CHAPTER

14

The Principal-Agent Problem

14.A Introduction

In Chapter 13, we considered situations in which asymmetries of information exist between individuals at the time of contracting. In this chapter, we shift our attention to asymmetries of information that develop *subsequent* to the signing of a contract.

Even when informational asymmetries do not exist at the time of contracting, the parties to a contract often anticipate that asymmetries will develop sometime after the contract is signed. For example, after an owner of a firm hires a manager, the owner may be unable to observe how much effort the manager puts into the job. Similarly, the manager will often end up having better information than the owner about the opportunities available to the firm.

Anticipating the development of such informational asymmetries, the contracting parties seek to design a contract that mitigates the difficulties they cause. These problems are endemic to situations in which one individual hires another to take some action for him as his "agent." For this reason, this contract design problem has come to be known as the *principal-agent problem*.

The literature has traditionally distinguished between two types of informational problems that can arise in these settings: those resulting from *hidden actions* and those resulting from *hidden information*. The hidden action case, also known as *moral hazard*, is illustrated by the owner's inability to observe how hard his manager is working; the manager's coming to possess superior information about the firm's opportunities, on the other hand, is an example of hidden information.[1]

Although many economic situations (and some of the literature) contain elements of both types of problems, it is useful to begin by studying each in isolation. In Section 14.B, we introduce and study a model of hidden actions. Section 14.C analyzes

1. The literature's use of the term *moral hazard* is not entirely uniform. The term originates in the insurance literature, which first focused attention on two types of informational imperfections: the "moral hazard" that arises when an insurance company cannot observe whether the insured exerts effort to prevent a loss and the "adverse selection" (see Section 13.B) that occurs when the insured knows more than the company at the time he purchases a policy about his likelihood of an accident. Some authors use moral hazard to refer to either of the hidden action or hidden information variants of the principal-agent problem [see, for example, Hart and Holmstrom (1987)]. Here, however, we use the term in the original sense.

a hidden information model. Then, in Section 14.D, we provide a brief discussion of hybrid models that contain both of these features. We shall see that the presence of postcontractual asymmetric information often leads to welfare losses for the contracting parties relative to what would be achievable in the absence of these informational imperfections.

It is important to emphasize the broad range of economic relationships that fit into the general framework of the principal-agent problem. The owner–manager relationship is only one example; others include insurance companies and insured individuals (the insurance company cannot observe how much care is exercised by the insured), manufacturers and their distributors (the manufacturer may not be able to observe the market conditions faced by the distributor), a firm and its workforce (the firm may have more information than its workers about the true state of demand for its products and therefore about the value of the workers' product), and banks and borrowers (the bank may have difficulty observing whether the borrower uses the loaned funds for the purpose for which the loan was granted). As would be expected given this diversity of examples, the principal-agent framework has found application in a broad range of applied fields in economics. Our discussion will focus on the owner–manager problem.

The analysis in this chapter, particularly that in Section 14.C, is closely related to that in two other chapters. First, the techniques developed in Section 14.C can be applied to the analysis of screening problems in which, in contrast with the case studied in Section 13.D, only one uninformed party screens informed individuals. We discuss the analysis of this *monopolistic screening problem* in small type at the end of Section 14.C. Second, the principal-agent problem is actually a special case of "mechanism design," the topic of Chapter 23. Thus, the material here constitutes a first pass at this more general issue. Mastery of the fundamentals of the principal-agent problem, particularly the material in Section 14.C, will be helpful when you study Chapter 23.

A good source for further reading on topics of this chapter is Hart and Holmstrom (1987).

14.B Hidden Actions (Moral Hazard)

Imagine that the owner of a firm (the *principal*) wishes to hire a manager (the *agent*) for a one-time project. The project's profits are affected, at least in part, by the manager's actions. If these actions were observable, the contracting problem between the owner and the manager would be relatively straightforward; the contract would simply specify the exact actions to be taken by the manager and the compensation (wage payment) that the owner is to provide in return.[2] When the manager's actions are not observable, however, the contract can no longer specify them in an effective manner, because there is simply no way to verify whether the manager has fulfilled his obligations. In this circumstance, the owner must design the manager's compensation scheme in a way that *indirectly* gives him the incentive to take the correct

2. Note that this requires not only that the manager's actions be observable to the owner but also that they be observable to any court that might be called upon to enforce the contract.

actions (those that would be contracted for if his actions were observable). In this section, we study this contract design problem.

To be more specific, let π denote the project's (observable) profits, and let e denote the manager's action choice. The set of possible actions is denoted by E. We interpret e as measuring managerial effort. In the simplest case that is widely studied in the literature, e is a one-dimensional measure of how "hard" the manager works, and so $E \subset \mathbb{R}$. More generally, however, managerial effort can have many dimensions—how hard the manager works to reduce costs, how much time he spends soliciting customers, and so on—and so e could be a vector with each of its elements measuring managerial effort in a distinct activity. In this case, $E \subset \mathbb{R}^M$ for some M.[3] In our discussion, we shall refer to e as the manager's *effort choice* or *effort level.*

For the nonobservability of managerial effort to have any consequence, the manager's effort must not be perfectly deducible from observation of π. Hence, to make things interesting (and realistic), we assume that although the project's profits are affected by e, they are not fully determined by it. In particular, we assume that the firm's profit can take values in $[\underline{\pi}, \bar{\pi}]$ and that it is stochastically related to e in a manner described by the conditional density function $f(\pi \mid e)$, with $f(\pi \mid e) > 0$ for all $e \in E$ and all $\pi \in [\underline{\pi}, \bar{\pi}]$. Thus, any potential realization of π can arise following any given effort choice by the manager.

In the discussion that follows, we restrict our attention to the case in which the manager has only two possible effort choices, e_H and e_L (see Appendix A for a discussion of the case in which the manager has many possible actions), and we make assumptions implying that e_H is a "high-effort" choice that leads to a higher profit level for the firm than e_L but entails greater difficulty for the manager. This fact will mean that there is a conflict between the interests of the owner and those of the manager.

More specifically, we assume that the distribution of π conditional on e_H first-order stochastically dominates the distribution conditional on e_L; that is, the distribution functions $F(\pi \mid e_L)$ and $F(\pi \mid e_H)$ satisfy $F(\pi \mid e_H) \leq F(\pi \mid e_L)$ at all $\pi \in [\underline{\pi}, \bar{\pi}]$, with strict inequality on some open set $\Pi \subset [\underline{\pi}, \bar{\pi}]$ (see Section 6.D). This implies that the level of expected profits when the manager chooses e_H is larger than that from e_L: $\int \pi f(\pi \mid e_H)\, d\pi > \int \pi f(\pi \mid e_L)\, d\pi$.

The manager is an expected utility maximizer with a Bernoulli utility function $u(w, e)$ over his wage w and effort level e. This function satisfies $u_w(w, e) > 0$ and $u_{ww}(w, e) \leq 0$ at all (w, e) (subscripts here denote partial derivatives) and $u(w, e_H) < u(w, e_L)$ at all w; that is, the manager prefers more income to less, is weakly risk averse over income lotteries, and dislikes a high level of effort.[4] In what follows, we focus on a special case of this utility function that has attracted much of the

3. In fact, more general interpretations are possible. For example, e could include non-effort-related managerial decisions such as what kind of inputs are purchased or the strategies that are adopted for appealing to buyers. We stick to the effort interpretation largely because it helps with intuition.

