of independent events, it is said to engage in pooling of risks. And this pooling often means that even though insurance companies protect people from risk, the owners of the insurance companies may not themselves face much risk.

Lloyd's of London wasn't just a way for wealthy individuals to get paid for taking on some of the risks of less wealthy merchants. It was also a vehicle for pooling some of those risks. The effect of that pooling was to shift the supply curve in Figure 18-5 outward: to make people willing to accept more risk, at a lower price, than would otherwise have been possible.

FOR INQUIRING MINDS

401(k) FOLLIES

It's important to save for your retirement. In the United States, many people save using what is known as a 401(k) plan—a plan that puts aside part of their salary in a special account, which can only be drawn down after retirement. The name comes from a special provision of the tax law, which exempts money set aside in these plans from income tax until the money is withdrawn.

The money in a 401(k) plan can be invested in stocks and other financial assets. Of course, this exposes the worker to the risk of losing some of his or her savings if the value of the stock falls.

The sensible thing is therefore to diversify 401(k) investments. But in 2001 and 2002, when the prices of some stocks plunged, it became clear that in some cases workers had been persuaded to invest most or all of their retirement savings in the stock of the companies that employed them. Most spectacularly, many employees at Enron, a seemingly solid company that suddenly collapsed in 2001, had their entire retirement savings in Enron stock—and were left with nothing. Other stock prices also fell at the time, but a broadly diversified 401(k) would have lost no more than 20 percent of its value.

The Limits of Diversification

Diversification can reduce risk. In some cases it can eliminate it. But these cases are not typical, because there are important limits to diversification. We can see the most important reason for these limits by returning to Lloyd's one more time.

During the period when Lloyd's was creating its legend, there was one important hazard facing British shipping other than pirates or storms: war. Between 1690 and 1815, Britain fought a series of wars, mainly with France (which, among other things, went to war with Britain in support of the American Revolution). Each time, France would sponsor "privateers"—basically pirates with official backing—to raid British shipping and thus indirectly damage Britain's war effort.

Whenever war broke out between Britain and France, losses of British merchant ships would suddenly increase. Unfortunately, merchants could not protect themselves against this eventuality by sending ships to different ports: the privateers would prey on British ships anywhere in the world. So the loss of a ship to French privateers in the West Indies and the loss of another ship to French privateers in the East Indies would *not* be independent events. It would be quite likely that they would happen in the same year.

When an event is more likely to occur if some other event occurs, these two events are said to be **positively correlated.** And like the risk of having a ship seized by French privateers, many financial risks are, alas, positively correlated.

Events are **positively correlated** if each of them is more likely to occur if the other one also occurs.

Here are some of the positively correlated financial risks that investors in the modern world face:

- *Severe weather*. Within any given region of the United States, losses due to weather are definitely not independent events. When a hurricane hits Florida, a lot of Florida homes will suffer hurricane damage. To some extent, insurance companies can diversify away this risk by insuring homes in many states. But events like El Niño (a recurrent temperature anomaly in the Pacific that disrupts weather around the world) can cause simultaneous flooding across the United States.
- *Political events*. Modern governments do not, thankfully, license privateers—although submarines served much the same function during World War II. Even today, however, some kinds of political events—say, a war or revolution in a key raw-material-producing area—can damage business around the globe.
- *Business cycles*. The causes of *business cycles*, fluctuations in the output of the economy as a whole, are a subject for macroeconomics. What we can say here is that if one company suffers a decline in business because of a nationwide economic slump, many other companies will also suffer such declines. So these events will be positively correlated.

When events are positively correlated, the risks they pose cannot be diversified away. An investor can protect herself from the risk that any one company will do badly by investing in many companies; she cannot use the same technique to protect against an economic slump in which *all* companies do badly. An insurance company can protect itself against the risk of losses from local flooding by insuring houses in many different places; but a global weather pattern that produces floods in many places will defeat this strategy.

