CHAPTER EIGHTEEN

Asymmetric Information

Markets may not be fully efficient when one side has information that the other side does not (asymmetric information). Contracts with more complex terms than simple per-unit prices may be used to help solve problems raised by such asymmetric information. The two important classes of asymmetric information problems studied in this chapter include moral hazard problems, in which one party's actions during the term of the contract are unobservable to the other, and adverse selection problems, in which a party obtains asymmetric information about market conditions before signing the contract. Carefully designed contracts may reduce such problems by providing incentives to reveal one's information and take appropriate actions. But these contracts seldom eliminate the inefficiencies entirely. Surprisingly, unbridled competition may worsen private information problems, although a carefully designed auction can harness competitive forces to the auctioneer's advantage.

COMPLEX CONTRACTS AS A RESPONSE TO ASYMMETRIC INFORMATION

So far, the transactions we have studied have involved simple contracts. We assumed that firms bought inputs from suppliers at constant per-unit prices and likewise sold output to consumers at constant per-unit prices. Many real-world transactions involve much more complicated contracts. Rather than an hourly wage, a corporate executive's compensation usually involves complex features such as the granting of stock, stock options, and bonuses. Insurance policies may cap the insurer's liability and may require the customer to bear costs in the form of deductibles and copayments. In this chapter, we will show that such complex contracts may arise as a way for transacting parties to deal with the problem of asymmetric information.

Asymmetric information

Transactions can involve a considerable amount of uncertainty. The value of a snow shovel will depend on how much snow falls during the winter season. The value of a hybrid car will depend on how much gasoline prices increase in the future. Uncertainty need not lead to inefficiency when both sides of a transaction have the same limited knowledge concerning the future, but it can lead to inefficiency when one side has better information. The side with better information is said to have *private information* or, equivalently, *asymmetric information*.

There are several sources of asymmetric information. Parties will often have "inside information" concerning themselves that the other side does not have. Consider the case of health insurance. A customer seeking insurance will often have private information about his or her own health status and family medical history that the insurance company does not. Consumers in good health may not bother to purchase health insurance at the prevailing rates. A consumer in poor health would have higher demand for insurance, wishing to shift the burden of large anticipated medical expenses to the insurer. A medical examination may help the insurer learn about a customer's health status, but examinations are costly and may not reveal all of the customer's private health information. The customer will be reluctant to report family medical history and genetic disease honestly if the insurer might use this information to deny coverage or increase premiums.

Other sources of asymmetric information arise when what is being bought is an agent's service. The buyer may not always be able to monitor how hard and well the agent is working. The agent may have better information about the requirements of the project because of his or her expertise, which is the reason the agent was hired in the first place. For example, a repairer called to fix a kitchen appliance will know more about the true severity of the appliance's mechanical problems than does the homeowner.

Asymmetric information can lead to inefficiencies. Insurance companies may offer less insurance and charge higher premiums than if they could observe the health of potential clients and could require customers to obey strict health regimens. The whole market may unravel as consumers who expect their health expenditures to be lower than the average insured consumer's withdraw from the market in successive stages, leaving only the few worst health risks as consumers. With appliance repair, the repairer may pad his or her bill by replacing parts that still function and may take longer than needed—a waste of resources.

The value of contracts

Contractual provisions can be added in order to circumvent some of these inefficiencies. An insurance company can offer lower health insurance premiums to customers who submit to medical exams or who are willing to bear the cost of some fraction of their own medical services. Lower-risk consumers may be more willing than high-risk consumers to submit to medical exams and to bear a fraction of their medical expenses. A homeowner may buy a service contract that stipulates a fixed fee for keeping the appliance running rather than a payment for each service call and part needed in the event the appliance breaks down.

Although contracts may help reduce the inefficiencies associated with asymmetric information, rarely do they eliminate the inefficiencies altogether. In the health insurance example, having some consumers undertake a medical exam requires the expenditure of real resources. Requiring low-risk consumers to bear some of their own medical expenditures means that they are not fully insured, which is a social loss to the extent that a risk-neutral insurance company would be a more efficient risk bearer than a risk-averse consumer. A fixed-fee contract to maintain an appliance may lead the repairer to supply too little effort, overlooking potential problems in the hope that nothing breaks until after the service contract expires (and so then the problems become the homeowner's).

PRINCIPAL-AGENT MODEL

Models of asymmetric information can quickly become quite complicated, and thus, before considering a full-blown market model with many suppliers and demanders, we will devote much of our analysis to a simpler model—called a principal-agent model—in which there is only one party on each side of the market. The party who proposes the contract is called the *principal*. The party who decides whether or not to accept the contract and then performs under the terms of the contract (if accepted) is called the *agent*. The agent is typically the party with the asymmetric information. We will use "she" for the principal and "he" for the agent to facilitate the exposition.

Two leading models

Two models of asymmetric information are studied most often. In a first model, the agent's actions taken during the term of the contract affect the principal, but the principal does not observe these actions directly. The principal may observe outcomes that are correlated with the agent's actions but not the actions themselves. This first model is called a hidden-action model. For historical reasons stemming from the insurance context, the hidden-action model is also called a *moral hazard model*.

In a second model, the agent has private information about the state of the world before signing the contract with the principal. The agent's private information is called his *type*, consistent with our terminology from games of private information studied in Chapter 8. The second model is thus called a hidden-type model. For historical reasons stemming from its application in the insurance context, which we discuss later, the hidden-type model is also called an *adverse selection model*.

As indicated by Table 18.1, the hidden-type and hidden-action models cover a wide variety of applications. Note that the same party might be a principal in one setting and an agent in another. For example, a company's CEO is the principal in dealings with the company's employees but is the agent of the firm's shareholders. We will study several of the applications from Table 18.1 in detail throughout the remainder of this chapter.

First, second, and third best

In a full-information environment, the principal could propose a contract to the agent that maximizes their joint surplus and captures all of this surplus for herself, leaving the agent with just enough surplus to make him indifferent between signing the contract or not. This outcome is called the *first best*, and the contract implementing this outcome is called the *first-best contract*. The first best is a theoretical benchmark that is unlikely to be achieved in practice because the principal is rarely fully informed. The outcome that maximizes the principal's surplus subject to the constraint that the principal is less well informed than the agent is called the *second best*, and the contract that implements this

TABLE 18.1 APPLICATIONS OF THE PRINCIPAL-AGENT MODEL

		Agent's Private Information	
Principal	**Agent**	**Hidden Type**	**Hidden Action**
Shareholders	Manager	Managerial skill	Effort, executive decisions
Manager	Employee	Job skill	Effort
Homeowner	Appliance repairer	Skill, severity of appliance malfunction	Effort, unnecessary repairs
Student	Tutor	Subject knowledge	Preparation, patience
Monopoly	Customer	Value for good	Care to avoid breakage
Health insurer	Insurance purchaser	Preexisting condition	Risky activity
Parent	Child	Moral fiber	Delinquency

outcome is called the *second-best contract*. Adding further constraints to the principal's problem besides the informational constraint—for example, restricting contracts to some simple form such as constant per-unit prices—leads to the third best, the fourth best, and so on, depending on how many constraints are added.

Since this chapter is in the part of the book that examines market failures, we will be interested in determining how important a market failure is asymmetric information. Comparing the first to the second best will allow us to quantify the reduction in total welfare due to asymmetric information.

Also illuminating is a comparison of the second and third best. This comparison will indicate how surpluses are affected when moving from simple contracts in the third best to potentially quite sophisticated contracts in the second best. Of course, the principal's surplus cannot decrease when she has access to a wider range of contracts with which to maximize her surplus. However, total welfare—the sum of the principal's and agent's surplus in a principal-agent model—may decrease. Figure 18.1 suggests why. In the example in panel (a) of the figure, the complex contract increases the total welfare "pie" that is divided between the principal and the agent. The principal likes the complex contract because it allows her to obtain a roughly constant share of a bigger pie. In panel (b), the principal likes the complex contract even though the total welfare pie is smaller with it

FIGURE 18.1

The Contracting "Pie"

The total welfare is the area of the circle ("pie"); the principal's surplus is the area of the shaded region. In panel (a), the complex contract increases total welfare and the principal's surplus along with it because she obtains a constant share. In panel (b), the principal offers the complex contract—even though this reduces total welfare—because the complex contract allows her to appropriate a larger share.

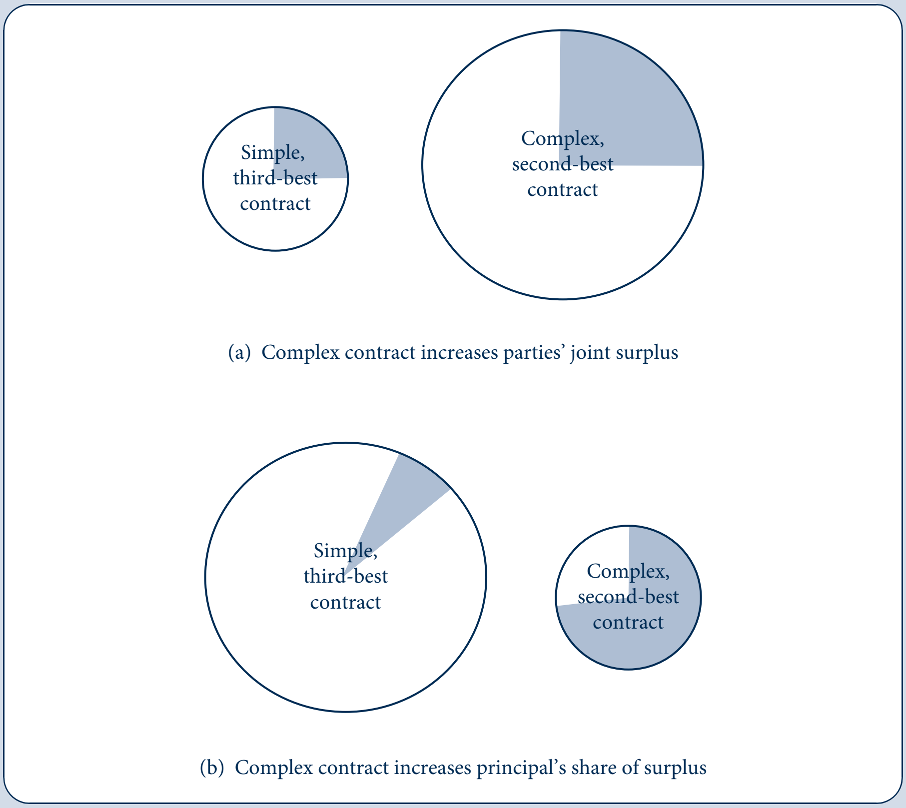

than with the simple contract. The complex contract allows her to appropriate a larger slice at the expense of reducing the pie's total size. The different cases in panels (a) and (b) will come up in the applications analyzed in subsequent sections.

HIDDEN ACTIONS

The first of the two important models of asymmetric information is the hidden-action model, also sometimes called the moral hazard model in insurance and other contexts. The principal would like the agent to take an action that maximizes their joint surplus (and given that the principal makes the contract offer, she would like to appropriate most of the surplus for herself). In the application to the owner-manager relationship that we will study, the owner would like the manager whom she hires to show up during business hours and work diligently. In the application to the accident insurance, the insurance company would like the insured individual to avoid accidents. The agent's actions may be unobservable to the principal. Observing the action may require the principal to monitor the agent at all times, and such monitoring may be prohibitively expensive. If the agent's action is unobservable, then he will prefer to shirk, choosing an action to suit himself rather than the principal. In the owner-manager application, shirking might mean showing up late for work and slacking off while on the job; in the insurance example, shirking might mean taking more risk than the insurance company would like.

Although contracts cannot prevent shirking directly by tying the agent's compensation to his action—because his action is unobservable—contracts can mitigate shirking by tying compensation to observable outcomes. In the owner-manager application, the relevant observable outcome might be the firm's profit. The owner may be able to induce the manager to work hard by tying the manager's pay to the firm's profit, which depends on the manager's effort. The insurance company may be able to induce the individual to take care by having him bear some of the cost of any accident.

Often, the principal is more concerned with the observable outcome than with the agent's unobservable action anyway, so it seems the principal should do just as well by conditioning the contract on outcomes as on actions. The problem is that the outcome may depend in part on random factors outside of the agent's control. In the owner-manager application, the firm's profit may depend on consumer demand, which may depend on unpredictable economic conditions. In the insurance application, whether an accident occurs depends in part on the care taken by the individual but also on a host of other factors, including other people's actions and acts of nature. Tying the agent's compensation to partially random outcomes exposes him to risk. If the agent is risk averse, then this exposure causes disutility and requires the payment of a risk premium before he will accept the contract (see Chapter 7). In many applications, the principal is less risk averse and thus is a more efficient risk bearer than the agent. In the owner-manager application, the owner might be one of many shareholders who each hold only a small share of the firm in a diversified portfolio. In the insurance application, the company may insure a large number of agents, whose accidents are uncorrelated, and thus face little aggregate risk. If there were no issue of incentives, then the agent's compensation should be independent of risky outcomes, completely insuring him against risk and shifting the risk to the efficient bearer: the principal. The second-best contract strikes the optimal balance between incentives and insurance, but it does not provide as strong incentives or as full insurance as the first-best contract.

In the following sections, we will study two specific applications of the hidden-action model. First, we will study employment contracts signed between a firm's owners and a manager who runs the firm on behalf of the owners. Second, we will study contracts offered by an insurance company to insure an individual against accident risk.