4. Note that in the multidimensional-effort case, it need not be that e_H has higher effort in every dimension; the only important thing for our analysis is that it leads to higher profits and entails a larger managerial disutility than does e_L.

attention in the literature: $u(w, e) = v(w) - g(e)$.[5] For this case, our assumptions on $u(w, e)$ imply that $v'(w) > 0$, $v''(w) \leq 0$, and $g(e_H) > g(e_L)$.

The owner receives the project's profits less any wage payments made to the manager. We assume that the owner is risk neutral and therefore that his objective is to maximize his expected return. The idea behind this simplifying assumption is that the owner may hold a well-diversified portfolio that allows him to diversify away the risk from this project. (Exercise 14.B.2 asks you to consider the case of a risk-averse owner.)

The Optimal Contract when Effort is Observable

It is useful to begin our analysis by looking at the optimal contracting problem when effort is observable.

Suppose that the owner chooses a contract to offer the manager that the manager can then either accept or reject. A contract here specifies the manager's effort $e \in \{e_L, e_H\}$ and his wage payment as a function of observed profits $w(\pi)$. We assume that a competitive market for managers dictates that the owner must provide the manager with an expected utility level of at least $\bar{u}$ if he is to accept the owner's contract offer ($\bar{u}$ is the manager's *reservation utility level*). If the manager rejects the owner's contract offer, the owner receives a payoff of zero.

We assume throughout that the owner finds it worthwhile to make the manager an offer that he will accept. The optimal contract for the owner then solves the following problem (for notational simplicity, we suppress the lower and upper limits of integration $\underline{\pi}$ and $\bar{\pi}$):

$$\underset{e \in \{e_L, e_H\},\, w(\pi)}{\text{Max}} \int (\pi - w(\pi))\, f(\pi \mid e)\, d\pi \tag{14.B.1}$$

$$\text{s.t.} \int v(w(\pi))\, f(\pi \mid e)\, d\pi - g(e) \geq \bar{u}.$$

It is convenient to think of this problem in two stages. First, for each choice of e that might be specified in the contract, what is the best compensation scheme $w(\pi)$ to offer the manager? Second, what is the best choice of e?

Given that the contract specifies effort level e, choosing $w(\pi)$ to maximize $\int (\pi - w(\pi)) f(\pi \mid e)\, d\pi = (\int \pi f(\pi \mid e)\, d\pi) - (\int w(\pi) f(\pi \mid e)\, d\pi)$ is equivalent to minimizing the expected value of the owner's compensation costs, $\int w(\pi) f(\pi \mid e)\, d\pi$, so (14.B.1) tells us that the optimal compensation scheme in this case solves

$$\underset{w(\pi)}{\text{Min}} \int w(\pi)\, f(\pi \mid e)\, d\pi \tag{14.B.2}$$

$$\text{s.t.} \int v(w(\pi))\, f(\pi \mid e)\, d\pi - g(e) \geq \bar{u}.$$

The constraint in (14.B.2) always binds at a solution to this problem; otherwise, the owner could lower the manager's wages while still getting him to accept the contract. Letting γ denote the multiplier on this constraint, at a solution to problem (14.B.2) the manager's wage $w(\pi)$ at each level of $\pi \in [\underline{\pi}, \bar{\pi}]$ must satisfy the first-order

5. Exercise 14.B.1 considers one implication of relaxing this assumption.

condition[6]

$$-f(\pi \mid e) + \gamma v'(w(\pi)) f(\pi \mid e) = 0,$$

or

$$\frac{1}{v'(w(\pi))} = \gamma. \tag{14.B.3}$$

If the manager is strictly risk averse [so that $v'(w)$ is strictly decreasing in w], the implication of condition (14.B.3) is that the optimal compensation scheme $w(\pi)$ is a constant; that is, the owner should provide the manager with a fixed wage payment. This finding is just a risk-sharing result: Given that the contract explicitly dictates the manager's effort choice and that there is no problem with providing incentives, the risk-neutral owner should fully insure the risk-averse manager against any risk in his income stream (in a manner similar to that in Example 6.C.1). Hence, given the contract's specification of e, the owner offers a fixed wage payment w_e^* such that the manager receives exactly his reservation utility level:

$$v(w_e^*) - g(e) = \bar{u}. \tag{14.B.4}$$

Note that since $g(e_H) > g(e_L)$, the manager's wage will be higher if the contract calls for effort e_H than if it calls for e_L.

On the other hand, when the manager is risk neutral, say with $v(w) = w$, condition (14.B.3) is necessarily satisfied for *any* compensation function. In this case, because there is no need for insurance, a fixed wage scheme is merely one of many possible optimal compensation schemes. Any compensation function $w(\pi)$ that gives the manager an expected wage payment equal to $\bar{u} + g(e)$ [the level derived from condition (14.B.4) when $v(w) = w$] is also optimal.

Now consider the optimal choice of e. The owner optimally specifies the effort level $e \in \{e_L, e_H\}$ that maximizes his expected profits less wage payments,

$$\int \pi f(\pi \mid e)\, d\pi - v^{-1}(\bar{u} + g(e)). \tag{14.B.5}$$

The first term in (14.B.5) represents the gross profit when the manager puts forth effort e; the second term represents the wages that must be paid to compensate the manager for this effort [derived from condition (14.B.4)]. Whether e_H or e_L is optimal depends on the incremental increase in expected profits from e_H over e_L compared with the monetary cost of the incremental disutility it causes the manager.

This is summarized in Proposition 14.B.1.

Proposition 14.B.1: In the principal-agent model with observable managerial effort, an optimal contract specifies that the manager choose the effort e^* that maximizes $[\int \pi f(\pi \mid e)\, d\pi - v^{-1}(\bar{u} + g(e))]$ and pays the manager a fixed wage $w^* = v^{-1}(\bar{u} + g(e^*))$. This is the uniquely optimal contract if $v''(w) < 0$ at all w.

6. The first-order condition for $w(\pi)$ is derived by taking the derivative with respect to the manager's wage at each level of π separately. To see this point, consider a discrete version of the model in which there is a finite number of possible profit levels $(\pi_1, \ldots, \pi_N)$ and associated wage levels $(w_1, \ldots, w_N)$. The first-order condition (14.B.3) is analogous to the condition one gets in the discrete model by examining the first-order conditions for each w_n, $n = 1, \ldots, N$ (note that we allow the wage payment to be negative). To be rigorous, we should add that when we have a continuum of possible levels of π, an optimal compensation scheme need only satisfy condition (14.B.3) at a set of profit levels that is of full measure.

The Optimal Contract when Effort is Not Observable

The optimal contract described in Proposition 14.B.1 accomplishes two goals: it specifies an efficient effort choice by the manager, and it fully insures him against income risk. When effort is not observable, however, these two goals often come into conflict because the only way to get the manager to work hard is to relate his pay to the realization of profits, which is random. When these goals come into conflict, the nonobservability of effort leads to inefficiencies.

To highlight this point, we first study the case in which the manager is risk neutral. We show that in this case, where the risk-bearing concern is absent, the owner can still achieve the same outcome as when effort is observable. We then study the optimal contract when the manager is risk averse. In this case, whenever the first-best (full observability) contract would involve the high-effort level, efficient risk bearing and efficient incentive provision come into conflict, and the presence of nonobservable actions leads to a welfare loss.

A risk-neutral manager

Suppose that $v(w) = w$. Applying Proposition 14.B.1, the optimal effort level e^* when effort is observable solves

$$\underset{e \in \{e_L, e_H\}}{\text{Max}} \quad \int \pi f(\pi \mid e)\, d\pi - g(e) - \bar{u}. \tag{14.B.6}$$

The owner's profit in this case is the value of expression (14.B.6), and the manager receives an expected utility of exactly $\bar{u}$.

