Microeconomic Theory

Andreu Mas-Colell   Michael D. Whinston

and

Jerry R. Green

New York   Oxford   OXFORD UNIVERSITY PRESS   1995

CHAPTER

13

Adverse Selection, Signaling, and Screening

13.A Introduction

One of the implicit assumptions of the fundamental welfare theorems is that the characteristics of all commodities are observable to all market participants. Without this condition, distinct markets cannot exist for goods having differing characteristics, and so the complete markets assumption cannot hold. In reality, however, this kind of information is often asymmetrically held by market participants. Consider the following three examples:

(i) When a firm hires a worker, the firm may know less than the worker does about the worker's innate ability.

(ii) When an automobile insurance company insures an individual, the individual may know more than the company about her inherent driving skill and hence about her probability of having an accident.

(iii) In the used-car market, the seller of a car may have much better information about her car's quality than a prospective buyer does.

A number of questions immediately arise about these settings of *asymmetric information*: How do we characterize market equilibria in the presence of asymmetric information? What are the properties of these equilibria? Are there possibilities for welfare-improving market intervention? In this chapter, we study these questions, which have been among the most active areas of research in microeconomic theory during the last twenty years.

We begin, in Section 13.B, by introducing asymmetric information into a simple competitive market model. We see that in the presence of asymmetric information, market equilibria often fail to be Pareto optimal. The tendency for inefficiency in these settings can be strikingly exacerbated by the phenomenon known as *adverse selection*. Adverse selection arises when an informed individual's trading decisions depend on her privately held information in a manner that adversely affects uninformed market participants. In the used-car market, for example, an individual is more likely to decide to sell her car when she knows that it is not very good. When adverse selection is present, uninformed traders will be wary of any informed trader who wishes to trade with them, and their willingness to pay for the product offered will be low. Moreover, this fact may even further exacerbate the adverse selection problem: If the price that can be received by selling a used car is very low, only sellers with *really* bad cars will offer them for sale. As a result, we may see little trade in markets in which adverse selection is present, even if a great deal of trade would occur were information symmetrically held by all market participants.

We also introduce and study in Section 13.B an important concept for the analysis of market intervention in settings of asymmetric information: the notion of a *constrained Pareto optimal allocation.* These are allocations that cannot be Pareto improved upon by a central authority who, like market participants, cannot observe individuals' privately held information. A Pareto-improving market intervention can be achieved by such an authority only when the equilibrium allocation fails to be a constrained Pareto optimum. In general, the central authority's inability to observe individuals' privately held information leads to a more stringent test for Pareto-improving market intervention.

In Sections 13.C and 13.D, we study how market behavior may adapt in response to these informational asymmetries. In Section 13.C, we consider the possibility that informed individuals may find ways to *signal* information about their unobservable knowledge through observable actions. For example, a seller of a used car could offer to allow a prospective buyer to take the car to a mechanic. Because sellers who have good cars are more likely to be willing to take such an action, this offer can serve as a signal of quality. In Section 13.D, we consider the possibility that uninformed parties may develop mechanisms to distinguish, or *screen*, informed individuals who have differing information. For example, an insurance company may offer two policies: one with no deductible at a high premium and another with a significant deductible at a much lower premium. Potential insureds then *self-select*, with high-ability drivers choosing the policy with a deductible and low-ability drivers choosing the no-deductible policy. In both sections, we consider the welfare characteristics of the resulting market equilibria and the potential for Pareto-improving market intervention.

For expositional purposes, we present all the analysis that follows in terms of the labor market example (i). We should nevertheless emphasize the wide range of settings and fields within economics in which these issues arise. Some of these examples are developed in the exercises at the end of the chapter.

13.B Informational Asymmetries and Adverse Selection

Consider the following simple labor market model adapted from Akerlof's (1970) pioneering work:[1] there are many identical potential firms that can hire workers. Each produces the same output using an identical constant returns to scale technology in which labor is the only input. The firms are risk neutral, seek to maximize their expected profits, and act as price takers. For simplicity, we take the price of the firms' output to equal 1 (in units of a numeraire good).

Workers differ in the number of units of output they produce if hired by a firm,

1. Akerlof (1970) used the example of a used-car market in which only the seller of a used car knows if the car is a "lemon." For this reason, this type of model is sometimes referred to as a *lemons* model.

then the manager gets exactly the same expected utility under $\bar{w}(\pi)$ as under $w(\pi, y)$ for any level of effort he chooses. Thus, the manager's effort choice will be unchanged, and he will still accept the contract. However, because the manager faces less risk, the expected wage payments are lower and the owner is better off (this again follows from Jensen's inequality: for all π, $v(E[w(\pi, y) \mid \pi]) > E[v(w(\pi, y)) \mid \pi]$, and so $\bar{w}(\pi) < E[w(\pi, y) \mid \pi]$).

This point can be pushed further. Note that we can always write

$$f(\pi, y \mid e) = f_1(\pi \mid e) f_2(y \mid \pi, e).$$

If $f_2(y \mid \pi, e)$ does not depend on e, then the $f_2(\cdot)$ terms in condition (14.B.11) again cancel out and the optimal compensation package does not depend on y. This condition on $f_2(y \mid \pi, e)$ is equivalent to the statistical concept that π is a *sufficient statistic* for y with respect to e. The converse is also true: As long as π is *not* a sufficient statistic for y, then wages *should* be made to depend on y, at least to some degree. See Holmstrom (1979) for further details.

A number of extensions of this basic analysis have been studied in the literature. For example, Holmstrom (1982), Nalebuff and Stiglitz (1983), and Green and Stokey (1983) examine cases in which many managers are being hired and consider the use of relative performance evaluation in such settings; Bernheim and Whinston (1986), on the other hand, extend the model in the other direction, examining settings in which a single agent is hired simultaneously by several principals; Dye (1986) considers cases in which effort may be observed through costly monitoring; Rogerson (1985a), Allen (1985), and Fudenberg, Holmstrom, and Milgrom (1990) examine situations in which the agency relationship is repeated over many periods, with a particular focus on the extent to which long-term contracts are more effective at resolving agency problems than is a sequence of short-term contracts of the type we analyzed in this section. (This list of extensions is hardly exhaustive.) Many of these analyses focus on the case in which effort is single-dimensional; Holmstrom and Milgrom (1991) discuss some interesting aspects of the more realistic case of multidimensional effort.

Holmstrom and Milgrom (1987) have pursued another interesting extension. Bothered by the simplicity of real-world compensation schemes relative to the optimal contracts derived in models like the one we have studied here, they investigate a model in which profits accrue incrementally over time and the manager is able to adjust his effort during the course of the project in response to early profit realizations. They identify conditions under which the owner can restrict himself without loss to the use of compensation schemes that are *linear* functions of the project's total profit. The optimality of linear compensation schemes arises because of the need to offer incentives that are "robust" in the sense that they continue to provide incentives regardless of how early profit realizations turn out. Their analysis illustrates a more general idea, namely, that complicating the nature of the incentive problem can actually lead to simpler forms for optimal contracts. For illustrations of this point, see Exercises 14.B.5 and 14.B.6.

The exercises at the end of the chapter explore some of these extensions.

14.C Hidden Information (and Monopolistic Screening)

In this section, we shift our focus to a setting in which the postcontractual informational asymmetry takes the form of hidden information.