OWNER-MANAGER RELATIONSHIP

Modern corporations may be owned by millions of dispersed shareholders who each own a small percentage of the corporation's stock. The shareholders—who may have little expertise in the line of business and who may own too little of the firm individually to devote much attention to it—delegate the operation of the firm to a managerial team consisting of the chief executive officer (CEO) and other officers. We will simplify the setting and suppose that the firm has one representative owner and one manager. The owner, who plays the role of the principal in the model, offers a contract to the manager, who plays the role of the agent. The manager decides whether to accept the employment contract and, if so, how much effort $e \geq 0$ to exert. An increase in e increases the firm's gross profit (not including payments to the manager) but is personally costly to the manager.[1]

Assume the firm's gross profit π_g takes the following simple form:

$$\pi_g = e + \varepsilon. \tag{18.1}$$

Gross profit is increasing in the manager's effort e and also depends on a random variable ε, which represents demand, cost, and other economic factors outside of the manager's control. Assume that ε is normally distributed with mean 0 and variance σ^2. The manager's personal disutility (or cost) of undertaking effort $c(e)$ is increasing $[c'(e) > 0]$ and convex $[c''(e) > 0]$.

Let s be the salary—which may depend on effort and/or gross profit, depending on what the owner can observe—offered as part of the contract between the owner and manager. Because the owner represents individual shareholders who each own a small share of the firm as part of a diversified portfolio, we will assume that she is risk neutral. Letting net profit π_n equal gross profit minus payments to the manager,

$$\pi_n = \pi_g - s, \tag{18.2}$$

the risk-neutral owner wants to maximize the expected value of her net profit:

$$E(\pi_n) = E(e + \varepsilon - s) = e - E(s). \tag{18.3}$$

To introduce a trade-off between incentives and risk, we will assume the manager is risk averse; in particular, we assume the manager has a utility function with respect to salary whose constant absolute risk aversion parameter is $A > 0$. We can use the results from Example 7.3 to show that his expected utility is

$$E(U) = E(s) - \frac{A}{2}\mathrm{Var}(s) - c(e). \tag{18.4}$$

We will examine the optimal salary contract that induces the manager to take appropriate effort e under different informational assumptions. We will study the first-best contract, when the owner can observe e perfectly, and then the second-best contract, when there is asymmetric information about e.

First best (full-information case)

With full information, it is relatively easy to design an optimal salary contract. The owner can pay the manager a fixed salary s^* if he exerts the first-best level of effort e^* (which we will compute shortly) and nothing otherwise. The manager's expected utility from the contract can be found by substituting the expected value $[E(s^*) = s^*]$ and variance $[\mathrm{Var}(s^*) = 0]$ of the fixed salary as well as the effort e^* into Equation 18.4. For the

[1]Besides effort, *(e)* could represent distasteful decisions such as firing unproductive workers.

manager to accept the contract, this expected utility must exceed what he would obtain from his next-best job offer:

$$E(U) = s^* - c(e^*) \geq 0, \tag{18.5}$$

where we have assumed for simplicity that he obtains 0 from his next-best job offer. In principal-agent models, a condition like Equation 18.5 is called a *participation constraint*, ensuring the agent's participation in the contract.

The owner optimally pays the lowest salary satisfying Equation 18.5: $s^* = c(e^*)$. The owner's net profit then is

$$E(\pi_n) = e^* - E(s^*) = e^* - c(e^*), \tag{18.6}$$

which is maximized for e^* satisfying the first-order condition

$$c'(e^*) = 1. \tag{18.7}$$

At an optimum, the marginal cost of effort, $c'(e^*)$, equals the marginal benefit, 1.

Second best (hidden-action case)

If the owner can observe the manager's effort, then she can implement the first best by simply ordering the manager to exert the first-best effort level. If she cannot observe effort, the contract cannot be conditioned on e. However, she can still induce the manager to exert some effort if the manager's salary depends on the firm's gross profit. The manager is given performance pay: the more the firm earns, the more the manager is paid.

Suppose the owner offers a salary to the manager that is linear in gross profit:

$$s(\pi_g) = a + b\pi_g, \tag{18.8}$$

where a is the fixed component of salary and b measures the slope, sometimes called the *power*, of the incentive scheme. If $b = 0$, then the salary is constant and, as we saw, provides no effort incentives. As b increases toward 1, the incentive scheme provides increasingly powerful incentives. The fixed component a can be thought of as the manager's base salary and b as the incentive pay in the form of stocks, stock options, and performance bonuses.

The owner-manager relationship can be viewed as a three-stage game. In the first stage, the owner sets the salary, which amounts to choosing a and b. In the second stage, the manager decides whether or not to accept the contract. In the third stage, the manager decides how much effort to exert conditional on accepting the contract. We will solve for the subgame-perfect equilibrium of this game by using backward induction, starting with the manager's choice of e in the last stage and taking as given that the manager was offered salary scheme $a + b\pi_g$ and accepted it. Substituting from Equation 18.8 into Equation 18.4, the manager's expected utility from the linear salary is

$$E(a + b\pi_g) - \frac{A}{2}\mathrm{Var}(a + b\pi_g) - c(e). \tag{18.9}$$

Reviewing a few facts about expectations and variances of a random variable will help us simplify Equation 18.9. First note that

$$E(a + b\pi_g) = E(a + be + b\varepsilon) = a + be + bE(\varepsilon) = a + be; \tag{18.10}$$

see Equation 2.179. Furthermore,

$$\mathrm{Var}(a + b\pi_g) = \mathrm{Var}(a + be + b\varepsilon) = b^2\ \mathrm{Var}(\varepsilon) = b^2\sigma^2; \tag{18.11}$$

see Equation 2.186. Therefore, Equation 18.9 reduces to

$$\text{manager's expected utility} = a + be - \frac{Ab^2\sigma^2}{2} - c(e). \tag{18.12}$$

The first-order condition for the e maximizing the manager's expected utility yields

$$c'(e) = b. \tag{18.13}$$

Because $c(e)$ is convex, the marginal cost of effort $c'(e)$ is increasing in e. Hence, as shown in Figure 18.2, the higher is the power b of the incentive scheme, the more effort e the manager exerts. The manager's effort depends only on the slope, b, and not on the fixed part, a, of his incentive scheme.

Now fold the game back to the manager's second-stage choice of whether to accept the contract. The manager accepts the contract if his expected utility in Equation 18.12 is non-negative or, upon rearranging, if

$$a \geq c(e) + \frac{Ab^2\sigma^2}{2} - be. \tag{18.14}$$

The fixed part of the salary, a, must be high enough for the manager to accept the contract.

Next, fold the game back to the owner's first-stage choice of the parameters a and b of the salary scheme. The owner's objective is to maximize her expected surplus, which (upon substituting from Equation 18.10 into 18.3) is

$$\text{owner's surplus} = e(1 - b) - a, \tag{18.15}$$

subject to two constraints. The first constraint (Equation 18.14) is that the manager must accept the contract in the second stage. As mentioned in the previous section, this is called a

FIGURE 18.2

Manager's Effort Responds to Increased Incentives

Because the manager's marginal cost of effort, $c'(e)$, slopes upward, an increase in the power of the incentive scheme from b_1 to b_2 induces the manager to increase his effort from e_1 to e_2.

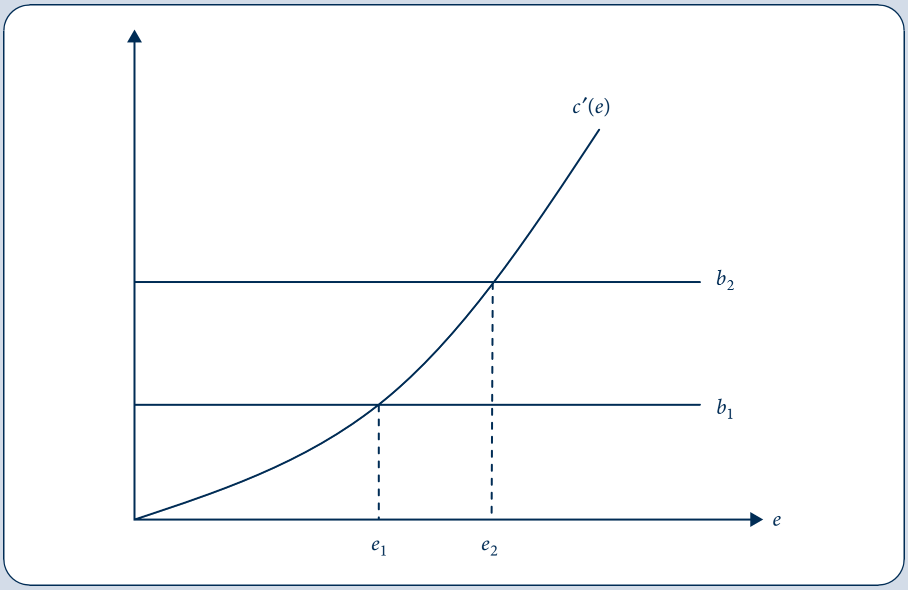

participation constraint. Although Equation 18.14 is written as an inequality, it is clear that the owner will keep lowering a until the condition holds with equality, since a does not affect the manager's effort and since the owner does not want to pay the manager any more than necessary to induce him to accept the contract. The second constraint (Equation 18.13) is that the manager will choose e to suit himself rather than the owner, who cannot observe e. This is called the *incentive compatibility constraint*. Substituting the constraints into Equation 18.15 allows us to express the owner's surplus as a function only of the manager's effort:

$$e - c(e) - \frac{A\sigma^2[c'(e)]^2}{2}. \tag{18.16}$$

The second-best effort e^{**} satisfies the first-order condition

$$c'(e^{**}) = \frac{1}{1 + A\sigma^2 c''(e^{**})}. \tag{18.17}$$

The right-hand side of Equation 18.17 is also equal to the power b^{**} of the incentive scheme in the second best, since $c'(e^{**}) = b^{**}$ by Equation 18.13.

The second-best effort is less than 1 and thus is less than the first-best effort $e^* = 1$. The presence of asymmetric information leads to lower equilibrium effort. If the owner cannot specify e in a contract, then she can induce effort only by tying the manager's pay to firm profit; however, doing so introduces variation into his pay for which the risk-averse manager must be paid a risk premium. This risk premium (the third term in Equation 18.16) adds to the owner's cost of inducing effort.

If effort incentives were not an issue, then the risk-neutral owner would be better-off bearing all risk herself and insuring the risk-averse manager against any fluctuations in profit by offering a constant salary, as we saw in the first-best problem. Yet if effort is unobservable then a constant salary will not provide any incentive to exert effort. The second-best contract trades off the owner's desire to induce high effort (which would come from setting b close to 1) against her desire to insure the risk-averse manager against variations in his salary (which would come from setting b close to 0). Hence the resulting value of b^{**} falls somewhere between 0 and 1.

In short, the fundamental trade-off in the owner-manager relationship is between incentives and insurance. The more risk averse is the manager (i.e., the higher is A), the more important is insurance relative to incentives. The owner insures the manager by reducing the dependence of his salary on fluctuating profit, reducing b^{**} and therefore e^{**}. For the same reason, the more that profit varies owing to factors outside of the manager's control (i.e., the higher is σ^2), the lower is b^{**} and e^{**}.[2]

EXAMPLE 18.1 Owner-Manager Relationship

As a numerical example of some of these ideas, suppose the manager's cost of effort has the simple form $c(e) = e^2/2$ and suppose $\sigma^2 = 1$.

First best. The first-best level of effort satisfies $c'(e^*) = e^* = 1$. A first-best contract specifies that the manager exerts first-best effort $e^* = 1$ in return for a fixed salary of 1/2, which leaves

[2]A study has confirmed that CEOs and other top executives receive more powerful incentives if they work for firms with less volatile stock prices. See R. Aggarwal and A. Samwick, "The Other Side of the Trade-off: The Impact of Risk on Executive Compensation," *Journal of Political Economy* 107 (1999): 65–105.

the manager indifferent between accepting the contract and pursuing his next-best available job (which we have assumed provides him with utility 0). The owner's net profit equals 1/2.

Second best. The second-best contract depends on the degree of the manager's risk aversion measured by A. Suppose first that $A = 1$.[3] Then, by Equation 18.17, the second-best level of effort is $e^{**} = 1/2$, and $b^{**} = 1/2$ as well. To compute the fixed part a^{**} of the manager's salary, recall that Equation 18.14 holds as an equality in the second best and substitute the variables computed so far, yielding $a^{**} = 0$. The manager receives no fixed pay but does receive incentive pay equal to 50 cents for every dollar of gross profit. Substituting the variables computed into Equation 18.15, we see that the owner's expected net profit is 1/4.

Now suppose $A = 2$, so that the manager is more risk averse. The second-best effort decreases to $e^{**} = 1/3$, and b^{**} decreases to 1/3 as well. The fixed part of the manager's salary increases to $a^{**} = 1/18$. The owner's expected net profit decreases to 1/6.

Empirical evidence. In an influential study of performance pay, Jensen and Murphy estimated that $b = 0.003$ for top executives in a sample of large U.S. firms, which is orders of magnitude smaller than the values of b^{**} we just computed.[4] The fact that real-world incentive schemes are less sensitive to performance than theory would indicate is a puzzle for future research to unravel.

QUERY: How would the analysis change if the owners did not perfectly observe gross profit but instead depended on the manager for a self-report? Could this explain the puzzle that top executives' incentives are unexpectedly low-powered?