Now consider the owner's payoff when the manager's effort is not observable. In Proposition 14.B.2, we establish that the owner can still achieve his full-information payoff.

Proposition 14.B.2: In the principal-agent model with unobservable managerial effort and a risk-neutral manager, an optimal contract generates the same effort choice and expected utilities for the manager and the owner as when effort is observable.

Proof: We show explicitly that there is a contract the owner can offer that gives him the same payoff that he receives under full information. This contract must therefore be an optimal contract for the owner because the owner can never do better when effort is not observable than when it is (when effort is observable, the owner is always free to offer the optimal nonobservability contract and simply leave the choice of an effort level up to the manager).

Suppose that the owner offers a compensation schedule of the form $w(\pi) = \pi - \alpha$, where α is some constant. This compensation schedule can be interpreted as "selling the project to the manager" because it gives the manager the full return π except for the fixed payment α (the "sales price"). If the manager accepts this contract, he chooses e to maximize his expected utility,

$$\int w(\pi) f(\pi \mid e)\, d\pi - g(e) = \int \pi f(\pi \mid e)\, d\pi - \alpha - g(e). \tag{14.B.7}$$

Comparing (14.B.7) with (14.B.6), we see that e^* maximizes (14.B.7). Thus, this contract induces the first-best (full observability) effort level e^*.

The manager is willing to accept this contract as long as it gives him an expected utility of at least $\bar{u}$, that is, as long as

$$\int \pi f(\pi \mid e^*)\, d\pi - \alpha - g(e^*) \geq \bar{u}. \tag{14.B.8}$$

Let α^* be the level of α at which (14.B.8) holds with equality. Note that the owner's payoff if the compensation scheme is $w(\pi) = \pi - \alpha^*$ is exactly α^* (the manager gets all of π except for the fixed payment α^*). Rearranging (14.B.8), we see that $\alpha^* = \int \pi f(\pi \mid e^*)\, d\pi - g(e^*) - \bar{u}$. Hence, with compensation scheme $w(\pi) = \pi - \alpha^*$, both the owner and the manager get exactly the same payoff as when effort is observable. ■

The basic idea behind Proposition 14.B.2 is straightforward. If the manager is risk neutral, the problem of risk sharing disappears. Efficient incentives can be provided without incurring any risk-bearing losses by having the manager receive the full marginal returns from his effort.

A risk-averse manager

When the manager is strictly risk averse over income lotteries, matters become more complicated. Now incentives for high effort can be provided only at the cost of having the manager face risk. To characterize the optimal contract in these circumstances, we again consider the contract design problem in two steps: first, we characterize the optimal incentive scheme for each effort level that the owner might want the manager to select; second, we consider which effort level the owner should induce.

The optimal incentive scheme for implementing a specific effort level e minimizes the owner's expected wage payment subject to two constraints. As before, the manager must receive an expected utility of at least $\bar{u}$ if he is to accept the contract. When the manager's effort is unobservable, however, the owner also faces a second constraint: The manager must actually *desire* to choose effort e when facing the incentive scheme. Formally, the optimal incentive scheme for implementing e must therefore solve

$$\underset{w(\pi)}{\text{Min}} \int w(\pi) f(\pi \mid e)\, d\pi \tag{14.B.9}$$

$$\text{s.t. (i)} \int v(w(\pi)) f(\pi \mid e)\, d\pi - g(e) \geq \bar{u}$$

$$\text{(ii) } e \text{ solves } \underset{\tilde{e}}{\text{Max}} \int v(w(\pi)) f(\pi \mid \tilde{e})\, d\pi - g(\tilde{e}).$$

Constraint (ii) is known as the *incentive constraint*: it insures that under compensation scheme $w(\pi)$ the manager's optimal effort choice is e.

How does the owner optimally implement each of the two possible levels of e? We consider each in turn.

Implementing e_L: Suppose, first, that the owner wishes to implement effort level e_L. In this case, the owner optimally offers the manager the fixed wage payment $w_e^* = v^{-1}(\bar{u} + g(e_L))$, the same payment he would offer if contractually specifying effort e_L when effort is observable. To see this, note that with this compensation scheme the manager selects e_L: His wage payment is unaffected by his effort, and so he will choose the effort level that involves the lowest disutility, namely e_L. Doing so, he earns exactly $\bar{u}$. Hence, this contract implements e_L at exactly the same cost as when effort is observable. But, as we noted in the proof of Proposition 14.B.2, the owner can never do better when effort is unobservable than when effort is observable [formally, in problem (14.B.9), the owner faces the additional constraint (ii) relative to problem (14.B.2)]; therefore, this must be a solution to problem (14.B.9).

Implementing e_H: The more interesting case arises when the owner decides to induce effort level e_H. In this case, constraint (ii) of (14.B.9) can be written as

$$(\text{ii}_H)\quad \int v(w(\pi))\, f(\pi \mid e_H)\, d\pi - g(e_H) \geq \int v(w(\pi))\, f(\pi \mid e_L)\, d\pi - g(e_L).$$

Letting $\gamma \geq 0$ and $\mu \geq 0$ denote the multipliers on constraints (i) and (ii$_H$), respectively, $w(\pi)$ must satisfy the following Kuhn–Tucker first-order condition at every $\pi \in [\underline{\pi}, \bar{\pi}]$:[7]

$$-f(\pi \mid e_H) + \gamma v'(w(\pi))\, f(\pi \mid e_H) + \mu[f(\pi \mid e_H) - f(\pi \mid e_L)]v'(w(\pi)) = 0$$

or

$$\frac{1}{v'(w(\pi))} = \gamma + \mu\left[1 - \frac{f(\pi \mid e_L)}{f(\pi \mid e_H)}\right]. \qquad (14.B.10)$$

We first establish that in any solution to problem (14.B.9), where $e = e_H$, both γ and μ are strictly positive.

Lemma 14.B.1: In any solution to problem (14.B.9) with $e = e_H$, both $\gamma > 0$ and $\mu > 0$.

Proof: Suppose that $\gamma = 0$. Because $F(\pi \mid e_H)$ first-order stochastically dominates $F(\pi \mid e_L)$, there must exist an open set of profit levels $\tilde{\Pi} \subset [\underline{\pi}, \bar{\pi}]$ such that $[f(\pi \mid e_L)/f(\pi \mid e_H)] > 1$ at all $\pi \in \tilde{\Pi}$. But if $\gamma = 0$, condition (14.B.10) then implies that $v'(w(\pi)) \leq 0$ at any such π (recall that $\mu \geq 0$), which is impossible. Hence, $\gamma > 0$.

On the other hand, if $\mu = 0$ in the solution to problem (14.B.9) then, by condition (14.B.10), the optimal compensation schedule gives a fixed wage payment for every profit realization. But we know that this would lead the manager to choose e_L rather than e_H, violating constraint (ii$_H$) of problem (14.B.9). Hence, $\mu > 0$. ■

7. Although problem (14.B.9) may not appear to be a convex programming problem, a simple transformation of the problem shows that (14.B.10) is both a necessary and a sufficient condition for a solution. To see this, reformulate (14.B.9) as a problem of choosing the manager's level of utility for each profit outcome π, say $\bar{v}(\pi)$. Letting $\phi(\cdot) = v^{-1}(\cdot)$, the objective function becomes $\int \phi(\bar{v}(\pi)) f(\pi \mid e_H)\, d\pi$, which is convex in $\bar{v}(\pi)$, and the constraints are then all linear in $\bar{v}(\pi)$. Thus, (Kuhn–Tucker) first-order conditions are both necessary and sufficient for a maximum of this reformulated problem (see Section M.K of the Mathematical Appendix). The first-order condition for this problem is

$$-\phi'(\bar{v}(\pi)) f(\pi \mid e_H) + \gamma f(\pi \mid e_H) + \mu[f(\pi \mid e_H) - f(\pi \mid e_L)] = 0 \quad \text{for all } \pi \in [\underline{\pi}, \bar{\pi}].$$

Defining $w(\pi)$ by $v(w(\pi)) = \bar{v}(\pi)$, and noting that $\phi'(v(w(\pi))) = 1/v'(w(\pi))$, this gives (14.B.10).