So institutions like insurance companies and stock markets cannot make risk go away completely. There is always an irreducible core of risk that cannot be diversified. Markets for risk, however, do accomplish two things: First, they enable the economy to eliminate the risk that *can* be diversified. Second, they allocate the risk that remains to the people most willing to bear it.

economics in action

When Lloyd's Almost Llost It

At the end of the 1980s, the venerable institution of Lloyd's found itself in severe trouble. Investors who had placed their capital at risk, believing that the risks were small and the return more or less assured, found themselves called upon to make large payments to satisfy enormous claims. A number of investors, including members of some very aristocratic families, found themselves pushed into bankruptcy.

What happened? Part of the answer is that ambitious managers at Lloyd's had persuaded investors to take on risks that were much larger than the investors realized. (Or to put it a different way, the premiums the investors accepted were too small for the true level of risk contained in the policies.)

But the biggest single problem was that many of the events against which Lloyd's had become a major insurer were *not* independent. In the 1970s and 1980s, Lloyd's had become a major provider of corporate liability insurance in the United States: it protected American corporations against the possibility that they might be sued for selling defective or harmful products. Everyone expected such suits to be more or less independent events. Why should one company's legal problems have much to do with another's?

The answer turned out to lie in one word: asbestos. For decades, the fireproofing material had been used in many products, which meant that many companies were responsible for its use. Then it turned out that asbestos can cause damage to the lungs, especially in children. The result was a torrent of lawsuits by people who believed themselves to have been injured by asbestos and billions of dollars in damage awards—many of them ultimately paid by Lloyd's investors. ■

>>QUICK REVIEW

- There are gains from trade in risk, leading to an *efficient allocation of risk*: those who are most willing to bear risk place their *capital at risk* to cover the financial losses of those least willing to bear risk.
- When *independent events* are involved, a strategy of *diversification* can substantially reduce risk. Diversification is made easier by the existence of institutions like the stock market, in which people trade *shares*. A strong form of diversification, relevant especially to insurance companies, is *pooling*.
- When events are *positively correlated*, there is a core risk that will not go away no matter how much individuals diversify.

>>CHECK YOUR UNDERSTANDING 18-2

1. Explain how each of the following events would change the equilibrium premium and quantity of insurance in the market, indicating any shifts in the supply and demand curves.
   a. An increase in the number of ships traveling the same trade routes and so facing the same kinds of risks
   b. An increase in the number of trading routes, with the same number of ships traveling a greater variety of routes and so facing different kinds of risk
   c. An increase in the degree of risk aversion among the most risk-averse people in the market
   d. An increase in the degree of risk aversion among the least risk-averse people in the market
   e. An increase in the risk affecting the economy as a whole
   f. A fall in the wealth level of the least risk-averse people in the market

Solutions appear at back of book.

Private Information: What You Don't Know Can Hurt You

Markets do very well at dealing with risk due to uncertainty: situations in which nobody knows what is going to happen, whose house will be flooded, or who will get sick. However, markets have much more trouble with situations in which *some people know things that other people don't know*—situations of **private information.** As we will see, private information can distort economic decisions and sometimes prevent mutually beneficial economic transactions from taking place. (Sometimes economists use the term *asymmetric information* rather than *private information*, but they are equivalent.)

Private information is information that some people have that others do not.

Why is some information private? The most important reason is that people generally know more about themselves than other people do. You know whether or not you are a careful driver; but unless you have already been in several accidents, your insurance company does not. You also probably know more about objects you deal with regularly than anyone else. So if you are selling me your used car, you are more likely to be aware of any problems with it than I am.

But why should such differences in who knows what be a problem? It turns out that there are two distinct sources of trouble: those arising from private information about *the way things are—adverse selection*—and those arising from private information about *what people do—moral hazard.*

Adverse Selection: The Economics of Lemons

Suppose that someone offers to sell you an almost brand-new car—purchased just three months ago, with only 2,000 miles on the odometer and no dents or scratches. Will you be willing to pay almost the same for it as for a car direct from the dealer?

Probably not, for one main reason: you cannot help but wonder why this car is being sold. Is it because the owner has discovered that something is wrong with it—that it is a "lemon"? Having driven the car for a while, the owner knows more about it than you do—and people are more likely to sell cars that give them trouble.