Comparison to standard model of the firm

It is natural to ask how the results with hidden information about the manager's action compare to the standard model of a perfectly competitive market with no asymmetric information. First, the presence of hidden information raises a possibility of shirking and inefficiency that is completely absent in the standard model. The manager does not exert as much effort as he would if effort were observable. Even if the owner does as well as she can in the presence of asymmetric information to provide incentives for effort, she must balance the benefits of incentives against the cost of exposing the manager to too much risk.

Second, although the manager can be regarded as an input like any other (capital, labor, materials, and so forth) in the standard model, he becomes a unique sort of input when his actions are hidden information. It is not enough to pay a fixed unit price for this input as a firm would the rental rate for capital or the market price for materials. How productive the manager is depends on how his compensation is structured. The same can be said for any sort of labor input: workers may shirk on the job unless monitored or given incentives not to shirk.

MORAL HAZARD IN INSURANCE

Another important context in which hidden actions lead to inefficiencies is the market for insurance. Individuals can take a variety of actions that influence the probability that a risky event will occur. Car owners can install alarms to deter theft; consumers can eat healthier foods to prevent illness. In these activities, utility-maximizing individuals will

[3]To make the calculations easier, we have scaled A up from its more realistic values in Chapter 7 and have rescaled several other parameters as well.

[4]M. Jensen and K. Murphy, "Performance Pay and Top-Management Incentives," *Journal of Political Economy* 98 (1990): 5–64.

pursue risk reduction up to the point at which marginal gains from additional precautions are equal to the marginal cost of these precautions.

In the presence of insurance coverage, however, this calculation may change. If a person is fully insured against losses, then he or she will have a reduced incentive to undertake costly precautions, which may increase the likelihood of a loss occurring. In the automobile insurance case, for example, a person who has a policy that covers theft may not bother to install a car alarm. This behavioral response to insurance coverage is termed *moral hazard*.

DEFINITION

Moral hazard. The effect of insurance coverage on an individual's precautions, which may change the likelihood or size of losses.

The use of the term "moral" to describe this response is perhaps unfortunate. There is nothing particularly "immoral" about the behavior being described, since individuals are simply responding to the incentives they face. In some applications, this response might even be desirable. For example, people with medical insurance may be encouraged to seek early treatment because the insurance reduces their out-of-pocket cost of medical care. But, because insurance providers may find it costly to measure and evaluate such responses, moral hazard may have important implications for the allocation of resources. To examine these, we need a model of utility-maximizing behavior by insured individuals.

Mathematical model

Suppose a risk-averse individual faces the possibility of incurring a loss (l) that will reduce his initial wealth (W_0). The probability of loss is π. An individual can reduce π by spending more on preventive measures (e).[5] Let $U(W)$ be the individual's utility given wealth W.

An insurance company (here playing the role of principal) offers an insurance contract involving a payment x to the individual if a loss occurs. The premium for this coverage is p. If the individual takes the coverage, then his wealth in state 1 (no loss) and state 2 (loss) are

$$\begin{aligned} W_1 &= W_0 - e - p \quad \text{and} \\ W_2 &= W_0 - e - p - l + x, \end{aligned} \tag{18.18}$$

and his expected utility is

$$(1-\pi)U(W_1) + \pi U(W_2). \tag{18.19}$$

The risk-neutral insurance company's objective is to maximize expected profit:

$$\text{expected insurance profit} = p - \pi x. \tag{18.20}$$

First-best insurance contract

In the first-best case, the insurance company can perfectly monitor the agent's precautionary effort e. It sets e and the other terms of the insurance contract (x and p) to maximize its expected profit subject to the participation constraint that the individual accepts the contract:

$$(1-\pi)U(W_1) + \pi U(W_2) \geq \bar{U}, \tag{18.21}$$

[5]For consistency, we use the same variable e as we did for managerial effort. In this context, since e is subtracted from the individual's wealth, e should be thought of as either a direct expenditure or the monetary equivalent of the disutility of effort.

where $\bar{U}$ is the highest utility the individual can attain in the absence of insurance. It is clear that the insurance company will increase the premium until the participation constraint holds with equality. Thus, the first-best insurance contract is the solution to a maximization problem subject to an equality constraint, which we can use Lagrange methods to solve. The associated Lagrangian is

$$\mathcal{L} = p - \pi x + \lambda[(1-\pi)U(W_1) + \pi U(W_2) - \bar{U}]. \tag{18.22}$$

The first-order conditions are

$$0 = \frac{\partial \mathcal{L}}{\partial p} = 1 - \lambda[(1-\pi)U'(W_0 - e - p) + \pi U'(W_0 - e - l + x)], \tag{18.23}$$

$$0 = \frac{\partial \mathcal{L}}{\partial x} = -\pi + \lambda \pi U'(W_0 - e - p - l + x), \tag{18.24}$$

$$\begin{aligned} 0 = \frac{\partial \mathcal{L}}{\partial e} = -\frac{\partial \pi}{\partial e}x - \lambda\Big\{&(1-\pi)U'(W_0 - e - p) + \pi U'(W_0 - e - p - l + x) \\ &+ \frac{\partial \pi}{\partial e}[U(W_0 - e - p) - U(W_0 - e - p - l + x)]\Big\}. \end{aligned} \tag{18.25}$$

These conditions may seem complicated, but they have simple implications. Equations 18.23 and 18.24 together imply

$$\begin{aligned} \frac{1}{\lambda} &= (1-\pi)U'(W_0 - e - p) + \pi U'(W_0 - e - p - l + x) \\ &= U'(W_0 - e - p - l + x), \end{aligned} \tag{18.26}$$

which in turn implies $x = l$. This is the familiar result that the first best involves full insurance. Substituting for λ from Equation 18.26 into Equation 18.25 and noting $x = l$, we have

$$-\frac{\partial \pi}{\partial e} l = 1. \tag{18.27}$$

At an optimum, the marginal social benefit of precaution (the reduction in the probability of a loss multiplied by the amount of the loss) equals the marginal social cost of precaution (which here is just 1). In sum, the first-best insurance contract provides the individual with full insurance but requires him to choose the socially efficient level of precaution.

Second-best insurance contract

To obtain the first best, the insurance company would need to monitor the insured individual to ensure that the person was constantly taking the first-best level of precaution, e^*. In the case of insurance for automobile accidents, the company would have to make sure that the driver never exceeds a certain speed, always keeps alert, and never drives while talking on his cell phone, for example. Even if a black-box recorder could be installed to constantly track the car's speed, it would still be impossible to monitor the driver's alertness. Similarly, for health insurance, it would be impossible to watch everything the insured party eats to make sure he doesn't eat anything unhealthy.

Assume for simplicity that the insurance company cannot monitor precaution e at all, so that e cannot be specified by the contract directly. This second-best problem is similar to the first-best except that a new constraint must to be added: an incentive compatibility constraint specifying that the agent is free to choose the level of precaution that suits him and maximizes his expected utility,

$$(1-\pi)U(W_1) + \pi U(W_2). \tag{18.28}$$

Unlike the first best, the second-best contract will typically not involve full insurance. Under full insurance, $x = l$ and (as Equation 18.18 shows) $W_1 = W_2$. But then the insured party's expected utility from Equation 18.28 is

$$U(W_1) = U(W_0 - e - p), \tag{18.29}$$

which is maximized by choosing the lowest level of precaution possible, $e = 0$.

To induce the agent to take precaution, the company should provide him only partial insurance. Exposing the individual to some risk induces him to take at least some precaution. The company will seek to offer just the right level of partial insurance: not too much insurance (else the agent's precaution drops too low) and not too little insurance (else the agent would not be willing to pay much in premiums). The principal faces the same trade-off in this insurance example as in the owner-manager relationship studied previously: incentives versus insurance.

The solution for the optimal second-best contract is quite complicated, given the general functional forms for utility that we are using.[6] Example 18.2 provides some further practice on the moral hazard problem with specific functional forms.

EXAMPLE 18.2 Insurance and Precaution against Car Theft

In Example 7.2 we examined the decision by a driver endowed with \$100,000 of wealth to purchase insurance against the theft of a \$20,000 car. Here we reexamine the market for theft insurance when he can also take the precaution of installing a car alarm that costs \$1,750 and that reduces the probability of theft from 0.25 to 0.15.

No insurance. In the absence of insurance, the individual can decide either not to install the alarm, in which case (as we saw from Example 7.2) his expected utility is 11.45714, or to install the alarm, in which case his expected utility is

$$0.85 \ln(100{,}000 - 1{,}750) + 0.15 \ln(100{,}000 - 1{,}750 - 20{,}000) = 11.46113. \tag{18.30}$$

He prefers to install the device.

First best. The first-best contract maximizes the insurance company's profit given that it requires the individual to install an alarm and can costlessly verify whether the individual has complied. The first-best contract provides full insurance, paying the full \$20,000 if the car is stolen. The highest premium p that the company can charge leaves the individual indifferent between accepting the full-insurance contract and going without insurance:

$$\ln(100{,}000 - 1{,}750 - p) = 11.46113. \tag{18.31}$$

Solving for p yields

$$98{,}250 - p = e^{11.46113}, \tag{18.32}$$

implying that $p = 3{,}298$. (Note that the e in Equation 18.32 is the number 2.7818…, not the individual's precaution.) The company's profit equals the premium minus the expected payout: $3{,}298 - (0.15 \times 20{,}000) = \$298$.

Second best. If the company cannot monitor whether the individual has installed an alarm, then it has two choices. It can induce him to install the alarm by offering only partial insurance, or it can disregard the alarm and provide him with full insurance.

[6] For more analysis see S. Shavell, "On Moral Hazard and Insurance," *Quarterly Journal of Economics* (November 1979): 541–62.

If the company offers full insurance, then the individual will certainly save the $1,750 by not installing the alarm. The highest premium that the company can charge him solves

$$\ln(100,000 - p) = 11.46113, \tag{18.33}$$

implying that $p = 5{,}048$. The company's profit is then $5{,}048 - (0.25 \times 20{,}000) = \$48$.

On the other hand, the company can induce the individual to install the alarm if it reduces the payment after theft from the full $20,000 down to $3,374 and lowers the premium to $602. (These second-best contractual terms were computed by the authors using numerical methods; we will forgo the complicated computations and just take these terms as given.) Let's check that the individual would indeed want to install the alarm. His expected utility if he accepts the contract and installs the alarm is

$$\begin{aligned} &0.85 \ln(100{,}000 - 1{,}750 - 602) \\ &\quad + 0.15 \ln(100{,}000 - 1{,}750 - 602 - 20{,}000 + 3{,}374) = 11.46113, \end{aligned} \tag{18.34}$$

the same as if he accepts the contract and does not install the alarm:

$$\begin{aligned} &0.75 \ln(100{,}000 - 602) \\ &\quad + 0.25 \ln(100{,}000 - 602 - 20{,}000 + 3{,}374) = 11.46113, \end{aligned} \tag{18.35}$$

also the same as he obtains if he goes without insurance. So he weakly prefers to accept the contract and install the alarm. The insurance company's profit is $602 - (0.15 \times 3{,}374) = \$96$. Thus, partial insurance is more profitable than full insurance when the company cannot observe precaution.

QUERY: What is the most that the insurance company would be willing to spend in order to monitor whether the individual has installed an alarm?

Competitive insurance market

So far in this chapter we have studied insurance using the same principal-agent framework as we used to study the owner-manager relationship. In particular, we have assumed that a monopoly insurance company (principal) makes a take-it-or-leave-it offer to the individual (agent). This is a different perspective than in Chapter 7, where we implicitly assumed that insurance is offered at fair rates—that is, at a premium that just covers the insurer's expected payouts for losses. Fair insurance would arise in a perfectly competitive insurance market.

With competitive insurers, the first best maximizes the insurance customer's expected utility given that the contract can specify his precaution level. The second best maximizes the customer's expected utility under the constraint that his precaution level must be induced by having the contract offer only partial insurance.

Our conclusions about the moral hazard problem remain essentially unchanged when moving from a monopoly insurer to perfect competition. The first best still involves full insurance and a precaution level satisfying Equation 18.27. The second best still involves partial insurance and a moderate level of precaution. The main difference is in the distribution of surplus: insurance companies no longer earn positive profits, since the extra surplus now accrues to the individual.

EXAMPLE 18.3 Competitive Theft Insurance

Return to Example 18.2, but now assume that car theft insurance is sold by perfectly competitive companies rather than by a monopolist.

First best. If companies can costlessly verify whether or not the individual has installed an alarm, then the first-best contract requires him to install the alarm and fully insures him for a premium of 3,000. This is a fair insurance premium because it equals the expected payout for a loss: $3{,}000 = 0.15 \times 20{,}000$. Firms earn zero profit at this fair premium. The individual's expected utility increases to 11.46426 from the 11.46113 of Example 18.2.

Second best. Suppose now that insurance companies cannot observe whether the individual has installed an alarm. The second-best contract is similar to that computed in Example 18.2 except that the \$96 earned by the monopoly insurer is effectively transferred back to the customer in the form of a reduced premium charged by competing insurers. The equilibrium premium is $p = 506$ and the payment for loss is $x = 3{,}374$.

QUERY: Which case—monopoly or perfect competition—best describes the typical insurance market? Which types of insurance (car, health, life, disability) and which countries do you think have more competitive markets?

HIDDEN TYPES

Next we turn to the other leading variant of principal-agent model: the model of hidden types. Whereas in the hidden-action model the agent has private information about a choice he has made, in the hidden-type model he has private information about an innate characteristic he cannot choose. For example, a student's type may be his innate intelligence as opposed to an action such as the effort he expends in studying for an exam.