Lemma 14.B.1 tells us that both constraints in problem (14.B.9) bind when $e = e_H$.[8] Moreover, given Lemma 14.B.1, condition (14.B.10) can be used to derive some useful insights into the shape of the optimal compensation schedule. Consider, for example, the fixed wage payment $\hat{w}$ such that $(1/v'(\hat{w})) = \gamma$. According to condition (14.B.10),

$$w(\pi) > \hat{w} \quad \text{if} \quad \frac{f(\pi \mid e_L)}{f(\pi \mid e_H)} < 1$$

and

$$w(\pi) < \hat{w} \quad \text{if} \quad \frac{f(\pi \mid e_L)}{f(\pi \mid e_H)} > 1.$$

This relationship is fairly intuitive. The optimal compensation scheme pays more than $\hat{w}$ for outcomes that are statistically relatively more likely to occur under e_H than under e_L in the sense of having a likelihood ratio $[f(\pi \mid e_L)/f(\pi \mid e_H)]$ less than 1. Similarly, it offers less compensation for outcomes that are relatively more likely when e_L is chosen. We should stress, however, that while this condition evokes a statistical interpretation, there is no actual statistical inference going on here; the owner *knows* what level of effort will be chosen given the compensation schedule he offers. Rather, the compensation package has this form because of its *incentive effects.* That is, by structuring compensation in this way, it provides the manager with an incentive for choosing e_H instead of e_L.

This point leads to what may at first seem a somewhat surprising implication: in an optimal incentive scheme, compensation is not necessarily monotonically increasing in profits. As is clear from examination of condition (14.B.10), for the optimal compensation scheme to be monotonically increasing, it must be that the likelihood ratio $[f(\pi \mid e_L)/f(\pi \mid e_H)]$ is decreasing in π; that is, as π increases, the likelihood of getting profit level π if effort is e_H relative to the likelihood if effort is e_L must increase. This property, known as the *monotone likelihood ratio property* [see Milgrom (1981)], is *not* implied by first-order stochastic dominance. Figures 14.B.1(a) and (b), for example, depict a case in which the distribution of π conditional on e_H stochastically dominates the distribution of π conditional on e_L but the monotone likelihood ratio property does not hold. In this example, increases in effort serve to convert low profit realizations into intermediate ones but have no effect on the likelihood of very high profit realizations. Condition (14.B.10) tells us that in this case, we should have higher wages at intermediate levels of profit than at very high ones because it is the likelihood of intermediate profit levels that is sensitive to increases in effort. The optimal compensation function for this example is shown in Figure 14.B.1(c).

8. A more direct argument for constraint (i) being binding goes as follows: Suppose that $w(\pi)$ is a solution to (14.B.9) in which constraint (i) is not binding. Consider a change in the compensation function that lowers the wage paid at each level of π in such a way that the resulting decrease in utility is equal at all π, that is, to a new function $\hat{w}(\pi)$ with $[v(w(\pi)) - v(\hat{w}(\pi))] = \Delta v > 0$ at all $\pi \in [\underline{\pi}, \bar{\pi}]$. This change does not affect the satisfaction of the incentive constraint (ii_H) since if the manager was willing to pick e_H when faced with $w(\pi)$, he will do so when faced with $\hat{w}(\pi)$. Furthermore, because constraint (i) is not binding, the manager will still accept this new contract if Δv is small enough. Lastly, the owner's expected wage payments will be lower than under $w(\pi)$. This yields a contradiction.

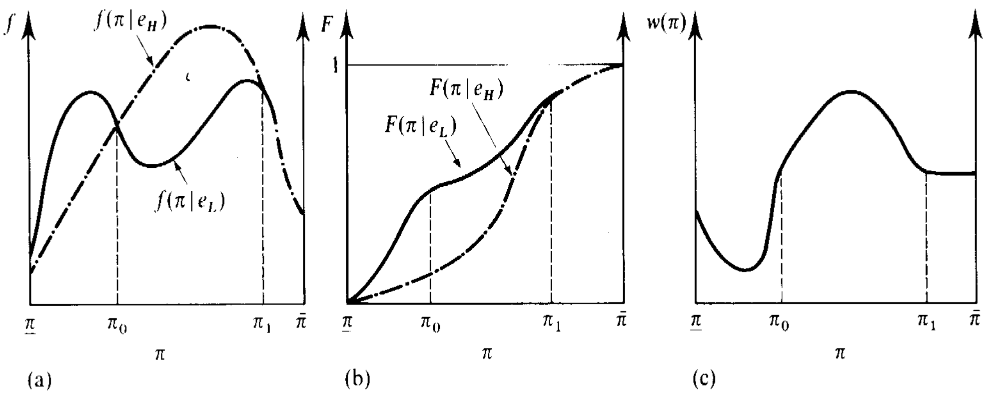

Figure 14.B.1
A violation of the monotone likelihood ratio property.
(a) Densities.
(b) Distribution functions. (c) Optimal wage scheme.

Condition (14.B.10) also implies that the optimal contract is not likely to take a simple (e.g., linear) form. The optimal shape of $w(\pi)$ is a function of the informational content of various profit levels (through the likelihood ratio), and this is unlikely to vary with π in a simple manner in most problems.

Finally, note that given the variability that is optimally introduced into the manager's compensation, the expected value of the manager's wage payment must be strictly greater than his (fixed) wage payment in the observable case, $w^*_{e_H} = v^{-1}(\bar{u} + g(e_H))$. Intuitively, because the manager must be assured an expected utility level of $\bar{u}$, the owner must compensate him through a higher average wage payment for any risk he bears. To see this point formally, note that since $E[v(w(\pi))\,|\,e_H] = \bar{u} + g(e_H)$ and $v''(\cdot) < 0$, Jensen's inequality (see Section M.C of the Mathematical Appendix) tells us that $v(E[w(\pi)\,|\,e_H]) > \bar{u} + g(e_H)$. But we know that $v(w^*_{e_H}) = \bar{u} + g(e_H)$, and so $E[w(\pi)\,|\,e_H] > w^*_{e_H}$. As a result, nonobservability increases the owner's expected compensation costs of implementing effort level e_H.

Given the preceding analysis, which effort level should the owner induce? As before, the owner compares the incremental change in expected profits from the two effort levels $[\int \pi f(\pi\,|\,e_H)\,d\pi - \int \pi f(\pi\,|\,e_L)\,d\pi]$ with the difference in expected wage payments in the contracts that optimally implement each of them, that is, with the difference in the value of problem (14.B.9) for $e = e_H$ compared with $e = e_L$.