You might think that the fact that sellers of used cars know more about them than the buyers represents an advantage to the sellers. But potential buyers know that potential sellers are likely to offer them lemons—they just don't know which car is a lemon. Because potential buyers of a used car know that potential sellers are more likely to sell lemons than good cars, buyers will offer a lower price than they would if they had a guarantee of the car's quality. Worse yet, this bad opinion of used cars tends to be self-reinforcing, precisely because it depresses the prices that buyers offer. Used cars sell at a discount because buyers expect a disproportionate share of those cars to be lemons. Even a used car that is not a lemon would sell only at a large discount, because buyers don't know whether it's a lemon or not. But potential sellers who have good cars are unwilling to sell them at a deep discount, except under exceptional circumstances. So good used cars are rarely offered for sale; and used cars that are offered for sale have a strong tendency to be lemons. (This is why people who have a compelling reason to sell a car, such as moving overseas, make a point of revealing that information to potential buyers—as if to say "This car is not a lemon!")

The end result, then, is not only that used cars sell for low prices and that there are a large number of used cars with hidden problems. Equally important, many potentially beneficial transactions—sales of good cars by people who would like to get rid of them to people who would like to drive them—end up being frustrated by the inability of potential sellers to convince potential buyers that their cars are actually a good value. So some mutually beneficial trade between those who want to sell used cars and those who want to buy them go unexploited.

Although economists sometimes refer to situations like this as the "lemons" problem (the issue was introduced in a famous 1970 paper by economist George Akerlof entitled "The Market for Lemons"), the more formal name of the problem is **adverse selection.** The reason for the name is obvious: because the potential sellers know more about the quality of what they are selling than the potential buyers, they deliberately select the worst things to sell.

Adverse selection occurs when an individual knows more about the way things are than other people do. Private information leads buyers to expect hidden problems in items offered for sale, leading to low prices and so to the best items being kept off the market.

Adverse selection does not apply only to used cars. It is a problem for many parts of the economy—most notably for insurance companies.

Suppose that an insurance company were to offer a standard policy for repairing cars after accidents, with the same premium for all drivers. This premium would reflect the *average* risk of accidents across all drivers. But that would make the policy look very expensive to drivers who know they are particularly careful or skillful and so less likely than the average driver to have an accident. So safe drivers would be less likely than less-safe drivers to buy this policy, leaving the insurance company with exactly the customers it doesn't want: people who are at high risk of having accidents. Mutually beneficial transactions between safe drivers and insurance companies go unexploited because with private information the insurance company cannot set premiums according to the true riskiness of the driver. Safe drivers get "priced out" of the market.

In practice, people or firms faced with the problem of adverse selection follow one of several well-established strategies for dealing with it.

One strategy is **screening:** using observable information to make inferences about private information. Auto insurance provides a very good example. An insurance company may not know whether you are a careful driver, but it has statistical data on the accident rates of people who resemble your profile—and it uses those data in setting premiums. A 19-year-old male who drives a sports car and has already had a fender-bender is likely to pay a very high premium. A 40-year-old female who drives

Adverse selection can be reduced through **screening:** using observable information about people to make inferences about their private information.

a minivan and has never had an accident is likely to pay much less. In some cases, this may be quite unfair: some adolescent males are very careful drivers, and some mature women drive their minivans as if they were F-16s. But nobody can deny that the insurance companies are right on average.

Adverse selection can be diminished by people **signaling** their private information through actions that credibly reveal what they know.

Another strategy is for people who are good prospects to do something **signaling** their private information—taking some action that wouldn't be worth taking unless they were indeed good prospects. Reputable used-car dealers often offer warranties—promises to repair any problems that may come up with the cars they sell. This isn't just a way of insuring their customers against possible expenses; it's a way of credibly showing that they are not selling lemons. As a result, more sales occur, and dealers can command higher prices for their used cars.

A long-term **reputation** allows an individual to reassure others that he or she isn't concealing private information.

Finally, in the face of adverse selection, it can be very valuable to establish a good **reputation:** a used-car dealer will often advertise how long it has been in business to show that it has continued to satisfy its customers. New customers, therefore, will be willing to purchase cars and to pay more for that dealer's used cars.