Once again, an owner wishes to hire a manager to run a one-time project. Now, however, the manager's effort level, denoted by e, is fully observable. What is not observable after the contract is signed is the random realization of the manager's disutility from effort. For example, the manager may come to find himself well suited to the tasks required at the firm, in which case high effort has a relatively low disutility associated with it, or the opposite may be true. However, only the manager comes to know which case obtains.[10]

Before proceeding, we note that the techniques we develop here can also be applied to models of *monopolistic screening* where, in a setting characterized by *precontractual* informational asymmetries, a single uninformed individual offers a menu of contracts in order to distinguish, or *screen*, informed agents who have differing information at the time of contracting (see Section 13.D for an analysis of a competitive screening model). We discuss this connection further in small type at the end of this section.

To formulate our hidden information principal-agent model, we suppose that effort can be measured by a one-dimensional variable $e \in [0, \infty)$. Gross profits (excluding any wage payments to the manager) are a simple deterministic function of effort, $\pi(e)$, with $\pi(0) = 0$, $\pi'(e) > 0$, and $\pi''(e) < 0$ for all e.

The manager is an expected utility maximizer whose Bernoulli utility function over wages and effort, $u(w, e, \theta)$, depends on a state of nature θ that is realized after the contract is signed and that only the manager observes. We assume that $\theta \in \mathbb{R}$, and we focus on a special form of $u(w, e, \theta)$ that is widely used in the literature:[11]

$$u(w, e, \theta) = v(w - g(e, \theta)).$$

The function $g(e, \theta)$ measures the disutility of effort in monetary units. We assume that $g(0, \theta) = 0$ for all θ and, letting subscripts denote partial derivatives, that

$$g_e(e, \theta) \begin{cases} > 0 & \text{for } e > 0 \\ = 0 & \text{for } e = 0 \end{cases}$$

$$g_{ee}(e, \theta) > 0 \qquad \text{for all } e$$

$$g_\theta(e, \theta) < 0 \qquad \text{for all } e$$

$$g_{e\theta}(e, \theta) \begin{cases} < 0 & \text{for } e > 0 \\ = 0 & \text{for } e = 0. \end{cases}$$

Thus, the manager is averse to increases in effort, and this aversion is larger the greater the current level of effort. In addition, higher values of θ are more productive states in the sense that both the manager's total disutility from effort, $g(e, \theta)$, and his marginal disutility from effort at any current effort level, $g_e(e, \theta)$, are lower when θ

10. A seemingly more important source of hidden information between managers and owners is that the manager of a firm often comes to know more about the potential profitability of various actions than does the owner. In Section 14.D, we discuss one hybrid hidden action–hidden information model that captures this alternative sort of informational asymmetry; its formal analysis reduces to that of the model studied here.

11. Exercise 14.C.3 asks you to consider an alternative form for the manager's utility function.

is greater. We also assume that the manager is strictly risk averse, with $v''(\cdot) < 0$.[12] As in Section 14.B, the manager's reservation utility level, the level of expected utility he must receive if he is to accept the owner's contract offer, is denoted by $\bar{u}$. Note that our assumptions about $g(e, \theta)$ imply that the manager's indifference curves have the single-crossing property discussed in Section 13.C.

Finally, for expositional purposes, we focus on the simple case in which θ can take only one of two values, θ_H and θ_L, with $\theta_H > \theta_L$ and $\text{Prob}(\theta_H) = \lambda \in (0, 1)$. (Exercise 14.C.1 asks you to consider the case of an arbitrary finite number of states.)

A contract must try to accomplish two objectives here: first, as in Section 14.B, the risk-neutral owner should insure the manager against fluctuations in his income; second, although there is no problem here in insuring that the manager puts in effort (because the contract can explicitly state the effort level required), a contract that maximizes the surplus available in the relationship (and hence, the owner's payoff) must make the level of managerial effort responsive to the disutility incurred by the manager, that is, to the state θ. To fix ideas, we first illustrate how these goals are accomplished when θ is observable; we then turn to an analysis of the problems that arise when θ is observed only by the manager.

The State θ is Observable

If θ is observable, a contract can directly specify the effort level and remuneration of the manager contingent on each realization of θ (note that these variables fully determine the economic outcomes for the two parties). Thus, a complete information contract consists of two wage–effort pairs: $(w_H, e_H) \in \mathbb{R} \times \mathbb{R}_+$ for state θ_H and $(w_L, e_L) \in \mathbb{R} \times \mathbb{R}_+$ for state θ_L. The owner optimally chooses these pairs to solve the following problem:

$$\begin{aligned} \underset{\substack{w_L, e_L \geq 0 \\ w_H, e_H \geq 0}}{\text{Max}} \quad & \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L] \\ \text{s.t.} \quad & \lambda\, v(w_H - g(e_H, \theta_H)) + (1 - \lambda) v(w_L - g(e_L, \theta_L)) \geq \bar{u}. \end{aligned} \tag{14.C.1}$$

In any solution $[(w_L^*, e_L^*), (w_H^*, e_H^*)]$ to problem (14.C.1) the reservation utility constraint must bind; otherwise, the owner could lower the level of wages offered and still have the manager accept the contract. In addition, letting $\gamma \geq 0$ denote the multiplier on this constraint, the solution must satisfy the following first-order conditions:

$$-\lambda + \gamma\lambda v'(w_H^* - g(e_H^*, \theta_H)) = 0. \tag{14.C.2}$$

$$-(1 - \lambda) + \gamma(1 - \lambda)\, v'(w_L^* - g(e_L^*, \theta_L)) = 0. \tag{14.C.3}$$

$$\lambda\pi'(e_H^*) - \gamma\lambda v'(w_H^* - g(e_H^*, \theta_H))\, g_e(e_H^*, \theta_H) \begin{cases} \leq 0, & \\ = 0 & \text{if } e_H^* > 0. \end{cases} \tag{14.C.4}$$

$$(1 - \lambda)\pi'(e_L^*) - \gamma(1 - \lambda)\, v'(w_L^* - g(e_L^*, \theta_L))\, g_e(e_L^*, \theta_L) \begin{cases} \leq 0, & \\ = 0 & \text{if } e_L^* > 0. \end{cases} \tag{14.C.5}$$

---

12. As with the case of hidden actions studied in Section 14.B, nonobservability causes no welfare loss in the case of managerial risk neutrality. As there, a "sellout" contract that faces the manager with the full marginal returns from his actions can generate the first-best outcome. (See Exercise 14.C.2.)

These conditions indicate how the two objectives of insuring the manager and making effort sensitive to the state are handled. First, rearranging and combining conditions (14.C.2) and (14.C.3), we see that

$$v'(w_H^* - g(e_H^*, \theta_H)) = v'(w_L^* - g(e_L^*, \theta_L)), \tag{14.C.6}$$

so the manager's marginal utility of income is equalized across states. This is the usual condition for a risk-neutral party optimally insuring a risk-averse individual. Condition (14.C.6) implies that $w_H^* - g(e_H^*, \theta_H) = w_L^* - g(e_L^*, \theta_L)$, which in turn implies that $v(w_H^* - g(e_H^*, \theta_H)) = v(w_L^* - g(e_L^*, \theta_L))$; that is, the manager's utility is equalized across states. Given the reservation utility constraint in (14.C.1), the manager therefore has utility level $\bar{u}$ in each state.