From the preceding analysis, we know that the wage payment when implementing e_L is exactly the same as when effort is observable, whereas the expected wage payment when the owner implements e_H under nonobservability is strictly larger than his payment in the observable case. Thus, in this model, nonobservability raises the cost of implementing e_H and does not change the cost of implementing e_L. The implication of this fact is that nonobservability of effort can lead to an inefficiently low level of effort being implemented. When e_L would be the optimal effort level if effort were observable, then it still is when effort is nonobservable. In this case, nonobservability causes no losses. In contrast, when e_H would be optimal if effort were observable, then one of two things may happen: it may be optimal to implement e_H using an incentive scheme that faces the manager with risk; alternatively, the risk-bearing costs may be high enough that the owner decides that it is better to simply implement e_L. In either case, nonobservability causes a welfare loss to the owner (the manager's expected utility is $\bar{u}$ in either case).[9]

These observations are summarized in Proposition 14.B.3.

Proposition 14.B.3: In the principal-agent model with unobservable manager effort, a risk-averse manager, and two possible effort choices, the optimal compensation scheme for implementing e_H satisfies condition (14.B.10), gives the manager expected utility $\bar{u}$, and involves a larger expected wage payment than is required when effort is observable. The optimal compensation scheme for implementing e_L involves the same fixed wage payment as if effort were observable. Whenever the optimal effort level with observable effort would be e_H, nonobservability causes a welfare loss.

The fact that nonobservability leads in this model only to *downward* distortions in the manager's effort level is a special feature of the two-effort-level specification. With many possible effort choices, nonobservability may still alter the level of managerial effort induced in an optimal contract from its level under full observability, but the direction of the bias can be upward as well as downward. (See Exercise 14.B.4 for an illustration.)

Imagine that another statistical signal of effort, say y, is available to the owner in addition to the realization of profits, and that the joint density of π and y given e is given by $f(\pi, y \mid e)$. In this case, the manager's compensation can, in principle, be made to depend on both π and y. When should compensation be made a function of this variable as well? That is, when does the optimal compensation function $w(\pi, y)$ actually depend on y?

To answer this question, suppose that the owner wishes to implement e_H. Following along the same lines as above, we can derive a condition analogous to condition (14.B.10):

$$\frac{1}{v'(w(\pi, y))} = \gamma + \mu\left[1 - \frac{f(\pi, y \mid e_L)}{f(\pi, y \mid e_H)}\right]. \tag{14.B.11}$$

Consider, first, the case in which y is simply a noisy random variable that is unrelated to e. Then we can write the density $f(\pi, y \mid e)$ as the product of two densities, $f_1(\pi \mid e)$ and $f_2(y)$: $f(\pi, y \mid e) = f_1(\pi \mid e) f_2(y)$. Substituting into (14.B.11), the $f_2(\cdot)$ terms cancel out, and so the optimal compensation package is independent of y.

The intuition behind this result is straightforward. Suppose that the owner is initially offering a contract that has wage payments dependent on y. Intuitively, this contract induces a randomness in the manager's wage that is unrelated to e and therefore makes the manager face risk without achieving any beneficial incentive effect. If the owner instead offers, for each realization of π, the certain payment $\bar{w}(\pi)$ such that

$$v(\bar{w}(\pi)) = E[v(w(\pi, y)) \mid \pi] = \int v(w(\pi, y)) f_2(y)\, dy,$$

9. Note, however, that although nonobservability leads to a welfare loss, the outcome here is a constrained Pareto optimum in the sense introduced in Section 13.B. To see this, note that the owner maximizes his profit subject to giving the manager an expected utility level no less than $\bar{u}$ and subject to constraints deriving from his inability to observe the manager's effort choice. As a result, no allocation that Pareto dominates this outcome can be achieved by a central authority who cannot observe the manager's effort choice. For market intervention by such an authority to generate a Pareto improvement, there must be externalities among the contracts signed by different pairs of individuals.

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, $v(E[w(\pi, y)\,|\,\pi]) > E[v(w(\pi, y))\,|\,\pi]$, and so $\bar{w}(\pi) < E[w(\pi, y)\,|\,\pi]$).

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

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

If $f_2(y\,|\,\pi, e)$ 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 $f_2(y\,|\,\pi, e)$ 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

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

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

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

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

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:

$$\underset{\substack{w_L, e_L \geq 0 \\ w_H, e_H \geq 0}}{\text{Max}} \quad \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L] \tag{14.C.1}$$

$$\text{s.t.} \quad \lambda\, v(w_H - g(e_H, \theta_H)) + (1 - \lambda) v(w_L - g(e_L, \theta_L)) \geq \bar{u}.$$

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

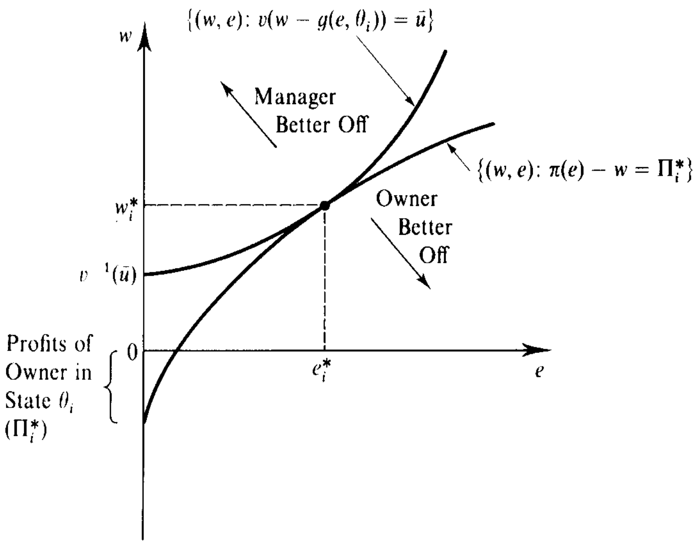

Figure 14.C.1
The optimal wage–effort pair for state θ_i when states are observable.

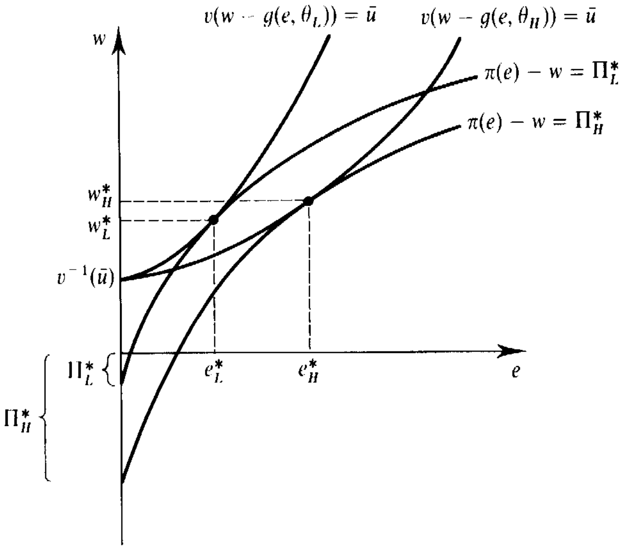

Figure 14.C.2
The optimal contract with full observability of θ.

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:

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

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

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

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 $\rho \to -\infty$, $[\lambda v_H^\rho + (1-\lambda)v_L^\rho]^{1/\rho} \to \text{Min}\,\{v_H, v_L\}$ (see Exercise 3.C.6). Hence, a contract gives the manager an expected utility greater than his (certain) reservation utility if and only if $\text{Min}\,\{v(w_H - g(e_H, \theta_H)), v(w_L - g(e_L, \theta_L))\} \geq \bar{u}$.