Moral Hazard

In the late 1970s, New York and other major cities experienced an epidemic of "suspicious" fires—fires that appeared to be deliberately set. Some of the fires were probably started by teenagers on a lark, others by gang members struggling over turf. But investigators eventually became aware of patterns in a number of the fires. Particular landlords, who owned several buildings, seemed to have an unusually large number of their buildings burn down. Although it was difficult to prove, police had few doubts that most of these "fire-prone" landlords were hiring professional arsonists to torch their own properties.

Why burn your own building? These buildings were typically in declining neighborhoods, where rising crime and middle-class flight had led to a decline in property values. But the insurance policies on the buildings were written to compensate owners based on historical property values, and so would actually pay the owner of a destroyed building more than the building was worth in the current market. For an unscrupulous landlord, who knew the right people, this presented a profit opportunity.

The arson epidemic became less severe during the 1980s, partly because insurance companies began making it difficult to overinsure properties, partly because a boom in real estate values made many previously arson-threatened buildings worth more unburned.

The episode makes it clear that it is a bad idea for insurance companies to let customers insure buildings for more than their value—it gives the customers some destructive incentives. You might think, however, that the incentive problem goes away as long as the insurance is no more than 100 percent of the value.

But, unfortunately, anything close to 100 percent insurance still distorts incentives—it induces policyholders to behave differently than they would in the absence of insurance. The reason is that preventing fires requires effort and cost on the part of a building's owner. Fire alarms and sprinkler systems have to be kept in good repair, fire safety rules have to be strictly enforced, and so on. All of this takes time and money—time and money that the owner may not find worth spending if the insurance policy will provide full compensation for any losses.

Of course, the insurance company could specify in the policy that it won't pay if basic safety precautions have not been taken. But it isn't always easy to tell how careful a building's owner has been—the owner knows, but the insurance company does not.

Moral hazard occurs when an individual knows more about his or her own actions than other people do. This leads to a distortion of incentives to take care or to expend effort, especially when the individual is insured.

The point is that the building's owner has private information about his or her own actions, about whether he or she has really taken all appropriate precautions. As a result, the insurance company is likely to face greater claims than if it were able to determine exactly how much effort a building owner extends to prevent a loss. The problem posed by private information an individual has about his own actions is known as **moral hazard**.

To deal with moral hazard, it is necessary to give individuals with private information some personal stake in what happens, a stake that gives them a reason to expend extra effort even if others cannot verify that they have done so. Moral hazard is the reason salespeople in many stores receive a commission on sales: it's hard for managers to be sure how hard the salespeople are really working, and if they were paid only straight salary, they would not have an incentive to exert effort to make those sales. As described in the Economics in Action below, similar logic explains why many stores and restaurants, even if they are part of national chains, are actually "franchises" owned by the people who run them.

Insurance companies deal with moral hazard by requiring a **deductible:** they compensate for losses only above a certain amount, so that coverage is always less than 100 percent. The insurance on your car, for example, may pay for repairs only after the first $500. This means that a careless driver who gets into a fender-bender will end up paying $500 for repairs even if he is insured, which provides at least some incentive to be careful and reduce moral hazard.

A **deductible** in an insurance policy is a sum that the insured individual must pay before being compensated for a claim.

In addition to reducing moral hazard, deductibles provide a partial solution to the problem of adverse selection. Your insurance premium often drops substantially if you are willing to accept a large deductible. This is an attractive option to people who know they are low-risk customers; it is less attractive to people who know they are high-risk—and so likely to have an accident and end up paying that deductible. By offering a menu of policies with different premiums and deductibles, insurance companies can screen their customers, inducing them to sort themselves out on the basis of their private information.

As the example of deductibles suggests, moral hazard limits the ability of the economy to allocate risks efficiently. You generally can't get full (100 percent) insurance on your home or car, even though you would like to buy it, and you bear the risk of large deductibles even though you would prefer not to. The following Economics in Action illustrates how in some cases moral hazard limits the ability of investors to diversify their investments. ■

economics in action

Franchise Owners Try Harder

When Americans go out for a quick meal, they often end up at one of the fast-food chains—McDonald's, Burger King, and so on. Because these are large corporations, most customers probably imagine that the people who serve them are themselves employees of large corporations. But usually they aren't. Most fast-food restaurants—for example, 85 percent of McDonald's outlets—are franchises. That is, some individual has paid the parent company for the right to operate a restaurant selling its product; he or she may look like an arm of a giant company but is in fact a small-business owner.