Now consider the optimal effort levels in the two states. Since $g_e(0, \theta) = 0$ and $\pi'(0) > 0$, conditions (14.C.4) and (14.C.5) must hold with equality and $e_i^* > 0$ for $i = 1, 2$. Combining condition (14.C.2) with (14.C.4), and condition (14.C.3) with (14.C.5), we see that the optimal level of effort in state θ_i, e_i^*, satisfies

$$\pi'(e_i^*) = g_e(e_i^*, \theta_i) \qquad \text{for } i = L, H. \tag{14.C.7}$$

This condition says that the optimal level of effort in state θ_i equates the marginal benefit of effort in terms of increased profit with its marginal disutility cost.

The pair (w_i^*, e_i^*) is illustrated in Figure 14.C.1 (note that the wage is depicted on the vertical axis and the effort level on the horizontal axis). As shown, the manager is better off as we move to the northwest (higher wages and less effort), and the owner is better off as we move toward the southeast. Because the manager receives utility level $\bar{u}$ in state θ_i, the owner seeks to find the most profitable point on the manager's state θ_i indifference curve with utility level $\bar{u}$. This is a point of tangency between the manager's indifference curve and one of the owner's isoprofit curves. At this point, the marginal benefit to additional effort in terms of increased profit is exactly equal to the marginal cost borne by the manager.

The owner's profit level in state θ_i is $\Pi_i^* = \pi(e_i^*) - v^{-1}(\bar{u}) - g(e_i^*, \theta_i)$. As shown in Figure 14.C.1, this profit is exactly equal to the distance from the origin to the point at which the owner's isoprofit curve through point (w_i^*, e_i^*) hits the vertical

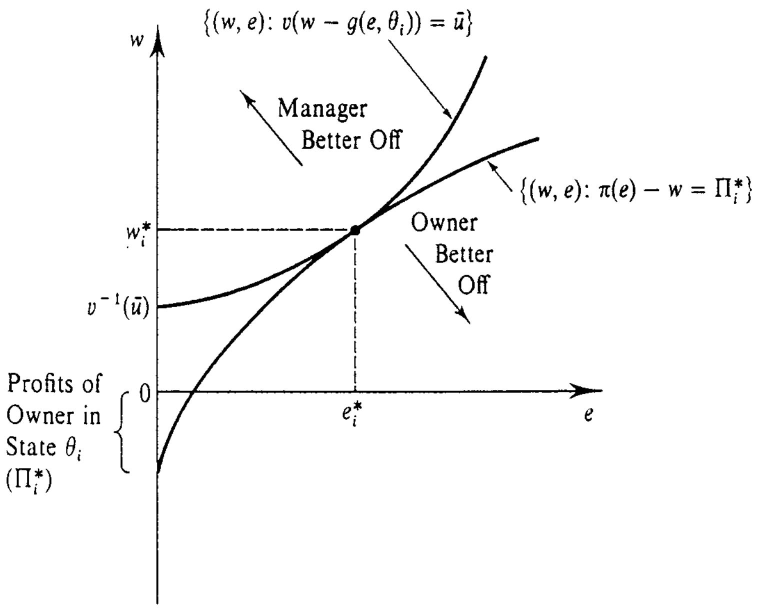

Figure 14.C.1
The optimal wage–effort pair for state θ_i when states are observable.

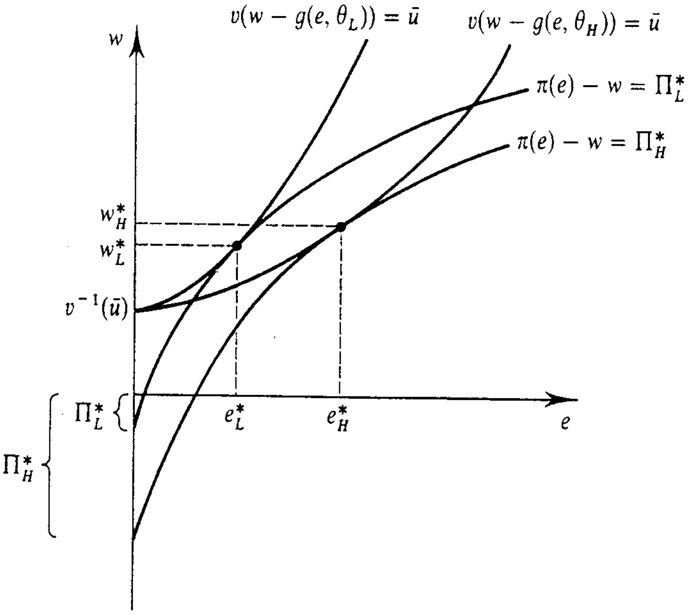

Figure 14.C.2
The optimal contract with full observability of θ.

axis [since $\pi(0) = 0$, if the wage payment at this point on the vertical axis is $\hat{w} < 0$, the owner's profit at (w^*_i, e^*_i) is exactly $-\hat{w}$].

From condition (14.C.7), we see that $g_{e\theta}(e, \theta) < 0$, $\pi''(e) < 0$, and $g_{ee}(e, \theta) > 0$ imply that $e^*_H > e^*_L$. Figure 14.C.2 depicts the optimal contract, $[(w^*_H, e^*_H), (w^*_L, e^*_L)]$.

These observations are summarized in Proposition 14.C.1.

Proposition 14.C.1: In the principal-agent model with an observable state variable θ, the optimal contract involves an effort level e^*_i in state θ_i such that $\pi'(e^*_i) = g_e(e^*_i, \theta_i)$ and fully insures the manager, setting his wage in each state θ_i at the level w^*_i such that $v(w^*_i - g(e^*_i, \theta_i)) = \bar{u}$.

Thus, with a strictly risk-averse manager, the first-best contract is characterized by two basic features: first, the owner fully insures the manager against risk; second, he requires the manager to work to the point at which the marginal benefit of effort exactly equals its marginal cost. Because the marginal cost of effort is lower in state θ_H than in state θ_L, the contract calls for more effort in state θ_H.

The State θ is Observed Only by the Manager

As in Section 14.B, the desire both to insure the risk-averse manager and to elicit the proper levels of effort come into conflict when informational asymmetries are present. Suppose, for example, that the owner offers a risk-averse manager the contract depicted in Figure 14.C.2 and relies on the manager to reveal the state voluntarily. If so, the owner will run into problems. As is evident in the figure, in state θ_H, the manager prefers point (w^*_L, e^*_L) to point (w^*_H, e^*_H). Consequently, in state θ_H he will *lie* to the owner, claiming that it is actually state θ_L. As is also evident in the figure, this misrepresentation lowers the owner's profit.

Given this problem, what is the optimal contract for the owner to offer? To answer this question, it is necessary to start by identifying the set of possible contracts that the owner can offer. One can imagine many different forms that a contract could conceivably take. For example, the owner might offer a compensation function $w(\pi)$ that pays the manager as a function of realized profit and that leaves the effort choice in each state to the manager's discretion. Alternatively, the owner could offer a compensation schedule $w(\pi)$ but restrict the possible effort choices by the manager to some degree. Another possibility is that the owner could offer compensation as a function of the observable effort level chosen by the manager, possibly again with some restriction on the allowable choices. Finally, more complicated arrangements might be imagined. For example, the manager might be required to make an announcement about what the state is and then be free to choose his effort level while facing a compensation function $w(\pi \mid \hat{\theta})$ that depends on his announcement $\hat{\theta}$.