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

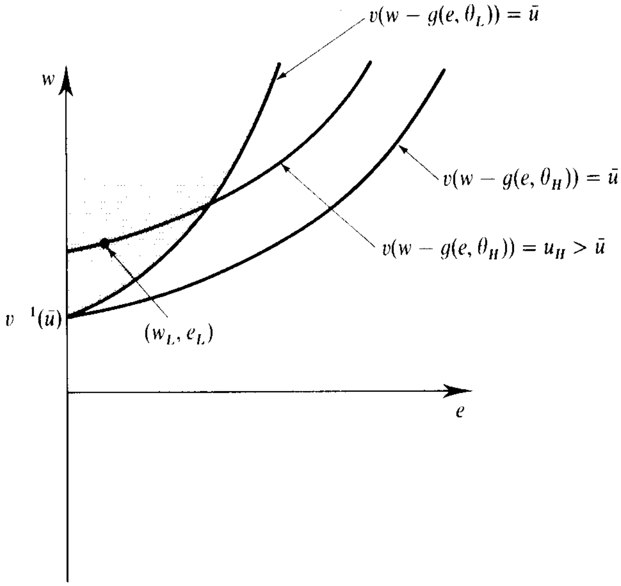

Figure 14.C.3
Constraint (ii) in problem (14.C.8) is satisfied by any contract satisfying constraints (i) and (iii).

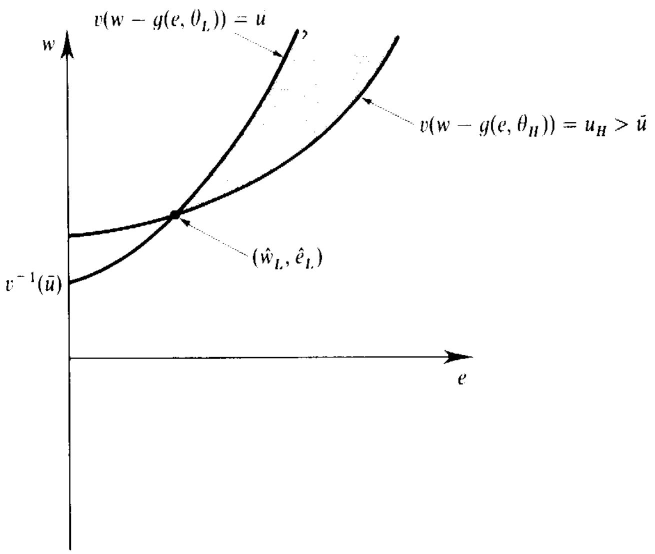

Figure 14.C.4 (left)
In a feasible contract offering $(\hat{w}_L, \hat{e}_L)$ for state θ_L, the pair (w_H, e_H) must lie in the shaded region.

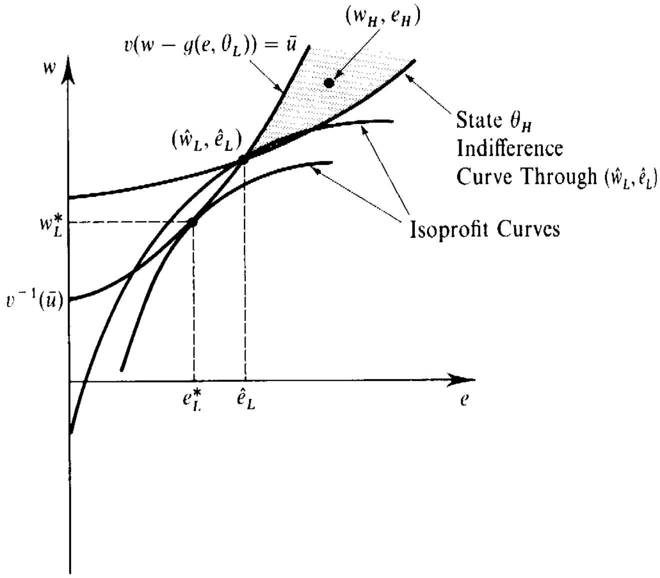

Figure 14.C.5 (right)
An optimal contract has $e_L \leq e_L^*$.

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. ■

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

(i) $e_L \leq e_L^*$; 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 $(\hat{w}_L, \hat{e}_L)$ 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,

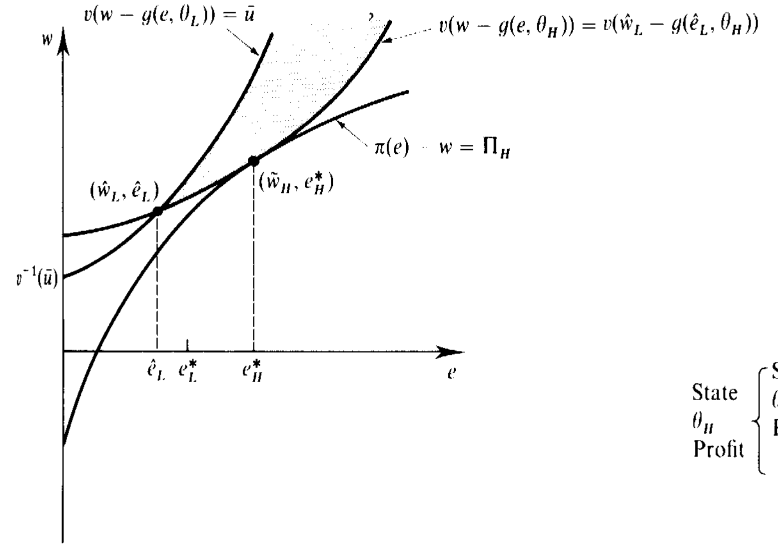

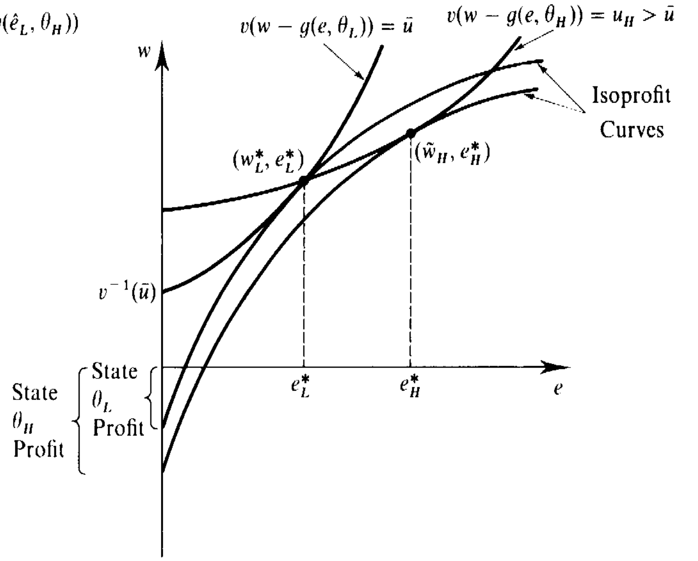

Figure 14.C.6 (left)
An optimal contract has $e_H = e_H^*$.

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 $\hat{e}_L \leq e_L^*$, 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 $\hat{e}_L \leq e_L^* < e_H^*$. ■

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.)