At first glance, it is not clear why there should be a fundamental economic difference between hidden types and hidden actions that requires us to construct a whole new model (and devote a whole new section to it). The fundamental economic difference is this: In a hidden-type model, the agent has private information *before* signing a contract with the principal; in a hidden-action model, the agent obtains private information afterward.

Having private information before signing the contract changes the game between the principal and the agent. In the hidden-action model, the principal shares symmetric information with the agent at the contracting stage and so can design a contract that extracts all of the agent's surplus. In the hidden-type model, the agent's private information at the time of contracting puts him in a better position. There is no way for the principal to extract all the surplus from all types of agents. A contract that extracts all the surplus from the "high" types (those who benefit more from a given contract) would provide the "low" types with negative surplus, and they would refuse to sign it. The principal will try to extract as much surplus as possible from agents through clever contract design. She will even be willing to shrink the size of the contracting pie, sacrificing some joint surplus in order to obtain a larger share for herself [as in panel (b) of Figure 18.1].

To extract as much surplus as possible from each type while ensuring that low types are not "scared off," the principal will offer a contract in the form of a cleverly designed menu that includes options targeted to each agent type. The menu of options will be more profitable for the principal than a contract with a single option, but the principal will still not be able to extract all the surplus from all agent types. Since the agent's type is hidden, he cannot be forced to select the option targeted at his type but is free to select any of the options, and this ability will ensure that the high types always end up with positive surplus.

To make these ideas more concrete, we will study two applications of the hidden-type model that are important in economics. First we will study the optimal nonlinear pricing problem, and then we will study private information in insurance.

NONLINEAR PRICING

In the first application of the hidden-type model, we consider a monopolist (the principal) who sells to a consumer (the agent) with private information about his own valuation for the good. Rather than allowing the consumer to purchase any amount he wants at a constant price per unit, the monopolist offers the consumer a nonlinear price schedule. The nonlinear price schedule is a menu of different-sized bundles at different prices, from which the consumer makes his selection. In such schedules, the larger bundle generally sells for a higher total price but a lower per-unit price than a smaller bundle.

Our approach builds on the analysis of second-degree price discrimination in Chapter 14. Here we analyze general nonlinear pricing schedules, the most general form of second-degree price discrimination. (In the earlier chapter, we limited our attention to a simpler form of second-degree price discrimination involving two-part tariffs.) The linear, two-part, and general nonlinear pricing schedules are plotted in Figure 18.3. The figure graphs the total tariff—the total cost to the consumer of buying q units—for the three different schedules. Basic and intermediate economics courses focus on the case of a constant per-unit price, which is called a linear pricing schedule. The linear pricing schedule is graphed as a straight line that intersects the origin (because nothing needs to be paid if no units are purchased). The two-part tariff is also a straight line, but its intercept—reflecting the fixed fee—is above the origin. The darkest curve is a general nonlinear pricing schedule.

Examples of nonlinear pricing schedules include a coffee shop selling three different sizes—say, a small (8-ounce) cup for \$1.50, a medium (12-ounce) cup for \$1.80, and a large (16-ounce) cup for \$2.00. Although larger cups cost more in total, they cost less per

FIGURE 18.3

Shapes of Various Pricing Schedules

The graph shows the shape of three different pricing schedules. Thicker curves are more complicated pricing schedules and so represent more sophisticated forms of second-degree price discrimination.

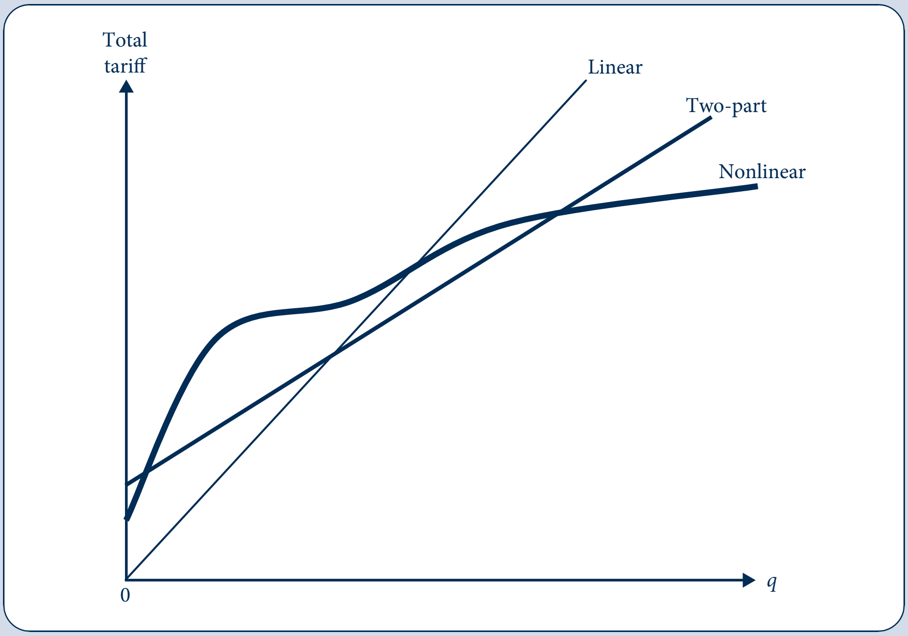

ounce (18.75 cents per ounce for the small, 15 for the medium, and 12.5 for the large). The consumer does not have the choice of buying as much coffee as he wants at a given per-ounce price; instead he must pick one of these three menu options, each specifying a particular bundled quantity. In other examples, the "q" that is bundled in a menu item is the *quality* of a single unit of the product rather than the *quantity* or number of units. For example, an airline ticket involves a single unit (i.e., a single flight) whose quality varies depending on the class of the ticket, which ranges from first class, with fancy drinks and meals and plush seats offering plenty of leg room, to coach class, with peanuts for meals and small seats having little leg room.

Mathematical model

To understand the economic principles involved in nonlinear pricing, consider a formal model in which a single consumer obtains surplus

$$U = \theta v(q) - T \qquad \textbf{(18.36)}$$

from consuming a bundle of q units of a good for which he pays a total tariff of T. The first term in the consumer's utility function, $\theta v(q)$, reflects the consumer's benefit from consumption. Assume $v'(q) > 0$ and $v''(q) < 0$, implying that the consumer prefers more of the good to less but that the marginal benefit of more units is decreasing. The consumer's type is given by θ, which can be high (θ_H) with probability β and low (θ_L) with probability $1 - \beta$. The high type enjoys consuming the good more than the low type: $0 < \theta_L < \theta_H$. The total tariff T paid by the consumer for the bundle is subtracted from his benefit to compute his net surplus.

For simplicity, we are assuming that there is a single consumer in the market. The analysis would likewise apply to markets with many consumers, a proportion β of which are high types and $1 - \beta$ of which are low types. The only complication in extending the model to many consumers is that we would need to assume that consumers cannot divide bundles into smaller packages for resale among themselves. (Of course, such repackaging would be impossible for a single unit of the good involving a bundle of quality; and reselling may be impossible even for quantity bundles if the costs of reselling are prohibitive.)

Suppose the monopolist has a constant marginal and average cost c of producing a unit of the good. Then the monopolist's profit from selling a bundle of q units for a total tariff of T is

$$\Pi = T - cq. \qquad \textbf{(18.37)}$$

First-best nonlinear pricing

In the first-best case, the monopolist can observe the consumer's type θ before offering him a contract. The monopolist chooses the contract terms q and T to maximize her profit subject to Equation 18.37 and subject to a participation constraint that the consumer accepts the contract. Setting the consumer's utility to 0 if he rejects the contract, the participation constraint may be written as

$$\theta v(q) - T \geq 0. \qquad \textbf{(18.38)}$$

The monopolist will choose the highest value of T satisfying the participation constraint: $T = \theta v(q)$. Substituting this value of T into the monopolist's profit function yields

$$\theta v(q) - cq. \qquad \textbf{(18.39)}$$

Taking the first-order condition and rearranging provides a condition for the first-best quantity:

$$\theta v'(q) = c. \tag{18.40}$$

This equation is easily interpreted. In the first best, the marginal social benefit of increased quantity on the left-hand side [the consumer's marginal private benefit, $\theta v'(q)$] equals the marginal social cost on the right-hand side [the monopolist's marginal cost, c].

The first-best quantity offered to the high type (q_H^*) satisfies Equation 18.40 for $\theta = \theta_H$, and that offered to the low type (q_L^*) satisfies the equation for $\theta = \theta_L$. The tariffs are set so as to extract all the type's surplus. The first best for the monopolist is identical to what we termed first-degree price discrimination in Chapter 14.

It is instructive to derive the monopolist's first best in a different way, using methods similar to those used to solve the consumer's utility maximization problem in Chapter 4. The contract (q, T) can be thought of as a bundle of two different "goods" over which the monopolist has preferences. The monopolist regards T as a good (more money is better than less) and q as a bad (higher quantity requires higher production costs). Her indifference curve (actually an isoprofit curve) over (q, T) combinations is a straight line with slope c. To see this, note that the slope of the monopolist's indifference curve is her marginal rate of substitution:

$$MRS = -\frac{\partial\Pi/\partial q}{\partial\Pi/\partial T} = -\frac{(-c)}{1} = c. \tag{18.41}$$

The monopolist's indifference curves are drawn as dashed lines in Figure 18.4. Because q is a bad for the monopolist, her indifference curves are higher as one moves toward the upper left.

FIGURE 18.4

First-Best Nonlinear Pricing

The consumer's indifference curves over the bundle of contractual terms are drawn as solid lines (the thicker one for the high type and thinner for the low type); the monopolist's isoprofits are drawn as dashed lines. Point A is the first-best contract option offered to the high type, and point B is that offered to the low type.

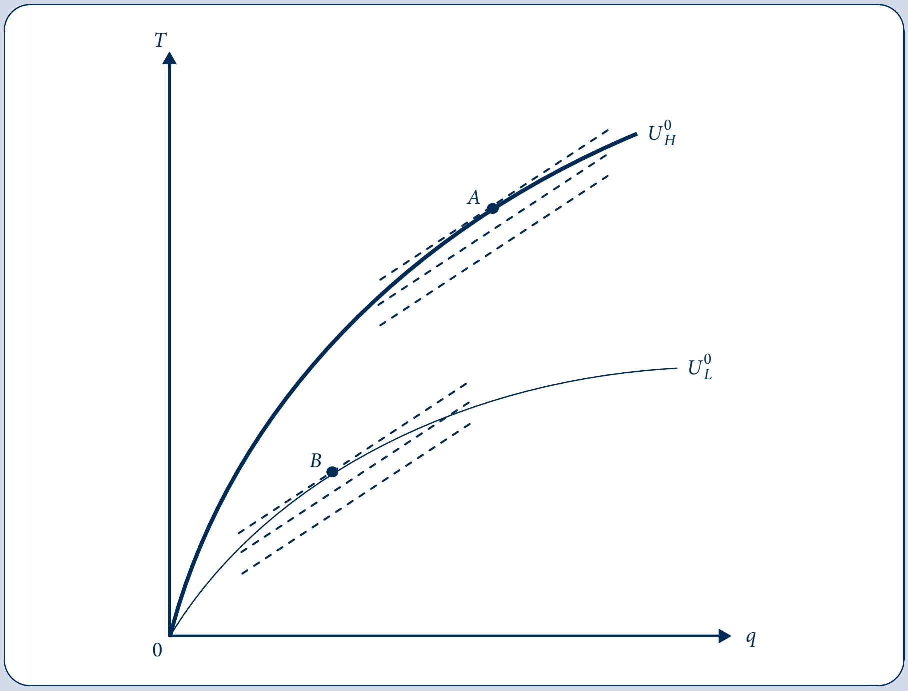

Figure 18.4 also draws indifference curves for the two consumer types: the high type's (labeled U_H^0) and the low type's (labeled U_L^0). Because T is a bad for consumers, higher indifference curves for both types of consumer are reached as one moves toward the lower right. The U_H^0 indifference curve for the high type is special because it intersects the origin, implying that the high type gets the same surplus as if he didn't sign the contract at all. The first-best contract offered by the monopolist to the high type is point A, at which the highest indifference curve for the monopolist still intersects the high type's U_H^0 indifference curve and thus still provides the high type with non-negative surplus. This is a point of tangency between the contracting parties' indifference curves—that is, a point at which the indifference curves have the same slope. The monopolist's indifference curves have slope c everywhere, as we saw in Equation 18.41. The slope of type θ's indifference curve is the marginal rate of substitution:

$$MRS = \frac{\partial U/\partial q}{\partial U/\partial T} = -\frac{\theta v'(q)}{-1} = \theta v'(q). \tag{18.42}$$

Equating the slopes gives the same condition for the first best as we found in Equation 18.40 (marginal social benefit equals marginal social cost of an additional unit). The same arguments imply that point B is the first-best contract offered to the low type, and we can again verify that Equation 18.40 is satisfied there.

To summarize, the first-best contract offered to each type specifies a quantity (q_H^* or q_L^*, respectively) that maximizes social surplus given the type of consumer and a tariff (T_H^* or T_L^*, respectively) that allows the monopolist to extract all of the type's surplus.