Although finding an optimal contract from among all these possibilities may seem a daunting task, an important result known as the *revelation principle* greatly simplifies the analysis of these types of contracting problems:[13]

Proposition 14.C.2: (*The Revelation Principle*) Denote the set of possible states by Θ. In searching for an optimal contract, the owner can without loss restrict himself to contracts of the following form:

(i) After the state θ is realized, the manager is required to announce which state has occurred.

(ii) The contract specifies an outcome $[w(\hat{\theta}), e(\hat{\theta})]$ for each possible announcement $\hat{\theta} \in \Theta$.

(iii) In every state $\theta \in \Theta$, the manager finds it optimal to report the state *truthfully*.

A contract that asks the manager to announce the state θ and associates outcomes with the various possible announcements is known as a *revelation mechanism*. The revelation principle tells us that the owner can restrict himself to using a revelation mechanism for which the manager always responds truthfully; revelation mechanisms with this truthfulness property are known as *incentive compatible* (or *truthful*) revelation mechanisms. The revelation principle holds in an extremely wide array of incentive problems. Although we defer its formal (and very general) proof to Chapter 23 (see Sections 23.C and 23.D), its basic idea is relatively straightforward.

For example, imagine that the owner is offering a contract with a compensation schedule $w(\pi)$ that leaves the choice of effort up to the manager. Let the resulting levels of effort in states θ_L and θ_H be e_L and e_H, respectively. We can now show that there is a truthful revelation mechanism that generates exactly the same outcome as this contract. In particular, suppose that the owner uses a revelation mechanism that assigns outcome $[w(\pi(e_L)), e_L]$ if the manager announces that the state is θ_L and outcome $[w(\pi(e_H)), e_H]$ if the manager announces that the state is θ_H. Consider the manager's incentives for truth telling when facing this revelation mechanism. Suppose, first, that the state is θ_L. Under the initial contract with compensation schedule $w(\pi)$, the manager could have achieved outcome $[w(\pi(e_H)), e_H]$ in state θ_L by choosing effort level e_H. Since he instead chose e_L, it must be that in state θ_L outcome $[w(\pi(e_L)), e_L]$ is at least as good for the manager as outcome $[w(\pi(e_H)), e_H]$. Thus, under the proposed revelation mechanism, the manager will find telling the truth to be an optimal response when the state is θ_L. A similar argument applies for state θ_H. We see therefore that this revelation mechanism results in truthful announcements

13. Two early discussions of the revelation principle are Myerson (1979) and Dasgupta, Hammond, and Maskin (1979).

by the manager and yields exactly the same outcome as the initial contract. In fact, a similar argument can be constructed for *any* initial contract (see Chapter 23), and so the owner can restrict his attention without loss to truthful revelation mechanisms.[14]

To simplify the characterization of the optimal contract, we restrict attention from this point on to a specific and extreme case of managerial risk aversion: *infinite* risk aversion. In particular, we take the expected utility of the manager to equal the manager's lowest utility level across the two states. Thus, for the manager to accept the owner's contract, it must be that the manager receives a utility of at least $\bar{u}$ in each state.[15] As above, efficient risk sharing requires that an infinitely risk-averse manager have a utility level equal to $\bar{u}$ in each state. If, for example, his utility is $\bar{u}$ in one state and $u' > \bar{u}$ in the other, then the owner's expected wage payment is larger than necessary for giving the manager an expected utility of $\bar{u}$.

Given this assumption about managerial risk preferences, the revelation principle allows us to write the owner's problem as follows:

$$\operatorname*{Max}_{w_H, e_H \geq 0,\, w_L, e_L \geq 0} \quad \lambda[\pi(e_H) - w_H] + (1-\lambda)[\pi(e_L) - w_L] \tag{14.C.8}$$

$$\text{s.t.} \quad \left.\begin{array}{l}\text{(i) } w_L - g(e_L, \theta_L) \geq v^{-1}(\bar{u}) \\ \text{(ii) } w_H - g(e_H, \theta_H) \geq v^{-1}(\bar{u})\end{array}\right\} \begin{array}{l}\textit{reservation utility} \\ \text{(or } \textit{individual rationality}\text{)} \\ \textit{constraint}\end{array}$$

$$\left.\begin{array}{l}\text{(iii) } w_H - g(e_H, \theta_H) \geq w_L - g(e_L, \theta_H) \\ \text{(iv) } w_L - g(e_L, \theta_L) \geq w_H - g(e_H, \theta_L)\end{array}\right\} \begin{array}{l}\textit{incentive compatibility} \\ \text{(or } \textit{truth-telling} \\ \text{or } \textit{self-selection}\text{)} \\ \textit{constraints}.\end{array}$$

The pairs (w_H, e_H) and (w_L, e_L) that the contract specifies are now the wage and effort levels that result from different *announcements* of the state by the manager; that is, the outcome if the manager announces that the state is θ_i is (w_i, e_i). Constraints (i) and (ii) make up the *reservation utility* (or *individual rationality*) *constraint* for the infinitely risk-averse manager; if he is to accept the contract, he must be guaranteed a utility of at least $\bar{u}$ in each state. Hence, we must have $v(w_i - g(e_i, \theta_i)) \geq \bar{u}$ for $i = L, H$ or, equivalently, $w_i - g(e_i, \theta_i) \geq v^{-1}(\bar{u})$ for $i = L, H$. Constraints (iii) and (iv) are the *incentive compatibility* (or *truth-telling* or *self-selection*) *constraints* for the manager in states θ_H and θ_L, respectively. Consider, for example, constraint (iii). The

---

14. One restriction that we have imposed here for expositional purposes is to limit the outcomes specified following the manager's announcement to being nonstochastic (in fact, much of the literature does so as well). Randomization can sometimes be desirable in these settings because it can aid in satisfying the incentive compatibility constraints that we introduce in problem (14.C.8). See Maskin and Riley (1984a) for an example.

15. This can be thought of as the limiting case in which, starting from the concave utility function $v(x)$, we take the concave transformation $v_\rho(v) = -v(x)^\rho$ for $\rho < 0$ as the manager's Bernoulli utility function and let $\rho \to -\infty$. To see this, note that the manager's expected utility over the random outcome giving $(w_H - g(e_H, \theta_H))$ with probability λ and $(w_L - g(e_L, \theta_L))$ with probability $(1-\lambda)$ is then $EU = -[\lambda v_H^\rho + (1-\lambda)v_L^\rho]$, where $v_i = v(w_i - g(e_i, \theta_i))$ for $i = L, H$. This expected utility is correctly ordered by $(-EU)^{1/\rho} = [\lambda v_H^\rho + (1-\lambda)v_L^\rho]^{1/\rho}$. Now as $\rho \to -\infty$, $[\lambda v_H^\rho + (1-\lambda)v_L^\rho]^{1/\rho} \to \operatorname{Min}\{v_H, v_L\}$ (see Exercise 3.C.6). Hence, a contract gives the manager an expected utility greater than his (certain) reservation utility if and only if $\operatorname{Min}\{v(w_H - g(e_H, \theta_H)), v(w_L - g(e_L, \theta_L))\} \geq \bar{u}$.

manager's utility in state θ_H is $v(w_H - g(e_H, \theta_H))$ if he tells the truth, but it is $v(w_L - g(e_L, \theta_H))$ if he instead claims that it is state θ_L. Thus, he will tell the truth if $w_H - g(e_H, \theta_H) \geq w_L - g(e_L, \theta_H)$. Constraint (iv) follows similarly.