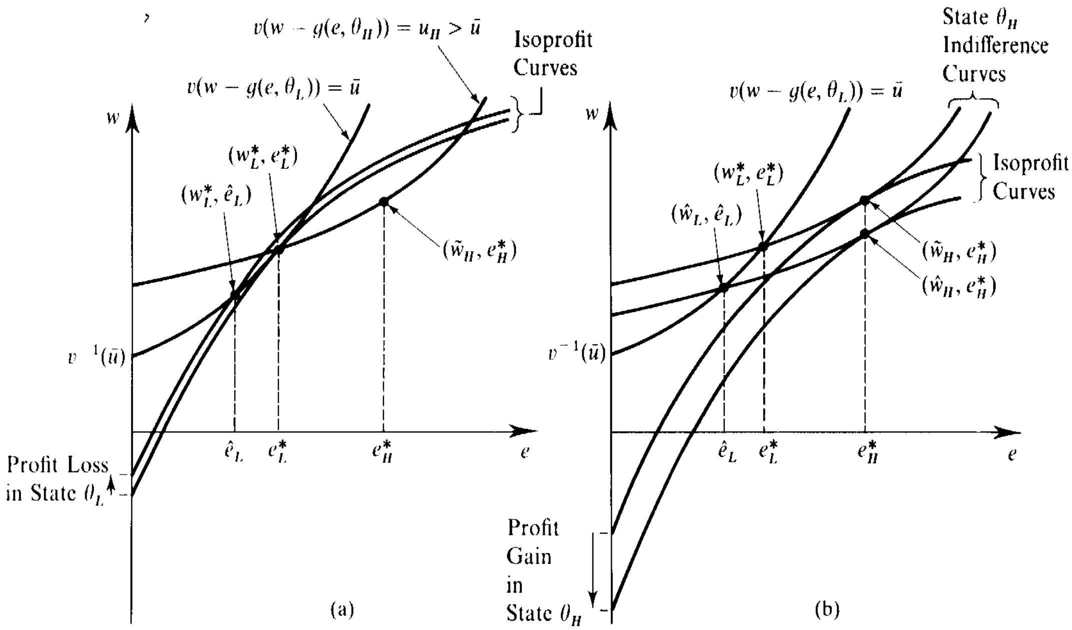

Figure 14.C.8 (a) The change in profits in state θ_L from lowering e_L slightly below e_L^*. (b) The change in profits in state θ_H from lowering e_L slightly below e_L^* and optimally adjusting w_H.

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:

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

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 $t_i = p_i$, $w_i = T_i$, $u = 0$, $g(t_i, \theta_i) = -(p_i - \theta_i)x(p_i)$, 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.

model considered in Section 14.B; the "truthfulness" constraints are generalizations of those considered in our hidden information model. See Myerson (1982) for details.

One special case of this hybrid model deserves particular mention because its analysis reduces to that of the pure hidden information model considered in Section 14.C. In particular, suppose that effort is unobservable but that the relationship between effort and profits is *deterministic*, given by the function $\pi(e)$. In that case, for any particular announcement $\hat{\theta}$, it is possible to induce any wage–effort pair that is desired, say $(\hat{w}, \hat{e})$, by use of a simple "forcing" compensation scheme: Just reward the manager with a wage payment of $\hat{w}$ if profits are $\pi(\hat{e})$, and give him a wage payment of $-\infty$ otherwise. Thus, the combination of the observability of π and the one-to-one relationship between π and e effectively allows the contract to specify e. The analysis of this model is therefore identical to that of the hidden information model considered in Section 14.C, where wage–effort pairs could be specified directly as functions of the manager's announcement.

To see this point in a slightly different way, note first that because of the ability to write forcing contracts, in this model an optimal contract can be thought of as specifying, for each announcement $\hat{\theta}$, a wage–profit pair $(w(\hat{\theta}), \pi(\hat{\theta}))$. Now, for any required profit level π, the effort level necessary to achieve a profit of π is $\tilde{e}$ such that $\pi(\tilde{e}) = \pi$. Let the function $\tilde{e}(\pi)$ describe this effort level. We can now think of the manager as having a disutility function defined directly over the profit level which is given by $\tilde{g}(\pi, \theta) = g(\tilde{e}(\pi), \theta)$. But this model looks just like a model with *observable* effort where the effort variable is π, the disutility function over this effort is $\tilde{g}(\pi, \theta)$, and the profit function is $\tilde{\pi}(\pi) = \pi$. Thus, the analysis of this model is identical to that in a pure hidden information model.

A similar point applies to a closely related hybrid model in which, instead of the manager's disutility of effort, it is the relation between profit and effort that depends on the state. In particular, suppose that the disutility of effort is given by the function $g(e)$ and profits are given by the function $\pi(e, \theta)$, where $\pi_e(\cdot) > 0$, $\pi_{ee}(\cdot) < 0$, $\pi_\theta(\cdot) > 0$, and $\pi_{e\theta}(\cdot) > 0$. Effort is not observable, but profits are. The idea is that the manager knows more than the owner does about the true profit opportunities facing the firm (e.g., the marginal productivity of effort). Again, we can think of a contract as specifying, for each announcement by the manager, a wage–profit pair (implicitly using forcing contracts). In this context, the effort needed to achieve any given level of profit π in state θ is given by some function $\hat{e}(\pi, \theta)$, and the disutility associated with this effort is then $\hat{g}(\pi, \theta) = g(\hat{e}(\pi, \theta))$. But this model is also equivalent to our basic hidden information model with observable effort: just let the effort variable be π, the disutility of this effort be $\hat{g}(\pi, \theta)$, and the profit function be $\hat{\pi}(\pi) = \pi$. Again, our results from Section 14.C apply.

APPENDIX A: MULTIPLE EFFORT LEVELS IN THE HIDDEN ACTION MODEL

In this appendix, we discuss additional issues that arise when the effort choice in the hidden action (moral hazard) model discussed in Section 14.B is more complex than the simple two-effort-choice specification $e \in \{e_L, e_H\}$ analyzed there. Here, we return

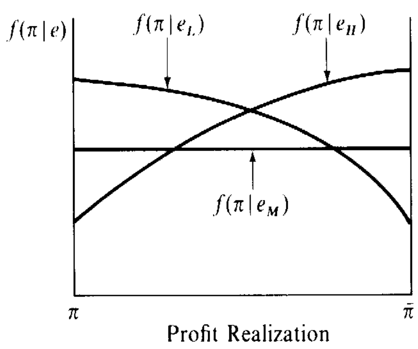

Figure 14.AA.1
Density functions for $E = \{e_L, e_M, e_H\}$: effort choice e_M may not be implementable.

to the more general specification initially introduced in Section 14.B in which E is the feasible set of effort choices.

As in Section 14.B, we can break up the principal's (the owner's) problem into several parts:

(a) What are the effort levels e that it is possible to induce?
(b) What is the optimal contract for inducing each specific effort level $e \in E$?
(c) Which effort level $e \in E$ is optimal?

In a multiple-action setting, each of these three parts becomes somewhat more complicated. For example, with just two actions, part (a) was trivial: e_L could be induced with a fixed wage contract, and e_H could always be induced by giving incentives that were sufficiently high at outcomes that were more likely to arise when e_H is chosen. With more than two actions, however, this may not be so. For example, consider the three-action case in which $E = \{e_L, e_M, e_H\}$ and the conditional density functions are those depicted in Figure 14.AA.1. As is suggested by the figure, it may be impossible to design incentives such that e_M is chosen because for any $w(\pi)$ the agent may prefer either e_L or e_H to e_M. (Exercise 14.B.4 provides an example along these lines.)

Part (b) also becomes more involved. The optimal contract for implementing effort choice e solves

$$\operatorname*{Min}_{w(\pi)} \int w(\pi)\, f(\pi \mid e)\, d\pi \qquad (14.AA.1)$$

$$\text{s.t. (i)} \int v(w(\pi))\, f(\pi \mid e)\, d\pi - g(e) \geq \bar{u}$$

$$\text{(ii) } e \text{ solves } \operatorname*{Max}_{\tilde{e} \in E} \int v(w(\pi))\, f(\pi \mid \tilde{e})\, d\pi - g(\tilde{e}).$$

If we have K possible actions in set E, the incentive constraints in problem (14.AA.1) [constraints (ii)] consist of $(K - 1)$ constraints that must be satisfied. In this case, with a change of variables in which we maximize over the level of utility that the manager gets conditional on π, say $\bar{v}(\pi)$, we have a problem with K linear constraints and a convex objective function [see Grossman and Hart (1983) and footnote 7 for more on this].