Second-best nonlinear pricing

Now suppose that the monopolist does not observe the consumer's type when offering him a contract but knows only the distribution ($\theta = \theta_H$ with probability β and $\theta = \theta_L$ with probability $1 - \beta$). As Figure 18.5 shows, the first-best contract would no longer "work" because the high type obtains more utility (moving from the indifference curve labeled U_H^0 to the one labeled U_H^2) by choosing the bundle targeted to the low type (B) rather than the bundle targeted to him (A). In other words, choosing A is no longer incentive compatible for the high type. To keep the high type from choosing B, the monopolist must reduce the high type's tariff, offering C instead of A.

The substantial reduction in the high type's tariff (indicated by the downward-pointing arrow) puts a big dent in the monopolist's expected profit. The monopolist can do better than offering the menu of contracts (B, C): she can distort the low type's bundle in order to make it less attractive to the high type. Then the high type's tariff need not be reduced as much to keep him from choosing the wrong bundle. Figure 18.6 shows how this new contract would work. The monopolist reduces the quantity in the low type's bundle (while reducing the tariff so that the low type stays on his U_L^0 indifference curve and thus continues to accept the contract), offering bundle D rather than B. The high type obtains less utility from D than B, as D reaches only his U_H^1 indifference curve and is short of his U_H^2 indifference curve. To keep the high type from choosing D, the monopolist need only lower the high type's tariff by the amount given by the vertical distance between A and E rather than all the way down to C.

Relative to (B, C), the second-best menu of contracts (D, E) trades off a distortion in the low type's quantity (moving from the first-best quantity in B to the lower quantity in D and destroying some social surplus in the process) against an increase in the tariff that can be extracted from the high type in moving from C to E. An attentive student might wonder why the monopolist would want to make this trade-off. After all, the monopolist must reduce the low type's tariff in moving from B to D or else the low

FIGURE 18.5

First Best Not Incentive Compatible

The first-best contract, involving points A and B, is not incentive compatible if the consumer has private information about his type. The high type can reach a higher indifference curve by choosing the bundle (B) that is targeted at the low type. To keep him from choosing B, the monopolist must reduce the high type's tariff by replacing bundle A with C.

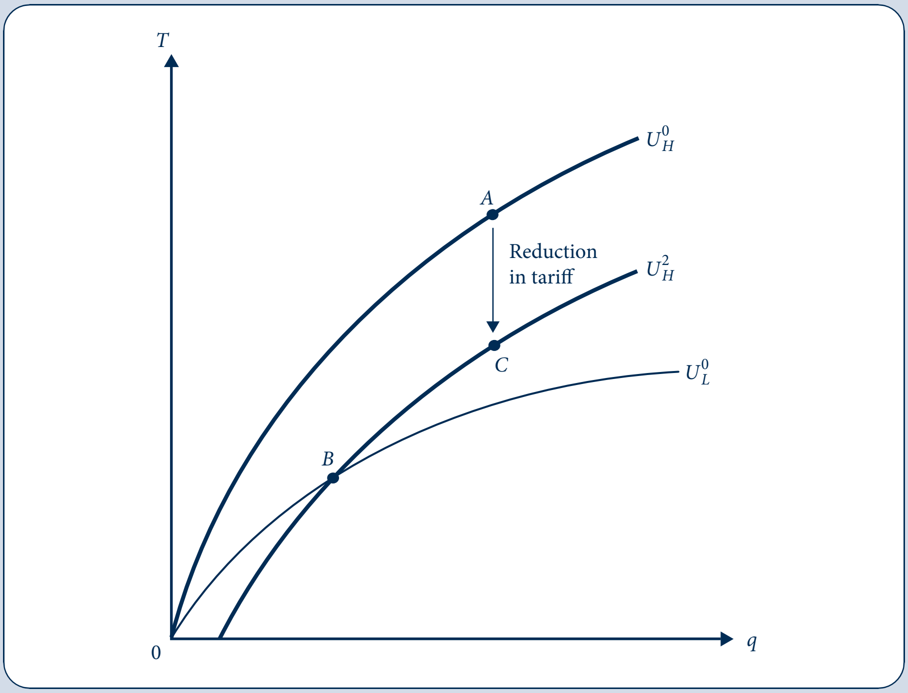

type would refuse to accept the contract. How can we be sure that this reduction in the low type's tariff doesn't more than offset any increase in the high type's tariff? The reason is that a reduction in quantity harms the high type more than it does the low type. As Equation 18.42 shows, the consumer's marginal rate of substitution between contractual terms (quantity and tariff) depends on his type θ and is higher for the high type. Since the high type values quantity more than does the low type, the high type would pay more to avoid the decrease in quantity in moving from B to D than would the low type.

Further insight can be gained from an algebraic characterization of the second best. The second-best contract is a menu that targets bundle (q_H, T_H) at the high type and (q_L, T_L) at the low type. The contract maximizes the monopolist's expected profit,

$$\beta(T_H - cq_H) + (1 - \beta)(T_L - cq_L), \tag{18.43}$$

subject to four constraints:

$$\theta_L v(q_L) - T_L \geq 0, \tag{18.44}$$

$$\theta_H v(q_H) - T_H \geq 0, \tag{18.45}$$

$$\theta_L v(q_L) - T_L \geq \theta_L v(q_H) - T_H, \tag{18.46}$$

$$\theta_H v(q_H) - T_H \geq \theta_H v(q_L) - T_L. \tag{18.47}$$

FIGURE 18.6

Second-Best Nonlinear Pricing

The second-best contract is indicated by the circled points D and E. Relative to the incentive-compatible contract found in Figure 18.5 (points B and C), the second-best contract distorts the low type's quantity (indicated by the move from B to D) in order to make the low type's bundle less attractive to the high type. This allows the principal to charge tariff to the high type (indicated by the move from C to E).

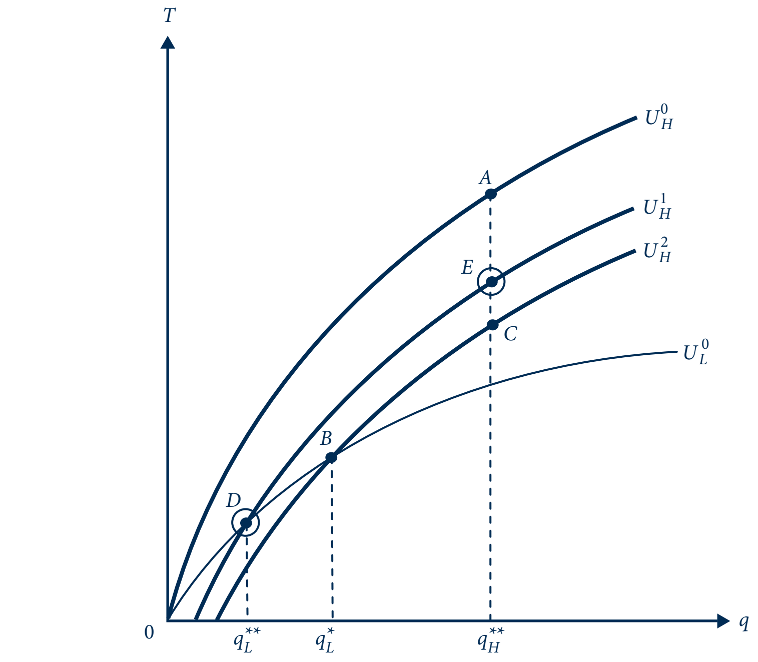

The first two are participation constraints for the low and high type of consumer, ensuring that they accept the contract rather than forgoing the monopolist's good. The last two are incentive compatibility constraints, ensuring that each type chooses the bundle targeted to him rather than the other type's bundle.

As suggested by the graphical analysis in Figure 18.6, only two of these constraints play a role in the solution. The most important constraint was to keep the high type from choosing the low type's bundle; this is Equation 18.47 (incentive compatibility constraint for the high type). The other relevant constraint was to keep the low type on his U_L^0 indifference curve to prevent him from rejecting the contract; this is Equation 18.44 (participation constraint for the low type). Hence, Equations 18.44 and 18.47 hold with equality in the second best.

The other two constraints can be ignored, as can be seen in Figure 18.6. The high type's second-best bundle E puts him on a higher indifference curve (U_H^1) than if he rejects the contract (U_H^0), so the high type's participation constraint (Equation 18.45) can be safely ignored. The low type would be on a lower indifference curve if he chose the high type's bundle (E) rather than his own (D), so the low type's incentive compatibility constraint (Equation 18.46) can also be safely ignored.

Treating Equations 18.44 and 18.47 as equalities and using them to solve for T_L and T_H yields

$$T_L = \theta_L v(q_L) \qquad \textbf{(18.48)}$$

and

$$\begin{aligned} T_H &= \theta_H[v(q_H) - v(q_L)] + T_L \\ &= \theta_H[v(q_H) - v(q_L)] + \theta_L v(q_L). \end{aligned} \tag{18.49}$$

By substituting these expressions for T_L and T_H into the monopolist's objective function (Equation 18.39), we convert a complicated maximization problem with four inequality constraints into the simpler unconstrained problem of choosing q_L and q_H to maximize

$$\beta\{\theta_H[v(q_H) - v(q_L)] + \theta_L v(q_L) - cq_H\} + (1-\beta)[\theta_L v(q_L) - cq_L]. \tag{18.50}$$

The low type's quantity satisfies the first-order condition with respect to q_L, which (upon considerable rearranging) yields

$$\theta_L v'(q_L^{**}) = c + \frac{\beta(\theta_H - \theta_L)v'(q_L^{**})}{1-\beta}. \tag{18.51}$$

The last term is clearly positive, and thus the equation implies that $\theta_L v'(q_L^{**}) > c$, whereas $\theta_L v'(q_L^*) = c$ in the first best. Since $v(q)$ is concave, we see that the second-best quantity is lower than the first best, verifying the insight from our graphical analysis that the low type's quantity is distorted downward in the second best to extract surplus from the high type.

The high type's quantity satisfies the first-order condition from the maximization of Equation 18.43 with respect to q_H; upon rearranging, this yields

$$\theta_H v'(q_H^{**}) = c. \tag{18.52}$$

This condition is identical to the first best, implying that there is no distortion of the high type's quantity in the second best. There is no reason to distort the high type's quantity because there is no higher type from whom to extract surplus. The result that the highest type is offered an efficient contract is often referred to as "no distortion at the top."

Returning to the low type's quantity, how much the monopolist distorts this quantity downward depends on the probabilities of the two consumer types or—equivalently, in a model with many consumers—on the relative proportions of the two types. If there are many low types (β is low), then the monopolist would not be willing to distort the low type's quantity very much, because the loss from this distortion would be substantial and there would be few high types from whom additional surplus could be extracted. The more high types (the higher is β), the more the monopolist is willing to distort the low type's quantity downward. Indeed, if there are enough high types, the monopolist may decide not to serve the low types at all and just offer one bundle that would be purchased by the high types. This would allow the monopolist to squeeze all the surplus from the high types because they would have no other option.

EXAMPLE 18.4 Monopoly Coffee Shop

The college has a single coffee shop whose marginal cost is 5 cents per ounce of coffee. The representative customer is equally likely to be a coffee hound (high type with $\theta_H = 20$) or a regular Joe (low type with $\theta_L = 15$). Assume $v(q) = 2\sqrt{q}$.

First best. Substituting the functional form $v(q) = 2\sqrt{q}$ into the condition for first-best quantities $[\theta v'(q) = c]$ and rearranging, we have $q = (\theta/c)^2$. Therefore, $q_L^* = 9$ and $q_H^* = 16$. The tariff extracts all of each type's surplus $[T = \theta v(q)]$, here implying that $T_L^* = 90$ and $T_H^* = 160$. The shop's expected profit is

$$\frac{1}{2}(T_H^* - cq_H^*) + \frac{1}{2}(T_L^* - cq_L^*) = 62.5 \tag{18.53}$$

cents per customer. The first best can be implemented by having the owner sell a 9-ounce cup for 90 cents to the low type and a 16-ounce cup for \$1.60 to the high type. (Somehow the barista can discern the customer's type just by looking at him as he walks in the door.)

Incentive compatibility when types are hidden. The first best is not incentive compatible if the barista cannot observe the customer's type. The high type obtains no surplus from the 16-ounce cup sold at \$1.60. If he instead paid 90 cents for the 9-ounce cup, he would obtain a surplus of $\theta_H v(9) - 90 = 30$ cents. Keeping the same cup sizes as in the first best, the price for the large cup would have to be reduced by 30 cents (to \$1.30) in order to keep the high type from buying the small cup. The shop's expected profit from this incentive compatible menu is

$$\frac{1}{2}(130 - 5 \cdot 16) + \frac{1}{2}(90 - 5 \cdot 9) = 47.5. \tag{18.54}$$

Second best. The shop can do even better by reducing the size of the small cup to make it less attractive to high demanders. The size of the small cup in the second best satisfies Equation 18.51, which, for the functional forms in this example, implies that

$$\theta_L q_L^{-1/2} = c + (\theta_H - \theta_L) q_L^{-1/2} \tag{18.55}$$

or, rearranging,

$$q_L^{**} = \left(\frac{2\theta_L - \theta_H}{c}\right)^2 = \left(\frac{2 \cdot 15 - 20}{5}\right)^2 = 4. \tag{18.56}$$

The highest price that can be charged without losing the low-type customers is

$$T_L^{**} = \theta_L v(q_L^{**}) = (15)(2\sqrt{4}) = 60. \tag{18.57}$$

The large cup is the same size as in the first best: 16 ounces. It can be sold for no more than \$1.40 or else the coffee hound would buy the 4-ounce cup instead. Although the total tariff for the large cup is higher at \$1.40 than for the small cup at 60 cents, the unit price is lower (8.75 cents versus 15 cents per ounce). Hence the large cup sells at a quantity discount.