Note that the first-best (full observability) contract depicted in Figure 14.C.2 does not satisfy the constraints of problem (14.C.8) because it violates constraint (iii).

We analyze problem (14.C.8) through a sequence of lemmas. Our arguments for these results make extensive use of graphical analysis to build intuition. An analysis of this problem using Kuhn–Tucker conditions is presented in Appendix B.

Lemma 14.C.1: We can ignore constraint (ii). That is, a contract is a solution to problem (14.C.8) if and only if it is the solution to the problem derived from (14.C.8) by dropping constraint (ii).

Proof: Whenever both constraints (i) and (iii) are satisfied, it must be that $w_H - g(e_H, \theta_H) \geq w_L - g(e_L, \theta_H) \geq w_L - g(e_L, \theta_L) \geq v^{-1}(\bar{u})$, and so constraint (ii) is also satisfied. This implies that the set of feasible contracts in the problem derived from (14.C.8) by dropping constraint (ii) is exactly the same as the set of feasible contracts in problem (14.C.8). ■

Lemma 14.C.1 is illustrated in Figure 14.C.3. By constraint (i), (w_L, e_L) must lie in the shaded region of the figure. But by constraint (iii), (w_H, e_H) must lie on or above the state θ_H indifference curve through point (w_L, e_L). As can be seen, this implies that the manager's state θ_H utility is at least $\bar{u}$, the utility he gets at point $(w, e) = (v^{-1}(\bar{u}), 0)$.

Therefore, from this point on we can ignore constraint (ii).

Lemma 14.C.2: An optimal contract in problem (14.C.8) must have $w_L - g(e_L, \theta_L) = v^{-1}(\bar{u})$.

Proof: Suppose not, that is, that there is an optimal solution $[(w_L, e_L), (w_H, e_H)]$ in which $w_L - g(e_L, \theta_L) > v^{-1}(\bar{u})$. Now, consider an alteration to the owner's contract

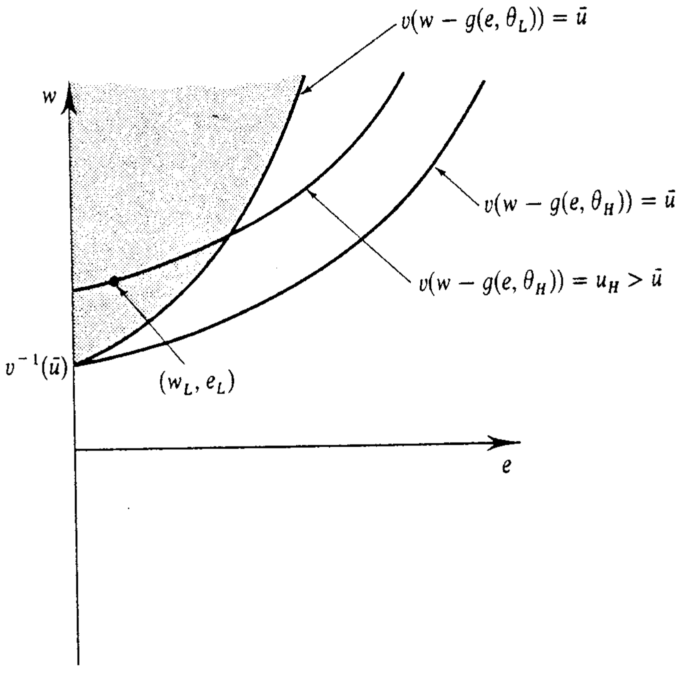

Figure 14.C.3 Constraint (ii) in problem (14.C.8) is satisfied by any contract satisfying constraints (i) and (iii).

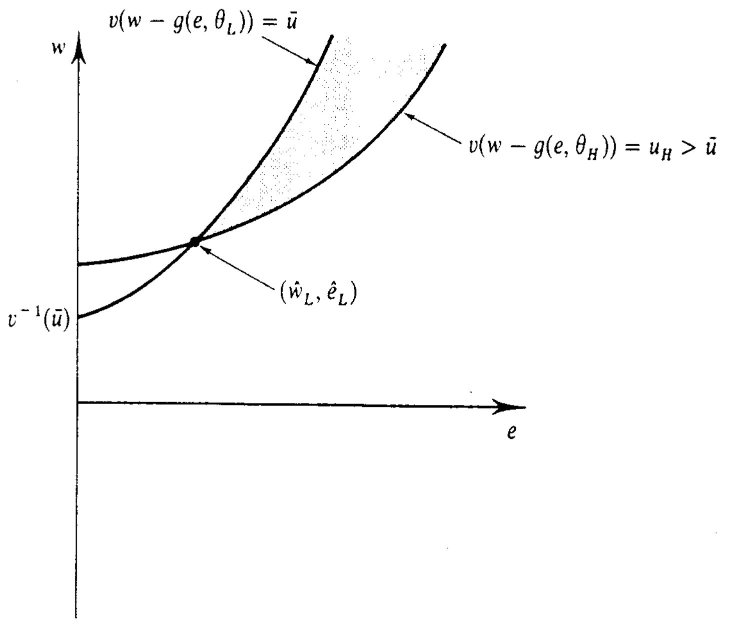

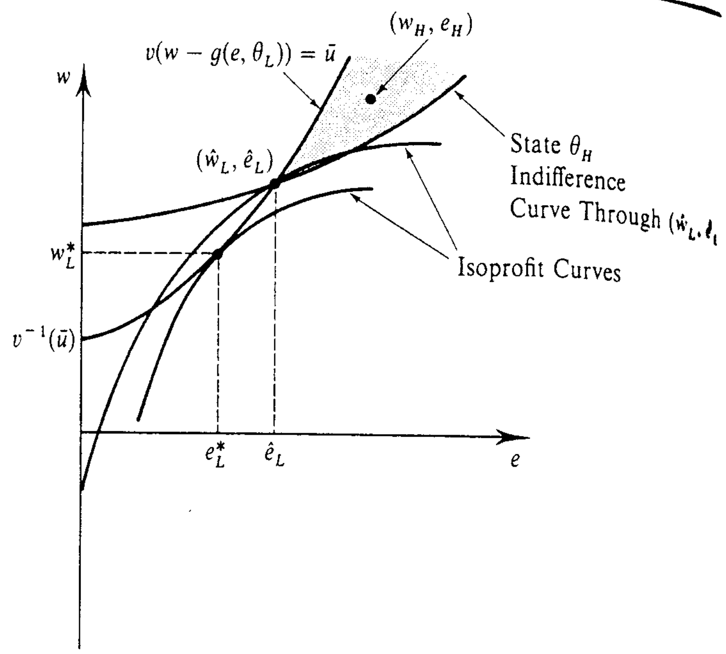

in which the owner pays wages in the two states of $\hat{w}_L = w_L - \varepsilon$ and $\hat{w}_H = w_H - \varepsilon$, where $\varepsilon > 0$ (i.e., the owner lowers the wage payments in both states by ε). This new contract still satisfies constraint (i) as long as ε is chosen small enough. In addition, the incentive compatibility constraints are still satisfied because this change just subtracts a constant, ε, from each side of these constraints. But if this new contract satisfies all the constraints, the original contract could not have been optimal because the owner now has higher profits, which is a contradiction. ■

Figure 14.C.4 (left) In a feasible contract offering $(\hat{w}_L, \hat{e}_L)$ for state θ_L, the pair (w_H, e_H) must lie in the shaded region.

Figure 14.C.5 (right) An optimal contract has $e_L \le e_L^*$.