Becoming a franchisee is not a guarantee of success. You must put up a large amount of money, both to buy the license and to set up the restaurant itself (to open a Taco Bell, for example, cost $1 million in 1997). And although McDonald's takes care that its franchises are not too close to each other, they often face stiff competition from rival chains and even from a few truly independent restaurants. Becoming a franchise owner, in other words, involves taking on quite a lot of risk.

But why should people need to take these risks? Didn't we just learn that it is better to diversify, to spread your wealth among many investments? The logic of diversification would seem to say that it's better for someone with $1 million to invest in a wide range of stocks rather than put it all into one Taco Bell. This implies that Taco Bell would find it hard to attract franchisees. More specifically, nobody would be willing to be a franchisee unless they expected to earn considerably more than they would as a simple hired manager. So wouldn't it be more profitable for McDonald's or Taco Bell simply to hire managers to run their restaurants?

It turns out that it isn't, because the success of a restaurant depends a lot on how hard the manager works, on the effort he or she puts into choosing the right employees, on keeping the place clean and attractive to customers, and so on. Could McDonald's get the right level of effort from a salaried manager? Probably not. The problem is moral hazard: the manager knows whether he or she is really putting 100 percent into the job, but company headquarters does not. So a salaried manager, who would get a salary even without doing everything possible to make the restaurant a success, would not have the incentive to do that extra bit—an incentive the owner would have if he or she had a substantial stake in the success of the restaurant.

In other words, there is a moral hazard problem in running a McDonald's, where the private information is how hard the manager works.

Franchising resolves this problem. A franchisee, whose wealth is tied up in the business and who stands to profit personally from its success, has every incentive to work extremely hard.

The result is that fast-food chains rely mainly on franchisees to operate their restaurants, even though the contracts with these owner-managers allow the franchisees on average to make much more than it would have cost the companies to hire store managers. The higher earnings of franchisees compensate them for the risk they accept. The moral hazard issue—the fact that it is hard to tell just how hard someone is working—prevents the elimination of risk through diversification. ■

>>QUICK REVIEW

- *Private information* can distort incentives and prevent mutually beneficial transactions from occurring. One source is *adverse selection*: sellers have private information about their goods and buyers offer low prices, leading the sellers of quality goods to drop out and leaving the market dominated by "lemons."
- Adverse selection can be reduced by revealing private information through *screening* or *signaling*, or by cultivating a long-term *reputation*.
- Another source of problems is *moral hazard*. In the case of insurance, it leads individuals to exert too little effort to prevent losses. This gives rise to features like *deductibles,* which limit the efficient allocation of risk.

>>CHECK YOUR UNDERSTANDING 18-3

1. Your car insurance premiums are lower if you have had no moving violations for several years. Explain how this feature tends to decrease the potential inefficiency caused by adverse selection.
2. A common feature of home construction contracts is that when it costs more to construct a building than was originally estimated, the contractor must absorb the additional cost. Explain how this feature reduces the problem of moral hazard but also forces the contractor to bear more risk than she would like.
3. True or false? Explain your answer, stating what concept analyzed in this chapter accounts for the feature.

   People with higher deductibles on their auto insurance:

   a. Generally drive more carefully
   b. Pay lower premiums
   c. Are generally wealthier

• A LOOK AHEAD •

We began this chapter with a discussion of yet another thing markets can do right: they allow individuals to trade risk, to their mutual benefit. But our discussion of private information shows that markets have difficulty handling some problems.

In the next two chapters, we will begin discussing in earnest the problem of *market failure*—situations in which markets, left to themselves, can lead to inefficient outcomes. As we'll see, one of the key roles of government is to correct market failure.