The shop's expected profit is

$$\frac{1}{2}(140 - 5 \cdot 16) + \frac{1}{2}(60 - 5 \cdot 4) = 50 \tag{18.58}$$

cents per consumer. Reducing the size of the small cup from 9 to 4 ounces allows the shop to recapture some of the profit lost when the customer's type cannot be observed.

QUERY: In the first-best menu, the price per ounce is the same (10 cents) for both the low and high type's cup. Can you explain why it is still appropriate to consider this a nonlinear pricing scheme?

ADVERSE SELECTION IN INSURANCE

For the second application of the hidden-type model, we will return to the insurance market in which an individual with state-independent preferences and initial income W_0 faces the prospect of loss l. Assume the individual can be one of two types: a high-risk type with probability of loss π_H or a low-risk type with probability π_L, where $\pi_H > \pi_L$. We will first assume the insurance company is a monopolist; later we will study the case of competitive insurers. The presence of hidden risk types in an insurance market is said to lead to *adverse selection*. Insurance tends to attract more risky than safe consumers (the "selection" in adverse selection) because it is more valuable to risky types, yet risky types are more expensive to serve (the "adverse" in adverse selection).

DEFINITION

Adverse selection. The problem facing insurers that risky types are both more likely to accept an insurance policy and more expensive to serve.

As we will see, if the insurance company is clever then it can mitigate the adverse selection problem by offering a menu of contracts. The policy targeted to the safe type offers only partial insurance so that it is less attractive to the high-risk type.

First best

In the first best, the insurer can observe the individual's risk type and offer a different policy to each. Our previous analysis of insurance makes it clear that the first best involves full insurance for each type, so the insurance payment x in case of a loss equals the full amount of the loss l. Different premiums are charged to each type and are set to extract all of the surplus that each type obtains from the insurance.

The solution is shown in Figure 18.7 (the construction of this figure is discussed further in Chapter 7). Without insurance, each type finds himself at point E. Point A (resp., B) is the first-best policy offered to the high-risk (resp., low-risk) type. Points A and B lie on the certainty line because both are fully insured. Since the premiums extract each type's surplus from insurance, both types are on their indifference curves through the no-insurance point E. The high type's premium is higher, so A is further down the certainty line toward the origin than is B.[7]

Second best

If the monopoly insurer cannot observe the agent's type, then the first-best contracts will not be incentive compatible: the high-risk type would claim to be low risk and take full insurance coverage at the lower premium. As in the nonlinear pricing problem, the second best will involve a menu of contracts. Other principles from the nonlinear pricing problem also carry over here. The high type continues to receive the first-best quantity (here, full insurance)—there is no distortion at the top. The low type's quantity is distorted downward from the first best, so he receives only partial insurance. Again we see that, with hidden types, the principal is willing to sacrifice some social surplus in order to extract some of the surplus the agent would otherwise derive from his private information.

[7]Mathematically, A appears further down the certainty line than B in Figure 18.7 because the high type's indifference curve through E is flatter than the low type's. To see this, note that expected utility equals $(1-\pi)U(W_1)+\pi U(W_2)$ and so the MRS is given by

$$-\frac{dW_1}{dW_2}=\frac{(1-\pi)U'(W_1)}{\pi U'(W_2)}.$$

At a given (W_1, W_2) combination on the graph, the marginal rates of substitution differ only because the underlying probabilities of loss differ. Since

$$\frac{1-\pi_H}{\pi_H}<\frac{1-\pi_L}{\pi_L},$$

it follows that the high-risk type's indifference curve will be flatter. This proof follows the analysis presented in M. Rothschild and J. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics* (November 1976): 629–50.

FIGURE 18.7

First Best for a Monopoly Insurer

In the first best, the monopoly insurer offers policy A to the high-risk type and B to the low-risk type. Both types are fully insured. The premiums are sufficiently high to keep each type on his indifference curve through the no-insurance point (E).

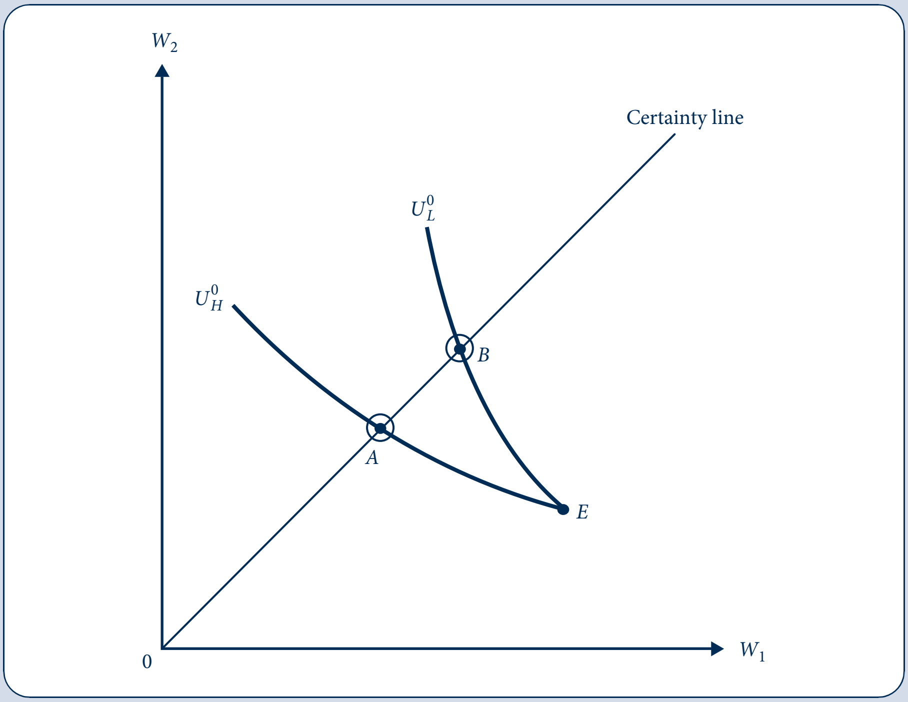

Figure 18.8 depicts the second best. If the insurer tried to offer a menu containing the first-best contracts A and B, then the high-risk type would choose B rather than A. To maintain incentive compatibility, the insurer distorts the low type's policy from B along its indifference curve U_L^0 down to D. The low type is only partially insured, and this allows the insurer to extract more surplus from the high type. The high type continues to be fully insured, but the increase in his premium shifts his policy down the certainty line to C.

EXAMPLE 18.5 Insuring the Little Red Corvette

The analysis of automobile insurance in Example 18.2 (which is based on Example 7.2) can be recast as an adverse selection problem. Suppose that the probability of theft depends not on the act of installing an antitheft device but rather on the color of the car. Because thieves prefer red to gray cars, the probability of theft is higher for red cars ($\pi_H = 0.25$) than for gray cars ($\pi_L = 0.15$).

First best. The monopoly insurer can observe the car color and offer different policies for different colors. Both colors are fully insured for the \$20,000 loss of the car. The premium is the maximum amount that each type would be willing to pay in lieu of going without insurance; as computed in Example 7.2, this amount is \$5,426 for the high type (red cars). Similar calculations show that a gray-car owner's expected utility if he is not insured is 11.4795, and the maximum premium he would be willing to pay for full insurance is \$3,287. Although the insurer pays more claims for red cars, the higher associated premium more than compensates, and thus the expected profit from a policy sold for a red car is $5{,}426 - 0.25 \cdot 20{,}000 = \$426$ versus $3{,}287 - 0.15 \cdot 20{,}000 = \$287$ for a gray car.

FIGURE 18.8

Second Best for a Monopoly Insurer

Second-best insurance policies are represented by the circled points: C for the high-risk type and D for the low-risk type.

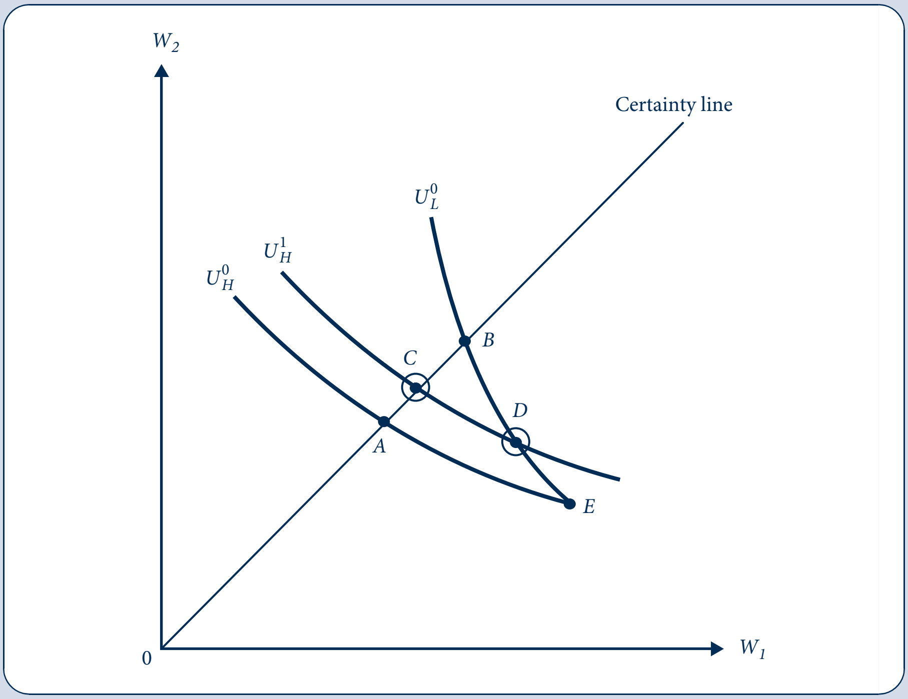

Second best. Suppose the insurer does not observe the color of the customer's car and knows only that 10 percent of all cars are red and the rest are gray. The second-best menu of insurance policies—consisting of a premium/insurance coverage bundle (p_H, x_H) targeted for high-risk, red cars and (p_L, x_L) for low-risk, gray cars—is indicated by the circled points in Figure 18.8. Red cars are fully insured: $x_H = 20{,}000$. To solve for the rest of the contractual parameters, observe that x_H, p_H, and p_L can be found as the solution to the maximization of expected insurer profit

$$0.1(p_H - 0.25 \cdot 20{,}000) + 0.9(p_L - 0.15x_L) \tag{18.59}$$

subject to a participation constraint for the low type,

$$0.85\ \ln(100{,}000 - p_L) + 0.15\ \ln(100{,}000 - p_L - 20{,}000 + x_L) \geq 11.4795, \tag{18.60}$$

and to an incentive compatibility constraint for the high type,

$$\begin{aligned} \ln(100{,}000 - p_H) \geq\ & 0.75\ \ln(100{,}000 - p_L) \\ &+\ 0.25\ \ln(100{,}000 - p_L - 20{,}000 + x_L). \end{aligned} \tag{18.61}$$

Participation and incentive compatibility constraints for the other types can be ignored, just as in the nonlinear pricing problem. This maximization problem is too difficult to solve by hand, but it can be solved numerically using popular spreadsheet programs or other mathematical software. The second-best values that result are $x_H^{**} = 20{,}000$, $p_H^{**} = 4{,}154$, $x_L^{**} = 11{,}556$, and $p_L^{**} = 1{,}971$.

QUERY: Look at the spreadsheet associated with this example on the website for this textbook. Play around with different probabilities of the two car colors. What happens when red cars become sufficiently common? (Even if you cannot access the spreadsheet, you should be able to guess the answer.)

Competitive insurance market

Assume now that insurance is provided not by a monopoly but rather by a perfectly competitive market, resulting in fair insurance. Figure 18.9 depicts the equilibrium in which insurers can observe each individual's risk type. Lines EF and EG are drawn with slopes $-(1-\pi_L)/\pi_L$ and $-(1-\pi_H)/\pi_H$, respectively, and show the market opportunities for each person to trade W_1 for W_2 by purchasing fair insurance.[8] The low-risk type is sold policy F, and the high-risk type is sold policy G. Each type receives full insurance at a fair premium.

However, the outcome in Figure 18.9 is unstable if insurers cannot observe risk types. The high type would claim to be low risk and take contract F. But then insurers that offered F would earn negative expected profit: at F, insurers break even serving only the low-risk types, so adding individuals with a higher probability of loss would push the company below the break-even point.

FIGURE 18.9

Competitive Insurance Equilibrium with Perfect Information

With perfect information, the competitive insurance market results in full insurance at fair premiums for each type. The high type is offered policy G; the low type, policy F.