Lemma 14.C.3: In any optimal contract:

(i) $e_L \le e_L^*$; that is, the manager's effort level in state θ_L is no more than the level that would arise if θ were observable.

(ii) $e_H = e_H^*$; that is, the manager's effort level in state θ_H is exactly equal to the level that would arise if θ were observable.

Proof: Lemma 14.C.3 can best be seen graphically. By Lemma 14.C.2, (w_L, e_L) lies on the locus $\{(w, e): v(w - g(e, \theta_L)) = \bar{u}\}$ in any optimal contract. Figure 14.C.4 depicts one possible pair $(\hat{w}_L, \hat{e}_L)$. In addition, the truth-telling constraints imply that the outcome for state θ_H, (w_H, e_H), must lie in the shaded region of Figure 14.C.4. To see this, note that by constraint (iv), (w_H, e_H) must lie on or below the state θ_L indifference curve through $(\hat{w}_L, \hat{e}_L)$. In addition, by constraint (iii), (w_H, e_H) must lie on or above the state θ_H indifference curve through $(\hat{w}_L, \hat{e}_L)$.

To see part (i), suppose that we have a contract with $\hat{e}_L > e_L^*$. Figure 14.C.5 depicts such a contract offer: $(\hat{w}_L, \hat{e}_L)$ lies on the manager's state θ_L indifference curve with utility level $\bar{u}$, and (w_H, e_H) lies in the shaded region defined by the truth-telling constraints. The state θ_L indifference curve for the manager and the isoprofit curve for the owner which go through point $(\hat{w}_L, \hat{e}_L)$ have the relation depicted at point $(\hat{\omega}_L, \hat{e}_L)$ because $\hat{e}_L > e_L^*$.

As can be seen in the figure, the owner can raise his profit level in state θ_L by moving the state θ_L wage–effort pair down the manager's indifference curve from $(\hat{w}_L, \hat{e}_L)$ to its first-best point (w_L^*, e_L^*). This change continues to satisfy all the constraints in problem (14.C.8): The manager's utility in each state is unchanged,

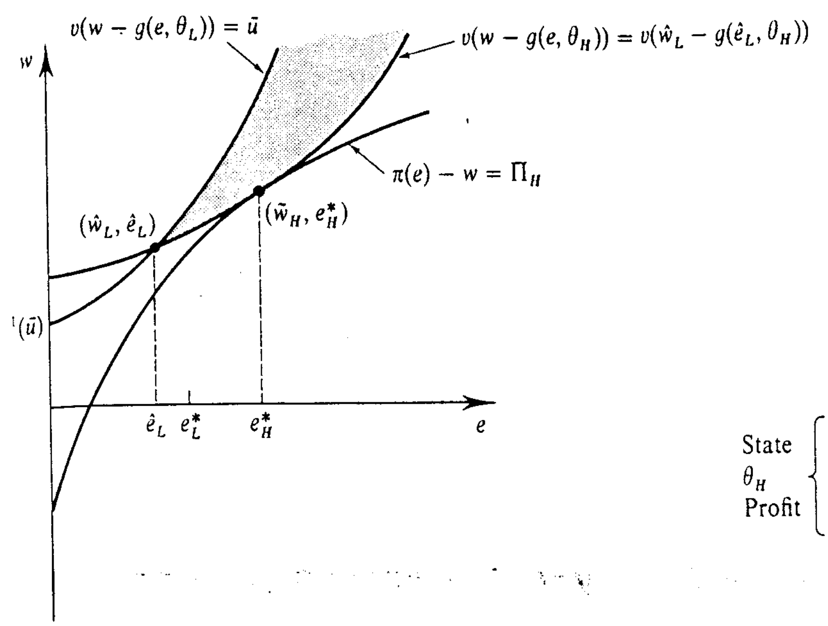

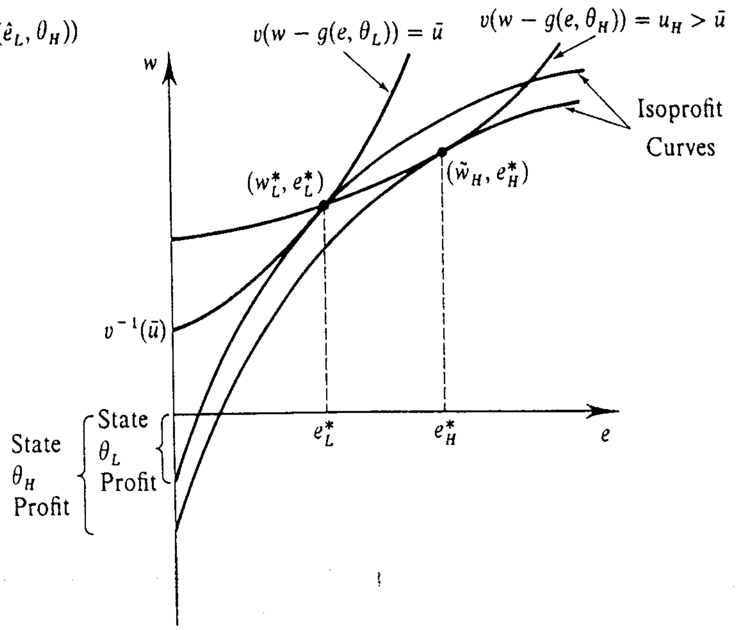

Figure 14.C.6 (left) An optimal contract has $e_H = e_H^*$.

Figure 14.C.7 (right) The best contract with $e_L = e_L^*$.

and, as is evident in Figure 14.C.5, the truth-telling constraints are still satisfied. Thus, a contract with $\hat{e}_L > e_L^*$ cannot be optimal.

Now consider part (ii). Given any wage–effort pair $(\hat{w}_L, \hat{e}_L)$ with $\hat{e}_L \le e_L^*$, such as that shown in Figure 14.C.6, the owner's problem is to find the location for (w_H, e_H) in the shaded region that maximizes his profit in state θ_H. The solution occurs at a point of tangency between the manager's state θ_H indifference curve through point $(\hat{w}_L, \hat{e}_H)$ and an isoprofit curve for the owner. This tangency occurs at point $(\tilde{w}_H, e_H^*)$ in the figure, and necessarily involves effort level e_H^* because all points of tangency between the manager's state θ_H indifference curves and the owner's isoprofit curves occur at effort level e_H^* [they are characterized by condition (14.C.7) for $i = H$]. Note that this point of tangency occurs strictly to the right of effort level $\hat{e}_L$ because $\hat{e}_L \le e_L^* < e_H^*$. ■

A secondary point emerging from the proof of Lemma 14.C.3 is that only the truth-telling constraint for state θ_H is binding in the optimal contract. This property is common to many of the other applications in the literature.[16]

Lemma 14.C.4: In any optimal contract, $e_L < e_L^*$; that is, the effort level in state θ_L is necessarily *strictly* below the level that would arise in state θ_L if θ were observable.

Proof: Again, this point can be seen graphically. Suppose we start with $(w_L, e_L) = (w_L^*, e_L^*)$, as in Figure 14.C.7. By Lemma 14.C.3, this determines the state θ_H outcome, denoted by $(\tilde{w}_H, e_H^*)$ in the figure. Note that by the definition of (w_L^*, e_L^*), the isoprofit curve through this point is tangent to the manager's state θ_L indifference curve.