SUMMARY

1. The *expected value* of a *random variable* is the weighted average of all possible values, where the weights correspond to the probability of that value occurring.
2. *Risk* is uncertainty about future events or *states of the world*. It is *financial risk* when the uncertainty is about monetary outcomes.
3. Under uncertainty, people maximize *expected utility*. A *risk-averse* person will choose to reduce risk when that reduction leaves the expected value of his or her income or wealth unchanged. A *fair insurance policy* has that feature: the *premium* is equal to the expected value of the claim.
4. Risk aversion arises from diminishing marginal utility: an additional dollar of income generates higher marginal utility in low-income states than in high-income states. A fair insurance policy increases a risk-averse person's utility because it transfers a dollar from a high-income state (a state when no loss occurs) to a low-income state (a state when a loss occurs).
5. Differences in preferences and income or wealth lead to differences in risk aversion. Depending on the size of the premium, a risk-averse person is willing to purchase unfair insurance, a policy for which the premium exceeds the expected value of the claim. The greater your risk aversion, the higher the premium you are willing to pay.
6. There are gains from trade in risk, leading to an *efficient allocation of risk*: those who are most willing to bear risk put their *capital at risk* to cover the losses of those least willing to bear risk.
7. Risk can also be reduced through *diversification*, investing in several different things that correspond to *independent events*. The stock market, where *shares* in companies are traded, offers one way to diversify. Insurance companies can engage in *pooling*, averaging many independent events so as to eliminate almost all risk. But when the underlying events are *positively correlated*, risk cannot be diversified away.
8. *Private information* can cause inefficiency. One problem is *adverse selection*, private information about the way things are. It creates the "lemons" problem in used car markets where sellers of high-quality cars drop out of the market. Adverse selection can be limited in several ways—through *screening* of individuals, through *signaling* that people use to reveal their private information, and through the cultivation of *reputation*.
9. A related problem is *moral hazard:* individuals have private information about their actions, which distorts their incentives and limits the ability of markets to allocate risk efficiently. Insurance companies try to limit moral hazard by imposing *deductibles,* placing more risk on the insured.

KEY TERMS

Random variable, p. 000
Expected value, p. 000
State of the world, p. 000
Risk, p. 000
Financial risk, p. 000
Expected utility, p. 000
Premium, p. 000
Fair insurance policy, p. 000
Risk-averse, p. 000
Capital at risk, p. 000
Efficient allocation of risk, p. 000
Independent events, p. 000
Diversification, p. 000
Share, p. 000
Pooling, p. 000
Positively correlated, p. 000
Private information, p. 000
Adverse selection, p. 000
Screening, p. 000
Signaling, p. 000
Reputation, p. 000
Moral hazard, p. 000
Deductible, p. 000

PROBLEMS

1. For each of the following situations, calculate the expected value.
   a. Sharon owns one share of IBM stock, which is currently trading at $80. There is a 50% chance that the share price will rise to $100 and a 50% chance that the share price will fall to $70. What is the expected value of the share price?
   b. Tanisha buys a ticket in a small lottery. There is a probability of 0.7 that she will win nothing, of 0.2 that she will win $10, and of 0.1 that she will win $50. What is the expected value of Tanisha's winnings?
   c. Fang is a farmer whose rice crop depends on the weather. If the weather is favorable, he will make a profit of $100. If the weather is unfavorable, he will make a profit of –$20 (that is, he will lose money). The weather forecast reports that the probability of weather being favorable is 0.9 and the probability of weather being unfavorable is 0.1. What is the expected value of Fang's profit?
2. Vicky N. Vestor is considering investing some of her money in a start-up company. She currently has income of $4,000, and she is considering investing $2,000 of that in the company. There is a 50% probability that the company will succeed and

will pay out $8,000 to Vicky (her original investment of $2,000 plus $6,000 of the company's profits). And there is a 50% probability that the company will fail and Vicky will get nothing (and lose her investment). The accompanying table illustrates Vicky's utility function.

Income	Total utility (utils)
$0	0
1,000	50
2,000	85
3,000	115
4,000	140
5,000	163
6,000	183
7,000	200
8,000	215
9,000	229
10,000	241

a. Calculate Vicky's marginal utility of income for each income level. Is Vicky risk-averse?

b. Calculate the expected value of Vicky's income from this investment.

c. Calculate Vicky's expected utility from making the investment.

d. What is Vicky's utility from not making the investment? Will Vicky therefore invest in the company?