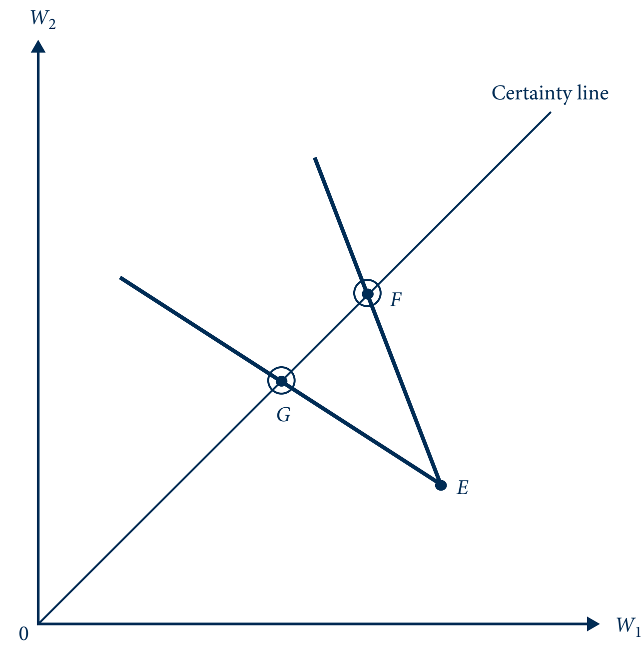

[8]To derive these slopes, called *odds ratios*, note that fair insurance requires the premium to satisfy $p = \pi x$. Substituting into W_1 and W_2 yields

$$W_1 = W_0 - p = W_0 - \pi x$$
$$W_2 = W_0 - p - l + x = W_0 - l + (1-\pi)x.$$

Hence a \$1 increase in the insurance payment (x) reduces W_1 by π and increases W_2 by $1-\pi$.

The competitive equilibrium with unobservable types is shown in Figure 18.10. The equilibrium is similar to the second best for a monopoly insurer. A set of policies is offered that separates the types. The high-risk type is fully insured at point *G*, the same policy as he was offered in the first best. The low-risk type is offered policy *J*, which features partial insurance. The low type would be willing to pay more for fuller insurance, preferring a policy such as *K*. Because *K* is below line *EF*, an insurer would earn positive profit from selling such a policy to low-risk types only. The problem is that *K* would also attract high-risk types, leading to insurer losses. Hence insurance is rationed to the low-risk type.

With hidden types, the competitive equilibrium must involve a set of separating contracts; it cannot involve a single policy that pools both types. This can be shown with the aid of Figure 18.11. To be accepted by both types and allow the insurer to at least break even, the pooling contract would have to be a point (such as *M*) within triangle *EFG*. But *M* cannot be a final equilibrium because at *M* there exist further trading opportunities. To see this, note that—as indicated in the figure and discussed earlier in the chapter—the indifference curve for the high type (U_H) is flatter than that for the low type (U_L). Consequently, there are insurance policies such as *N* that are unattractive to high-risk types, attractive to low-risk types, and profitable to insurers (because such policies lie below *EF*).

FIGURE 18.10

Competitive Insurance Equilibrium with Hidden Types

With hidden types, the high-risk type continues to be offered first-best policy *G* but the low-risk type is rationed, receiving only partial insurance at *J* in order to keep the high-risk type from pooling.

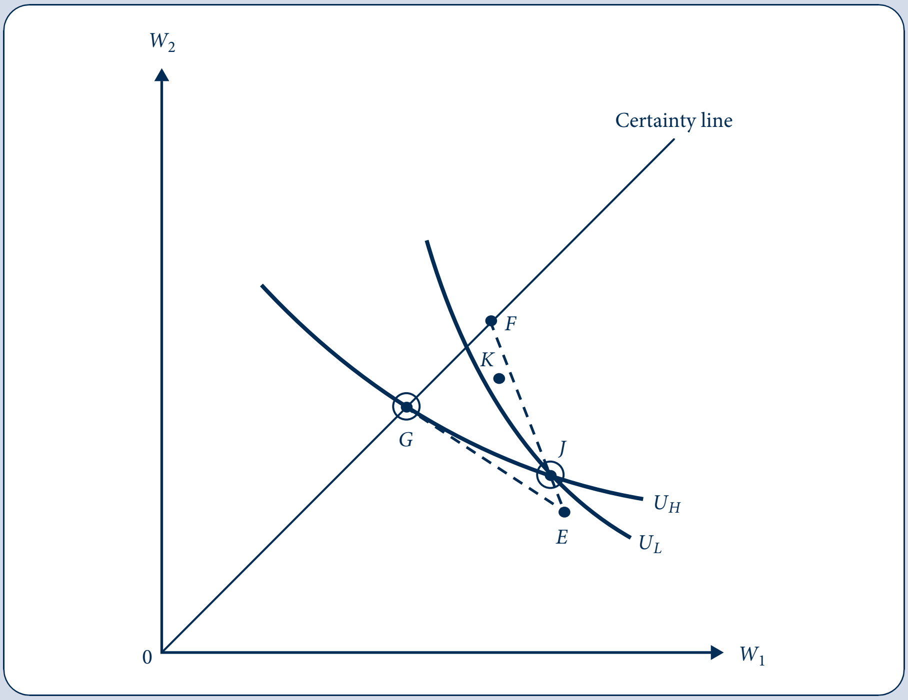

FIGURE 18.11

Impossibility of a Competitive Pooling Equilibrium

Pooling contract M cannot be an equilibrium because there exist insurance policies such as N that are profitable to insurers and are attractive to low-risk types but not to high-risk types.

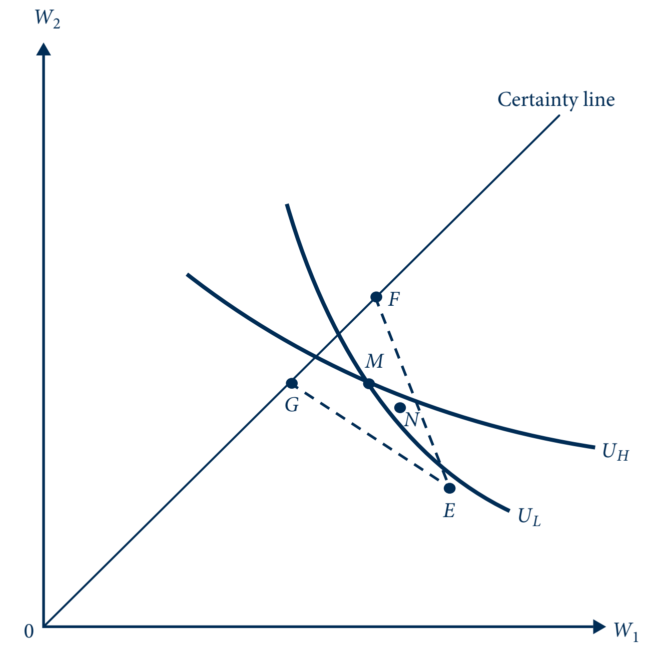

Assuming that no barriers prevent insurers from offering new contracts, policies such as N will be offered and will "skim the cream" of low-risk individuals from any pooling equilibrium. Insurers that continue to offer M are left with the "adversely selected" individuals, whose risk is so high that insurers cannot expect to earn any profit by serving them.

EXAMPLE 18.6 Competitive Insurance for the Little Red Corvette

Recall the automobile insurance analysis in Example 18.5, but now assume that insurance is provided by a competitive market rather than a monopolist. Under full information, the competitive equilibrium involves full insurance for both types at a fair premium of (0.25)(20,000) = \$5,000 for high-risk, red cars and (0.15)(20,000) = \$3,000 for low-risk, gray cars.

If insurers cannot observe car colors, then in equilibrium the coverage for the two types will still be separated into two policies. The policy targeted for red cars is the same as under full information. The policy targeted for gray cars involves a fair premium

$$p_L = 0.15x_L \tag{18.62}$$

and an insurance level that does not give red-car owners an incentive to deviate by pooling on the gray-car policy:

$$0.75\ \ln(100{,}000 - p_L) + 0.25\ \ln(100{,}000 - p_L - 20{,}000\ +\ x_L) = \ln(95{,}000). \tag{18.63}$$

Equations 18.62 and 18.63 can be solved numerically, yielding $p_L = 453$ and $x_L = 3{,}020$.

QUERY: How much more would gray-car owners be willing to pay for full insurance? Would an insurer profit from selling full insurance at this higher premium if it sold only to owners of gray cars? Why then do the companies ration insurance to gray cars by insuring them partially?

MARKET SIGNALING

In all the models studied so far, the uninformed principal moved first—making a contract offer to the agent, who had private information. If the information structure is reversed and the informed player moves first, then the analysis becomes much more complicated, putting us in the world of signaling games studied in Chapter 8. When the signaler is a principal who is offering a contract to an agent, the signaling games become complicated because the strategy space of contractual terms is virtually limitless. Compare the simpler strategy space of Spence's education signaling game in Chapter 8, where the worker chose one of just two actions: to obtain an education or not. We do not have space to delve too deeply into complex signaling games here nor to repeat Chapter 8's discussion of simpler signaling games. We will be content to gain some insights from a few simple applications.

Signaling in competitive insurance markets

In a competitive insurance market with adverse selection (i.e., hidden risk types), we saw that the low-risk type receives only partial insurance in equilibrium. He would benefit from report of his type, perhaps hiring an independent auditor to certify that type so the reporting would be credible. The low-risk type would be willing to pay the difference between his equilibrium and his first-best surplus in order to issue such a credible signal.

It is important that there be some trustworthy auditor or other way to verify the authenticity of such reports, because a high-risk individual would now have an even greater incentive to make false reports. The high-risk type may even be willing to pay a large bribe to the auditor for a false report.

EXAMPLE 18.7 Certifying Car Color

Return to the competitive market for automobile insurance from Example 18.6. Let R be the most that the owner of a gray car would be willing to pay to have his car color (and thus his type) certified and reported to the market. He would then be fully insured at a fair premium of \$3,000, earning surplus $\ln(100{,}000 - 3{,}000 - R)$. In the absence of such a certified report, his expected surplus is

$$\begin{aligned} &0.85 \ln(100{,}000 - 453) + 0.15 \ln(100{,}000 - 453 - 20{,}000 + 3{,}020) \\ &\quad = 11.4803. \end{aligned} \tag{18.64}$$

Solving for R in the equation

$$\ln(100{,}000 - 453 - R) = 11.4803 \tag{18.65}$$

yields $R = 207$. Thus the low-risk type would be willing to pay up to \$207 to have a credible report of his type issued to the market.

The owner of the red car would pay a bribe as high as \$2,000—the difference between his fair premium with full information (\$5,000) and the fair premium charged to an individual known to be of low risk (\$3,000). Therefore, the authenticity of the report is a matter of great importance.

QUERY: How would the equilibrium change if reports are not entirely credible (i.e., if there is some chance the high-risk individual can successfully send a false report about his type)? What incentives would an auditor have to maintain his or her reputation for making honest reports?

Market for lemons

Markets for used goods raise an interesting possibility for signaling. Cars are a leading example: having driven the car over a long period of time, the seller has much better information about its reliability and performance than a buyer, who can take only a short test drive. Yet even the mere act of offering the car for sale can be taken as a signal of car quality by the market. The signal is not positive: the quality of the good must be below the threshold that would have induced the seller to keep it. As George Akerlof showed in the article for which he won the Nobel Prize in economics, the market may unravel in equilibrium so that only the lowest-quality goods, the "lemons," are sold.[9]

To gain more insight into this result, consider the used-car market. Suppose there is a continuum of qualities from low-quality lemons to high-quality gems and that only the owner of a car knows its type. Because buyers cannot differentiate between lemons and gems, all used cars will sell for the same price, which is a function of the average car quality. A car's owner will choose to keep it if the car is at the upper end of the quality spectrum (since a good car is worth more than the prevailing market price) but will sell the car if it is at the low end (since these are worth less than the market price). This reduction in average quality of cars offered for sale will reduce market price, leading would-be sellers of the highest-quality remaining cars to withdraw from the market. The market continues to unravel until only the worst-quality lemons are offered for sale.

The lemons problem leads the market for used cars to be much less efficient than it would be under the standard competitive model in which quality is known. (Indeed, in the standard model the issue of quality does not arise, because all goods are typically assumed to be of the same quality.) Whole segments of the market disappear—along with the gains from trade in these segments—because higher-quality items are no longer traded. In the extreme, the market can simply break down with nothing (or perhaps just a few of the worst items) being sold. The lemons problem can be mitigated by trustworthy used-car dealers, by development of car-buying expertise by the general public, by sellers providing proof that their cars are trouble-free, and by sellers offering money-back guarantees. But anyone who has ever shopped for a used car knows that the problem of potential lemons is a real one.

EXAMPLE 18.8 Used-Car Market

Suppose the quality q of used cars is uniformly distributed between 0 and 20,000. Sellers value their cars at q. Buyers (equal in number to the sellers) place a higher value on cars, $q + b$, so there are gains to be made from trade in the used-car market. Under full information about quality, all used cars would be sold. But this does not occur when sellers have private information about quality and buyers know only the distribution. Let p be the market price. Sellers offer their cars for sale if and only if $q \leq p$. The quality of a car offered for sale is thus uniformly distributed between 0 and p, implying that expected quality is

$$\int_0^p q\left(\frac{1}{p}\right)dq = \frac{p}{2} \qquad \textbf{(18.66)}$$

[9] G. A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* (August 1970): 488–500.

(see Chapter 2 for background on the uniform distribution). Hence, a buyer's expected net surplus is

$$\frac{p}{2} + b - p = b - \frac{p}{2}. \tag{18.67}$$

There may be multiple equilibria, but the one with the most sales involves the highest value of p for which Equation 18.67 is non-negative: $b - p/2 = 0$, implying that $p^* = 2b$. Only a fraction $2b/20{,}000$ of the cars are sold. As b decreases, the market for used cars dries up.

QUERY: What would the equilibrium look like in the full-information case?