Recall that the absolute distance between the origin and the point where each state's isoprofit curve hits the vertical axis represents the profit the owner earns in that state. The owner's overall expected profit with this contract offer is therefore

16. In models with more than two types, this property takes the form that only the incentive constraints between adjacent types bind, and they do so only in one direction. (See Exercise 14.C.1.)

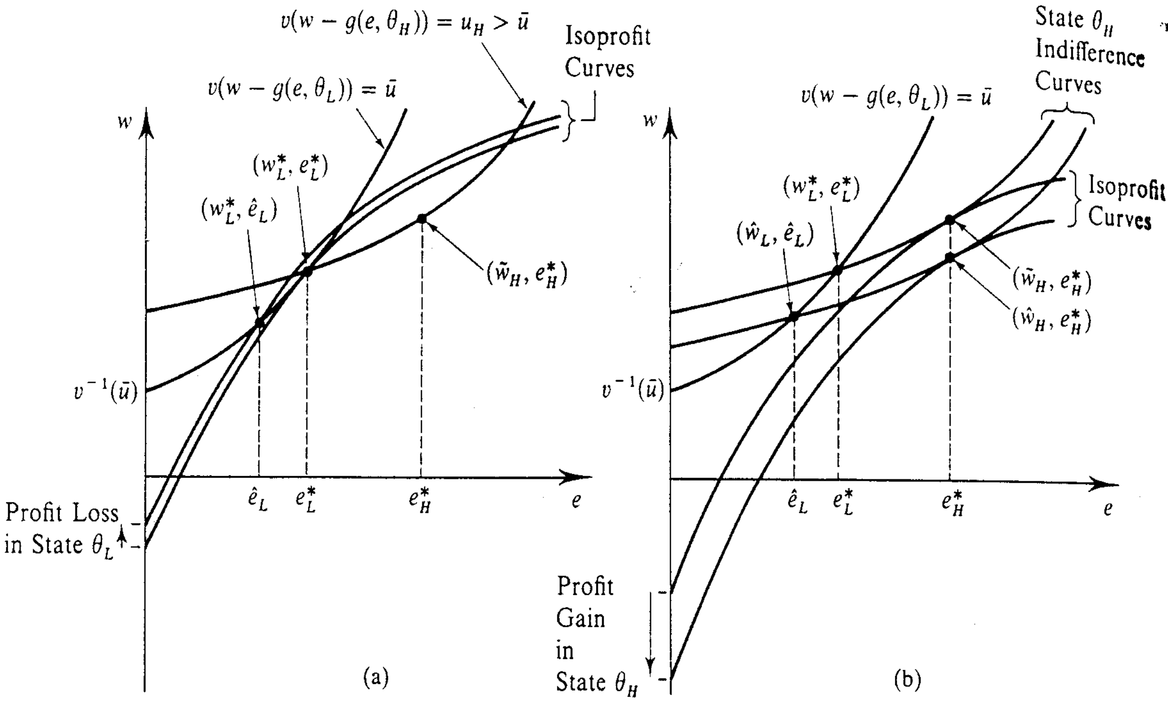

Figure 14.C.8 (a) The change in profits in state θ_L from lowering e_L slightly below e_L^*. (b) The change in profits in state from lowering e_L slightly below e_L^* and optimally adjusting w_H.

equal to the average of these two profit levels (with weights equal to the relative probabilities of the two states).

We now argue that a change in the state θ_L outcome that lowers this state's effort level to one slightly below e_L^* necessarily raises the owner's expected profit. To see this, start by moving the state θ_L outcome to a slightly lower point, $(\hat{w}_L, \hat{e}_L)$, on the manager's state θ_L indifference curve. This change is illustrated in Figure 14.C.8, along with the owner's isoprofit curve through this new point. As is evident in Figure 14.C.8(a), this change lowers the profit that the owner earns in state θ_L. However, it also relaxes the incentive constraint on the state θ_H outcome and, by doing so, it allows the owner to offer a lower wage in that state. Figure 14.C.8(b) shows the new state θ_H outcome, say $(\hat{w}_H, e_H^*)$, and the new (higher-profit) isoprofit curve through this point.

Overall, this change results in a lower profit for the owner in state θ_L and a higher profit for the owner in state θ_H. Note, however, that because we started at a point of tangency at (w_L^*, e_L^*), the profit loss in state θ_L is small relative to the gain in state θ_H. Indeed, if we were to look at the derivative of the owner's profit in state θ_L with respect to an *infinitesimal* change in that state's outcome, we would find that it is *zero*. In contrast, the derivative of profit in state θ_H with respect to this infinitesimal change would be strictly positive. The zero derivative in state θ_L is an envelope theorem result: because we started out at the first-best level of effort in state θ_L, a small change in (w_L, e_L) that keeps the manager's state θ_L utility at $\bar{u}$ has no first-order effect on the owner's profit in that state; but because it relaxes the state θ_H incentive constraint, for a small-enough change the owner's expected profit is increased. ■

How far should the owner go in lowering e? In answering this question, the owner must weigh the marginal loss in profit in state θ_L against the marginal gain in state

θ_H [note that once we move away from (w_L^*, e_L^*), the envelope result no longer holds and the marginal reduction in state θ_L's profit is strictly positive]. It should not be surprising that the extent to which the owner wants to make this trade-off depends on the relative probabilities of the two states. In particular, the greater the likelihood of state θ_H, the more the owner is willing to distort the state θ_L outcome to increase profit in state θ_H. In the extreme case in which the probability of state θ_L gets close to zero, the owner may set $e_L = 0$ and hire the manager to work only in state θ_H.[17]

The analysis in Appendix B confirms this intuition. There we show that the optimal level of e_L satisfies the following first-order condition:

$$[\pi'(e_L) - g_e(e_L, \theta_L)] + \frac{\lambda}{1-\lambda}[g_e(e_L, \theta_H) - g_e(e_L, \theta_L)] = 0. \tag{14.C.9}$$

The first term of this expression is zero at $e_L = e_L^*$ and is strictly positive at $e_L < e_L^*$; the second term is always strictly negative. Thus, we must have $e_L < e_L^*$ to satisfy this condition, confirming our finding in Lemma 14.C.4. Differentiating this expression reveals that the optimal level of e_L falls as $\lambda/(1-\lambda)$ rises.

These findings are summarized in Proposition 14.C.3.

Proposition 14.C.3: In the hidden information principal-agent model with an infinitely risk-averse manager the optimal contract sets the level of effort in state θ_H at its first-best (full observability) level e_H^*. The effort level in state θ_L is distorted downward from its first-best level e_L^*. In addition, the manager is inefficiently insured, receiving a utility greater than $\bar{u}$ in state θ_H and a utility equal to $\bar{u}$ in state θ_L. The owner's expected payoff is strictly lower than the expected payoff he receives when θ is observable, while the infinitely risk-averse manager's expected utility is the same as when θ is observable (it equals $\bar{u}$).[18,19]

A basic, and very general, point that emerges from this analysis is that the optimal contract for the owner in this setting of hidden information necessarily *distorts* the effort choice of the manager in order to ameliorate the costs of asymmetric information, which here take the form of the higher expected wage payment that the owner makes because the manager has a utility in state θ_H in excess of $\bar{u}$.