3. Vicky N. Vestor's utility function was given in Problem 2. As in Problem 2, Vicky currently has income of $4,000. She is considering investing in a start-up company, but the investment now costs $4,000 to make. If the company fails, Vicky will get nothing from the company. But if the company succeeds, she will get $10,000 from the company (her original investment of $4,000 plus $6,000 of the company's profits). Will Vicky invest in the company?

4. You have $1,000 that you can invest. If you buy Ford stock, then, in one year's time: with a probability of 0.2 you will get $1,500; with a probability of 0.4 you will get $1,100; and with a probability of 0.4 you will get $900. If you put the money into the bank, in one year's time you will get $1,100 for certain.

a. What is the expected value of your earnings from investing in Ford stock?

b. Suppose you are risk-averse. Can we say for sure whether you will invest in Ford stock or put your money into the bank?

5. You have $1,000 that you can invest. If you buy General Motors stock, then, in one year's time: with a probability of 0.4 you will get $1,600; with a probability of 0.4 you will get $1,100; and with a probability of 0.2 you will get $800. If you put the money into the bank, in one year's time you will get $1,100 for certain.

a. What is the expected value of your earnings from investing in General Motors stock?

b. Suppose you prefer putting your money into the bank to investing it in General Motors stock. What does that tell us about your attitude to risk?

6. Wilbur is an airline pilot who currently has income of $60,000. If he gets sick and loses his flight medical certificate, he loses his job and has only $10,000 income. His probability of staying healthy is 0.6, and his probability of getting sick is 0.4. Wilbur's utility function is given in the accompanying table.

Income	Total utility (utils)
$0	0
10,000	60
20,000	110
30,000	150
40,000	180
50,000	200
60,000	210

a. What is the expected value of Wilbur's income?

b. What is Wilbur's expected utility?

Wilbur thinks about buying "loss-of-license" insurance that will compensate him if he loses his flight medical certificate.

c. One insurance company offers Wilbur full compensation for his income loss (that is, the insurance company pays Wilbur $50,000 if he loses his flight medical certificate), and it charges a premium of $40,000. That is, regardless of whether he loses his flight medical certificate, Wilbur's income will be $20,000. What is Wilbur's utility? Will he buy the insurance?

d. What is the highest premium Wilbur would just be willing to pay for full insurance (insurance that completely compensates him for the income loss)?

7. Hugh has $5,000 income right now. His utility function is shown in the accompanying table.

Income	Total utility (utils)
$0	0
1,000	100
2,000	140
3,000	166
4,000	185
5,000	200
6,000	212
7,000	222
8,000	230
9,000	236
10,000	240

a. Calculate Hugh's marginal utility of income. What is his attitude toward risk?

b. Hugh is thinking about gambling in a casino. With a probability of 0.5 he loses $4,000, and with a probability of 0.5 he wins $4,000. What is the expected value of Hugh's income? What is Hugh's expected utility? Will he decide to gamble? (Suppose that he gets no extra utility from going to the casino.)

c. Suppose that the "spread" (how much he can win versus how much he can lose) of the gamble narrows, so that with a probability of 0.5 Hugh loses $2,000, and with a probability of 0.5 he wins $2,000. What is the expected value of Hugh's income? What is his expected utility? Is this gamble better for him than the gamble in part b? Will he decide to gamble?

8. Eva is risk-averse. Currently her income is $50,000. She faces the following choice: she can invest in the stock of a dot-com company, or she can invest in IBM stock. If she invests in the dot-com company, then with probability 0.5 she will lose $30,000, but with probability 0.5 she will gain $50,000. If she invests in IBM stock, then with probability 0.5 she will lose only $10,000, but with probability 0.5 she will gain only $30,000. Can you tell which investment she will prefer to make?

9. Suppose you have $1,000 that you can invest in Ted and Larry's Ice Cream Parlor and Ethel's House of Cocoa. The price of a share of stock in either company is $100. The fortunes of each company are closely linked to the weather. When it is warm, the value of Ted and Larry's stock rises to $150 but the value of Ethel's stock falls to $60. When it is cold, the value of Ethel's stock rises to $150 but the value of Ted and Larry's stock falls to $60. There is an equal chance of the weather being warm or cold.

a. If you invest all your money in Ted and Larry's, what is your expected payoff? What if you invested all of your money in Ethel's?

b. Suppose you diversify and invest half of your $1,000 in each company. How much money will you have if the weather is warm? What if it is cold?

c. Suppose you are risk-averse. Would you prefer to put all your money in Ted and Larry's, as in part a? Or would you prefer to diversify, as in part b? Explain your reasoning.