AUCTIONS

The monopolist has difficulty extracting surplus from the agent in the nonlinear pricing problem because high-demand consumers could guarantee themselves a certain surplus by choosing the low demanders' bundle. A seller can often do better if several consumers compete against each other for her scarce supplies in an auction. Competition among consumers in an auction can help the seller solve the hidden-type problem, because high-value consumers are then pushed to bid high so they don't lose the good to another bidder. In the setting of an auction, the principal's "offer" is no longer a simple contract or menu of contracts as in the nonlinear pricing problem; instead, her offer is the format of the auction itself. Different formats might lead to substantially different outcomes and more or less revenue for the seller, so there is good reason for sellers to think carefully about how to design the auction. There is also good reason for buyers to think carefully about what bidding strategies to use.

Auctions have received a great deal of attention in the economics literature ever since William Vickery's seminal work, for which he won the Nobel Prize in economics.[10] Auctions continue to grow in significance as a market mechanism and are used for selling such goods as airwave spectrum, Treasury bills, foreclosed houses, and collectibles on the Internet auction site eBay.

There are a host of different auction formats. Auctions can involve sealed bids or open outcries. Sealed-bid auctions can be first price (the highest bidder wins the object and must pay the amount bid) or second price (the highest bidder still wins but need only pay the next-highest bid). Open-outcry auctions can be either ascending, as in the so-called English auction where buyers yell out successively higher bids until no one is willing to top the last, or descending, as in the so-called Dutch auction where the auctioneer starts with a high price and progressively lowers it until one of the participants stops the auction by accepting the price at that point. The seller can decide whether or not to set a "reserve clause," which requires bids to be over a certain threshold else the object will not be sold. Even more exotic auction formats are possible. In an "all-pay" auction, for example, bidders pay their bids even if they lose.

A powerful and somewhat surprising result due to Vickery is that, in simple settings (risk-neutral bidders who each know their valuation for the good perfectly, no collusion, etc.), many of the different auction formats listed here (and more besides) provide the monopolist with the same expected revenue in equilibrium. To see why this result is

[10] W. Vickery, "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance* (March 1961): 8–37.

surprising, we will analyze two auction formats in turn—a first-price and a second-price sealed-bid auction—supposing that a single object is to be sold.

In the first-price sealed-bid auction, all bidders simultaneously submit secret bids. The auctioneer unseals the bids and awards the object to the highest bidder, who pays his or her bid. In equilibrium, it is a weakly dominated strategy to submit a bid b greater than or equal to the buyer's valuation v.

DEFINITION

Weakly dominated strategy. A strategy is *weakly dominated* if there is another strategy that does at least as well against all rivals' strategies and strictly better against at least one.

A buyer receives no surplus if he bids $b = v$ no matter what his rivals bid: if the buyer loses, he gets no surplus; if he wins, he must pay his entire surplus back to the seller and again gets no surplus. By bidding less than his valuation, there is a chance that others' valuations (and consequent bids) are low enough that the bidder wins the object and derives a positive surplus. Bidding more than his valuation is even worse than just bidding his valuation. There is good reason to think that players avoid weakly dominated strategies, meaning here that bids will be below buyers' valuations.

In a second-price sealed-bid auction, the highest bidder pays the next-highest bid rather than his own. This auction format has a special property in equilibrium. All bidding strategies are weakly dominated by the strategy of bidding exactly one's valuation. Vickery's analysis of second-price auctions and of the property that they induce bidders to reveal their valuations has led them to be called *Vickery auctions*.

We will prove that, in this kind of auction, bidding something other than one's true valuation is weakly dominated by bidding one's valuation. Let v be a buyer's valuation and b his bid. If the two variables are not equal, then there are two cases to consider: either $b < v$ or $b > v$. Consider the first case ($b < v$). Let $\tilde{b}$ be the highest rival bid. If $\tilde{b} > v$, then the buyer loses whether his bid is b or v, so there is a tie between the strategies. If $\tilde{b} < b$, then the buyer wins the object whether his bid is b or v and his payment is the same (the second-highest bid, $\tilde{b}$) in either case, so again we have a tie. We no longer have a tie if $\tilde{b}$ lies between b and v. If the buyer bids b, then he loses the object and obtains no surplus. If he bids v, then he wins the object and obtains a net surplus of $v - \tilde{b} > 0$, so bidding v is strictly better than bidding $b < v$ in this case. Similar logic shows that bidding v weakly dominates bidding $b > v$.

The reason that bidding one's valuation is weakly dominant is that the winner's bid does not affect the amount he has to pay, for that depends on someone else's (the second-highest bidder's) bid. But bidding one's valuation ensures the buyer wins the object when he should.

With an understanding of equilibrium bidding in second-price auctions, we can compare first- and second-price sealed-bid auctions. Each format has plusses and minuses with regard to the revenue the seller earns. On the one hand, bidders shade their bids below their valuations in the first-price auction but not in the second-price auction, a "plus" for second-price auctions. On the other hand, the winning bidder pays the highest bid in the first-price auction but only the second-highest bid in the second-price auction, a "minus" for second-price auctions. The surprising result proved by Vickery is that these plusses and minuses balance perfectly, so that both auction types provide the seller with the same expected revenue. Rather than working through a general proof of this *revenue equivalence* result, we will show in Example 18.9 that it holds in a particular case.

EXAMPLE 18.9 Art Auction

Suppose two buyers (1 and 2) bid for a painting in a first-price sealed-bid auction. Buyer i's valuation, v_i, is a random variable that is uniformly distributed between 0 and 1 and is independent of the other buyer's valuation. Buyers' valuations are private information. We will look for a symmetric equilibrium in which buyers bid a constant fraction of their valuations, $b_i = kv_i$. The remaining step is to solve for the equilibrium value of k.

Symmetric equilibrium. Given that buyer 1 knows his own type v_1 and knows buyer 2's equilibrium strategy $b_2 = kv_2$, buyer 1 best responds by choosing the bid b_1 maximizing his expected surplus

$$\begin{aligned}&\Pr(1 \text{ wins auction})(v_1 - b_1) + \Pr(1 \text{ loses auction})(0)\\ &\quad = \Pr(b_1 > b_2)(v_1 - b_1)\\ &\quad = \Pr(b_1 > kv_2)(v_1 - b_1)\\ &\quad = \Pr(v_2 < b_1/k)(v_1 - b_1)\\ &\quad = \frac{b_1}{k}(v_1 - b_1).\end{aligned} \tag{18.68}$$

We have ignored the possibility of equal bids, because they would only occur in equilibrium if buyers had equal valuations yet the probability is zero that two independent and continuous random variables equal each other.

The only tricky step in Equation 18.68 is the last one. The discussion of cumulative distribution functions in Chapter 2 shows that the probability $\Pr(v_2 < x)$ can be written as

$$\Pr(v_2 < x) = \int_{-\infty}^{x} f(v_2)dv_2, \tag{18.69}$$

where f is the probability density function. But for a random variable uniformly distributed between 0 and 1 we have

$$\int_0^x f(v_2)dv_2 = \int_0^x (1)dv_2 = x, \tag{18.70}$$

so $\Pr(v_2 < b_1/k) = b_1/k$.

Taking the first-order condition of Equation 18.68 with respect to b_1 and rearranging yields $b_1 = v_1/2$. Hence $k^* = 1/2$, implying that buyers shade their valuations down by half in forming their bids.

Order statistics. Before computing the seller's expected revenue from the auction, we will introduce the notion of an order statistic. If n independent draws are made from the same distribution and if they are arranged from smallest to largest, then the kth lowest draw is called the *kth-order statistic*, denoted $X_{(k)}$. For example, with n random variables, the nth-order statistic $X_{(n)}$ is the largest of the n draws; the $(n-1)$th-order statistic $X_{(n-1)}$ is the second largest; and so on. Order statistics are so useful that statisticians have done a lot of work to characterize their properties. For instance, statisticians have computed that if n draws are taken from a uniform distribution between 0 and 1, then the expected value of the kth-order statistic is

$$E(X_{(k)}) = \frac{k}{n+1}. \tag{18.71}$$

This formula may be found in many standard statistical references.

Expected revenue. The expected revenue from the first-price auction equals

$$E(\max(b_1, b_2)) = \frac{1}{2}E(\max(v_1,\ v_2)). \tag{18.72}$$

But $\max(v_1,\ v_2)$ is the largest-order statistic from two draws of a uniform random variable between 0 and 1, the expected value of which is 2/3 (according to Equation 18.71). Therefore, the expected revenue from the auction equals $(1/2)(2/3) = 1/3$.

Second-price auction. Suppose that the seller decides to use a second-price auction to sell the painting. In equilibrium, buyers bid their true valuations: $b_i = v_i$. The seller's expected revenue is $E(\min(b_1,\ b_2))$ because the winning bidder pays an amount equal to the loser's bid. But $\min(b_1,\ b_2) = \min(v_1,\ v_2)$, and the latter is the first-order statistic for two draws from a random variable uniformly distributed between 0 and 1 whose expected value is 1/3 (according to Equation 18.71). This is the same expected revenue generated by the first-price auction.

QUERY: In the first-price auction, could the seller try to boost bids up toward buyers' valuations by specifying a reservation price r such that no sale is made if the highest bid does not exceed r? What are the trade-offs involved for the seller from such a reservation price? Would a reservation price help boost revenue in a second-price auction?

In more complicated economic environments, the many different auction formats do not necessarily yield the same revenue. One complication that is frequently considered is supposing that the good has the same value to all bidders but that they do not know exactly what that value is: each bidder has only an imprecise estimate of what his or her valuation might be. For example, bidders for oil tracts may have each conducted their own surveys of the likelihood that there is oil below the surface. All bidders' surveys taken together may give a clear picture of the likelihood of oil, but each one separately may give only a rough idea. For another example, the value of a work of art depends in part on its resale value (unless the bidder plans on keeping it in the family forever), which in turn depends on others' valuations; each bidder knows his or her own valuation but perhaps not others'. An auction conducted in such an economic environment is called a *common values* auction.

The most interesting issue that arises in a common values setting is the so-called winner's curse. The winning bidder realizes that every other bidder probably thought the good was worth less, meaning that he or she probably overestimated the value of the good. The winner's curse sometimes leads inexperienced bidders to regret having won the auction. Sophisticated bidders take account of the winner's curse by shading down their bids below their (imprecise) estimates of the value of the good, so they never regret having won the auction in equilibrium.

Analysis of the common values setting is rather complicated, and the different auction formats previously listed no longer yield equivalent revenue. Roughly speaking, auctions that incorporate other bidders' information in the price paid tend to provide the seller with more revenue. For example, a second-price auction tends to be better than a first-price auction because the price paid in a second-price auction depends on what other bidders think the object is worth. If other bidders thought the object was not worth much, then the second-highest bid will be low and the price paid by the winning bidder will be low, precluding the winner's curse.

Summary

In this chapter we have provided a survey of some issues that arise in modeling markets with asymmetric information. Asymmetric information can lead to market inefficiencies relative to the first-best benchmark, which assumes perfect information. Cleverly designed contracts can often help recover some of this lost surplus. We examined some of the following specific issues.

- Asymmetric information is often studied using a principal-agent model in which a principal offers a contract to an agent who has private information. The two main variants of the principal-agent model are the models of hidden actions and of hidden types.
- In a hidden-action model (called a moral hazard model in an insurance context), the principal tries to induce the agent to take appropriate actions by tying the agent's payments to observable outcomes. Doing so exposes the agent to random fluctuations in these outcomes, which is costly for a risk-averse agent.
- In a hidden-type model (called an adverse selection model in an insurance context), the principal cannot extract all the surplus from high types because they can always gain positive surplus by pretending to be a low type. In an effort to extract the most surplus possible, the principal offers a menu of contracts from which different types of agent can select. The principal distorts the quantity in the contract targeted to low types in order to make this contract less attractive to high types, thus extracting more surplus in the contract targeted to the high types.
- Most of the insights gained from the basic form of the principal-agent model, in which the principal is a monopolist, carry over to the case of competing principals. The main change is that agents obtain more surplus.
- The lemons problem arises when sellers have private information about the quality of their goods. Sellers whose goods are higher than average quality may refrain from selling at the market price, which reflects the average quality of goods sold on the market. The market may collapse, with goods of only the lowest quality being offered for sale.
- The principal can extract more surplus from agents if several of them are pitted against each other in an auction setting. In a simple economic environment, a variety of common auction formats generate the same revenue for the seller. Differences in auction format may generate different levels of revenue in more complicated settings.

Problems

18.1

A personal-injury lawyer works as an agent for his injured plaintiff. The expected award from the trial (taking into account the plaintiff's probability of prevailing and the damage award if she prevails) is l, where l is the lawyer's effort. Effort costs the lawyer $l^2/2$.

a. What is the lawyer's effort, his surplus, and the plaintiff's surplus in equilibrium when the lawyer obtains the customary 1/3 contingency fee (i.e., the lawyer gets 1/3 of the award if the plaintiff prevails)?
b. Repeat part (a) for a general contingency fee of c.
c. What is the optimal contingency fee from the plaintiff's perspective? Compute the associated surpluses for the lawyer and plaintiff.
d. What would be the optimal contingency fee from the plaintiff's perspective if she could "sell" the case to her lawyer [i.e., if she could ask him for an up-front payment in return for a specified contingency fee, possibly higher than in part (c)]? Compute the up-front payment (assuming that the plaintiff makes the offer to the lawyer) and the associated surpluses for the lawyer and plaintiff. Do they do better in this part than in part (c)? Why do you think selling cases in this way is outlawed in many countries?

18.2

Solve for the optimal linear price per ounce of coffee that the coffee shop would charge in Example 18.4. How does the shop's profit compare to when it uses nonlinear prices? *Hint:* Your first step should be to compute each type's demand at a linear price p.