Note that nothing would change if the profit level π were not publicly observable (and so could not be contracted on), since our analysis relied only on the fact that the effort level e was observable. Moreover, in the case in which π is not publicly observable, we can extend the model to allow the relationship between profits and effort to depend on the state; that is, the owner's profits in states θ_L and θ_H given effort level e might be given by the functions $\pi_L(e)$ and $\pi_H(e)$.[20] As long as

17. In fact, this can happen only if $g_e(0, \theta_L) > 0$.

18. Recall that an infinitely risk-averse manager's expected utility is equal to his lowest utility level across the two states.

19. Note, however, that while the outcome here is Pareto inefficient, it is a constrained Pareto optimum in the sense introduced in Section 13.B; the reasons parallel those given in footnote 9 of Section 14.B for the hidden action model (although here it is θ that the authority cannot observe rather than e).

20. The nonobservability of profits is important for this extension because if π could be contracted upon, the manager could be punished for misrepresenting the state by simply comparing the realized profit level with the profit level that should have been realized in the announced state for the specified level of effort.

$\pi'_H(e) \geq \pi'_L(e) > 0$ for all $e \geq 0$, the analysis of this model follows exactly along the lines of the analysis we have just conducted (see Exercise 14.C.5).

As in the case of hidden action models, a number of extensions of this basic hidden information model have been explored in the literature. Some of the most general treatments appear in the context of the "mechanism design" literature associated with social choice theory. A discussion of these models can be found in Chapter 23.

The Monopolistic Screening Model

In Section 13.D, we studied a model of *competitive screening* in which firms try to design their employment contracts in a manner that distinguishes among workers who, at the time of contracting, have different unobservable productivity levels (i.e., there is *precontractual* asymmetric information). The techniques that we have developed in our study of the principal-agent model with hidden information enable us to formulate and solve a model of *monopolistic screening* in which, in contrast with the analysis in Section 13.D, only a single firm offers employment contracts (actually, this might more properly be called a *monopsonistic* screening model because the single firm is on the demand side of the market).

To see this, suppose that, as in the model in Section 13.D, there are two possible types of workers who differ in their productivity. A worker of type θ has utility $u(w, t \mid \theta) = w - g(t, \theta)$ when he receives a wage of w and faces task level t. His reservation utility level is $\bar{u}$. The productivities of the two types of workers are θ_H and θ_L, with $\theta_H > \theta_L > 0$. The fraction of workers of type θ_H is $\lambda \in (0, 1)$. We assume that the firm's profits, which are not publicly observable, are given by the function $\pi_H(t)$ for a type θ_H worker and by $\pi_L(t)$ for a type θ_L worker, and that $\pi'_H(t) \geq \pi'_L(t) > 0$ for all $t \geq 0$ [e.g., as in Exercise 13.D.1, we could have $\pi_i(t) = \theta_i(1 - \mu t)$ for $\mu > 0$].[21]

The firm's problem is to offer a set of contracts that maximizes its profits given worker self-selection among, and behavior within, its offered contracts. Once again, the revelation principle can be invoked to greatly simplify the firm's problem. Here the firm can restrict its attention to offering a menu of wage–task pairs $[(w_H, t_H), (w_L, t_L)]$ to solve

$$\max_{w_H, t_H \geq 0,\, w_L, t_L \geq 0} \quad \lambda[\pi_H(t_H) - w_H] + (1 - \lambda)[\pi_L(t_L) - w_L] \tag{14.C.10}$$

$$\begin{aligned}
\text{s.t.}\quad &\text{(i) } w_L - g(t_L, \theta_L) \geq \bar{u} \\
&\text{(ii) } w_H - g(t_H, \theta_H) \geq \bar{u} \\
&\text{(iii) } w_H - g(t_H, \theta_H) \geq w_L - g(t_L, \theta_H) \\
&\text{(iv) } w_L - g(t_L, \theta_L) \geq w_H - g(t_H, \theta_L).
\end{aligned}$$

This problem has exactly the same structure as (14.C.8) but with the principal's (here the firm's) profit being a function of the state. As noted above, the analysis of this problem follows exactly the same lines as our analysis of problem (14.C.8).

This class of models has seen wide application in the literature (although often with a continuum of types assumed). Maskin and Riley (1984b), for example, apply this model to the study of monopolistic price discrimination. In their model, a consumer of type θ has utility $v(x, \theta) - T$ when he consumes x units of a monopolist's good and makes a total payment of T to the monopolist, and can earn a reservation utility level of $v(0, \theta) = 0$ by not purchasing from the monopolist. The monopolist has a constant unit cost of production equal to $c > 0$

---

21. The model studied in Section 13.D with $\pi_i(t) = \theta_i$ corresponds to the limiting case where $\mu \to 0$.

and seeks to offer a menu of (x_i, T_i) pairs to maximize its profit. The monopolist's problem then takes the form in (14.C.10) where we take $t_i = x_i$, $w_i = -T_i$, $\bar{u} = 0$, $g(t_i, \theta_i) = -v(x_i, \theta_i)$, and $\pi_i(t_i) = -cx_i$.

Baron and Myerson's (1982) analysis of optimal regulation of a monopolist with unknown costs provides another example. There, a regulated firm faces market demand function $x(p)$ and has unobservable unit costs of θ. The regulator, who seeks to design a regulatory policy that maximizes consumer surplus, faces the monopolist with a choice among a set of pairs (p_i, T_i), where p_i is the allowed retail price and T_i is a transfer payment from the regulator to the firm. The regulated firm is able to shut down if it cannot earn profits of at least zero from any of the regulator's offerings. The regulator's problem then corresponds to (14.C.10) with $t_i = p_i$, $w_i = T_i$, $\bar{u} = 0$, $g(t_i, \theta_i) = -(p_i - \theta_i)x(p_i)$, and $\pi_i(t_i) = \int_{p_i}^{\infty} x(s)\, ds$.[22]

Exercises 14.C.7 to 14.C.9 ask you to study some examples of monopolistic screening models.

14.D Hidden Actions and Hidden Information: Hybrid Models

Although the hidden action–hidden information dichotomization serves as a useful starting point for understanding principal-agent models, many real-world situations (and some of the literature as well) involve elements of both problems.

To consider an example of such a model, suppose that we augment the simple hidden information model considered in Section 14.C in the following manner: let the level of effort e now be unobservable, and let profits be a stochastic function of effort, described by conditional density function $f(\pi \mid e)$. In essence, what we now have is a hidden action model, but one in which the owner also does not know something about the disutility of the manager (which is captured in the state variable θ).

Formal analysis of this model is beyond the scope of this chapter, but the basic thrust of the revelation principle extends to the analysis of these types of hybrid problems. In particular, as Myerson (1982) shows, the owner can now restrict attention to contracts of the following form:

(i) After the state θ is realized, the manager announces which state has occurred.

(ii) The contract specifies, for each possible announcement $\hat{\theta} \in \Theta$, the effort level $e(\hat{\theta})$ that the manager should take and a compensation scheme $w(\pi \mid \hat{\theta})$.

(iii) In every state θ, the manager is willing to be both *truthful* in stage (i) and *obedient* following stage (ii) [i.e., he finds it optimal to choose effort level $e(\theta)$ in state θ].

This contract can be thought of as a revelation game, but one in which the outcome of the manager's announcement about the state is a hidden action-style contract, that is, a compensation scheme and a "recommended action." The requirement of "obedience" amounts to an incentive constraint that is like that in the hidden action

22. The regulator's objective function can be generalized to allow a weighted average of consumer and producer surplus, with greater weight on consumers. In this case, the function $\pi_i(\cdot)$ will depend on θ_i.