10. You are considering buying a second-hand Volkswagen. From reading car magazines, you know that half of all Volkswagens have problems of some kind (they are "lemons") and the other half run just fine (they are "plums"). If you knew that you were getting a plum, you would be willing to pay $10,000 for it: this is how much a plum is worth to you. You would also be willing to buy a lemon, but only if its price was no more than $4,000: this is how much a lemon is worth to you. And someone who owns a plum would be willing to sell it at any price above $8,000. Someone who owns a lemon would be willing to sell it for any price above $2,000.

a. For now, suppose that you can immediately tell whether the car that you are being offered is a lemon or a plum. Suppose someone offers you a plum. Will there be trade?

Now suppose that the seller has private information about the car she is selling: the seller knows whether she has a lemon or a plum. But when the seller offers you a Volkswagen, you do not know whether it is a lemon or a plum. So this is a situation of adverse selection.

b. Since you do not know whether you are being offered a plum or a lemon, you base your decision on the expected value to you of a Volkswagen: how much the average Volkswagen is worth to you. Calculate this expected value.

c. Suppose, from driving the car, the seller knows she has a plum. However, you don't know whether this particular car is a lemon or a plum, so the most you are willing to pay is your expected value. Will there be trade?

11. You own a company that produces chairs, and you are thinking about hiring one more employee. Each chair produced gives you revenue of $10. There are two potential employees, Fred Ast and Sylvia Low. Fred is a fast worker who produces ten chairs per day, creating revenue for you of $100. Fred knows that he is fast and so will work for you only if you pay him more than $80 per day. Sylvia is a slow worker who produces only five chairs per day, creating revenue for you of $50. Sylvia knows that she is slow and so will work for you if you pay her more than $40 per day. Although Sylvia knows she is slow and Fred knows he is fast, you do not know who is fast and who is slow. So this is a situation of adverse selection.

a. Since you do not know which type of worker you get, you think about what the expected value of your revenue will be if you hire one of the two. What is that expected value?

b. Suppose you offered to pay a daily wage equal to the expected revenue you calculated in part a. Who would you be able to hire: Fred, or Sylvia, or either?

12. For each of the following situations, do the following: First describe whether it is a situation of moral hazard or of adverse selection. Then explain what inefficiency can arise from this situation and explain how the proposed solution reduces the inefficiency.

a. When you buy a second-hand car, you do not know whether it is a lemon (low quality) or a plum (high quality), but the seller knows. A solution is for sellers is to offer a warranty with the car that pays for repair costs.

b. Health maintenance organizations do not know how urgently you need a doctor. A solution is for patients to have to make a co-payment of a certain dollar amount (for example, $10) each time they visit a health care provider.

c. When airlines sell tickets, they do not know whether a buyer is a business traveler (who is willing to pay a lot for a seat), or a leisure traveler (who has a low willingness to pay). A solution is for the airline to offer an expensive ticket that is very flexible (it allows date and route changes) and a cheap ticket that is very inflexible (it has to be booked in advance and cannot be changed).

d. Shareholders in a company do not know whether the CEO in fact maximizes the firm's profit. A solution is for shareholders to pay part of the CEO's salary in stock options, which are worth a lot if the company's profits are high and very little if the company's profits are low.

e. A manager does not know whether workers on an assembly line work hard or whether they shirk. A solution is to pay the workers "piece rates," that is, pay them according to how much they have produced each day.

f. When making a decision about hiring you, prospective employers do not know whether you are a productive or unproductive worker. A solution is for productive workers to acquire more education than unproductive workers.

>web... To continue your study and review of concepts in this chapter, please visit the Krugman/Wells website for quizzes, animated graph tutorials, web links to helpful resources, and more.

www.worthpublishers.com/krugmanwells