However, if E is a continuous set of possible actions, say $E = [0, \bar{e}] \subset \mathbb{R}$, then we have an *infinity* of incentive constraints. One trick sometimes used in this case to simplify problem (14.AA.1) is to replace constraint (ii) with a *first-order condition* (this is sometimes called the *first-order approach*). For example, if e is a one-dimensional measure of effort, then the manager's first-order condition is

$$\int v(w(\pi))\, f_e(\pi \mid e)\, d\pi - g'(e) = 0, \tag{14.AA.2}$$

where $f_e(\pi \mid e) = \partial f(\pi \mid e)/\partial e$. If we replace constraint (ii) with (14.AA.2) and solve the resulting problem, we can derive a condition for $w(\pi)$ that parallels condition (14.B.10):

$$\frac{1}{v'(w(\pi))} = \lambda + \mu\left[\frac{f_e(\pi \mid e)}{f(\pi \mid e)}\right]. \tag{14.AA.3}$$

The condition that ratio $[f_e(\pi \mid e)/f(\pi \mid e)]$ be increasing in π is the differential version of the monotone likelihood ratio property (see Exercise 14.AA.1).

In general, however, a solution to the problem resulting from this substitution is not necessarily a solution to the actual problem (14.AA.1). The reason is that the agent may satisfy first-order condition (14.AA.2) even when effort level e is not his optimal effort choice. First, effort level e could be a *minimum* rather than a maximum; therefore, we at least want the agent to also be satisfying a local second-order condition. But even this will not be sufficient. In general, we need to be sure that the agent's objective function is concave in e. Note that this is not a simple matter because the concavity of his objective function in e will depend both on the shape of $f(\pi \mid e)$ *and* on the shape of the incentive contract $w(\pi)$ that is offered. The known conditions which insure that this condition is met are very restrictive. See Grossman and Hart (1983) and Rogerson (1985b) for details. Exercise 14.AA.2 provides a very simple example.

Finally, to answer part (c), we need to compute the optimal contract from part (b) for each action that part (a) reveals is implementable and then compare their relative profits for the principal. With more than two effort choices, two features of the two-effort-choice case fail to generalize. First, nonobservability can lead to an upward distortion in effort. (Exercise 14.B.4 provides an example.) Second, at the optimal contract under nonobservability we can get *both* an inefficient effort choice and inefficiencies resulting from managerial risk bearing.

APPENDIX B: A FORMAL SOLUTION OF THE PRINCIPAL-AGENT PROBLEM WITH HIDDEN INFORMATION

Recall problem (14.C.8):

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

Using Lemma 14.C.1 we can restate problem (14.C.8) as

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

$$\text{s.t.} \quad \text{(i) } w_L - g(e_L, \theta_L) \geq v^{-1}(\bar{u})$$

$$\text{(iii) } w_H - g(e_H, \theta_H) \geq w_L - g(e_L, \theta_H)$$

$$\text{(iv) } w_L - g(e_L, \theta_L) \geq w_H - g(e_H, \theta_L).$$

Letting $(\gamma, \phi_H, \phi_L) \geq 0$ be the multipliers on constraints (i), (iii), and (iv), respectively, the Kuhn–Tucker conditions for this problem can be written (see Section M.K of the Mathematical Appendix)

$$-\lambda + \phi_H - \phi_L = 0. \qquad (14.\text{BB}.2)$$

$$-(1-\lambda) + \gamma - \phi_H + \phi_L = 0. \qquad (14.\text{BB}.3)$$

$$\lambda\pi'(e_H) - \phi_H g_e(e_H, \theta_H) + \phi_L g_e(e_H, \theta_L) \begin{cases} \leq 0 & \\ = 0 & \text{if } e_H > 0 \end{cases} \qquad (14.\text{BB}.4)$$

$$(1-\lambda)\pi'(e_L) - (\gamma + \phi_L) g_e(e_L, \theta_L) + \phi_H g_e(e_L, \theta_H) \begin{cases} \leq 0 & \\ = 0 & \text{if } e_L > 0 \end{cases} \qquad (14.\text{BB}.5)$$

along with the complementary slackness conditions for constraints (i), (iii), and (iv) [conditions (M.K.7)].

Let us break up the analysis of these conditions into several steps.

Step 1: Condition (14.BB.2) implies that $\phi_H > 0$. Thus, constraint (iii) must bind (hold with equality) at an optimal solution.

Step 2: Adding conditions (14.BB.2) and (14.BB.3) implies that $\gamma = 1$. Hence, constraint (i) must bind at an optimal solution.

Step 3: Both e_L and e_H are strictly positive. To see this, note that condition (14.BB.4) cannot hold at $e_H = 0$ because $\pi'(0) > 0$ and $g_e(0, \theta_i) = 0$ for $i = L, H$. Similarly for condition (14.BB.5) and e_L.

Step 4: Steps 1 to 3 imply that $\phi_L = 0$. Suppose not: i.e., that $\phi_L > 0$. Then constraint (iv) must be binding. We shall now derive a contradiction. First, substitute for ϕ_H in conditions (14.BB.4) and (14.BB.5) using the fact that $\phi_H = \phi_L + \lambda$ from condition (14.BB.2). Then, using the fact that $(e_L, e_H) \gg 0$, we can write conditions (14.BB.4) and (14.BB.5) as

$$\lambda[\pi'(e_H) - g_e(e_H, \theta_H)] + \phi_L[g_e(e_H, \theta_L) - g_e(e_H, \theta_H)] = 0$$

and

$$(1-\lambda)[\pi'(e_L) - g_e(e_L, \theta_H)] + (1 + \phi_L)[g_e(e_L, \theta_H) - g_e(e_L, \theta_L)] = 0.$$

But $\phi_L > 0$ then implies that

$$\pi'(e_L) - g_e(e_L, \theta_H) > 0 > \pi'(e_H) - g_e(e_H, \theta_H),$$

which implies $e_H > e_L$ since $\pi(e) - g(e, \theta_H)$ is concave in e. But if $e_H > e_L$ and constraint (iii) binds (which it does from Step 1), then constraint (iv) must be slack because we then have

$$\begin{aligned}(w_H - w_L) &= g(e_H, \theta_H) - g(e_L, \theta_H) \\ &= \int_{e_L}^{e_H} g_e(e, \theta_H)\, de \\ &< \int_{e_L}^{e_H} g_e(e, \theta_L)\, de \\ &= g(e_H, \theta_L) - g(e_L, \theta_L).\end{aligned}$$

This is our desired contradiction.

Step 5: Since $\phi_L = 0$, we know from (14.BB.2) that $\phi_H = \lambda$. Substituting these two values into conditions (14.BB.4) and (14.BB.5) we have

$$\pi'(e_H) - g_e(e_H, \theta_H) = 0 \qquad (14.BB.6)$$

and

$$[\pi'(e_L) - g_e(e_L, \theta_L)] + \frac{\lambda}{1-\lambda}[g_e(e_L, \theta_H) - g_e(e_L, \theta_L)] = 0. \qquad (14.BB.7)$$

Conditions (14.BB.6) and (14.BB.7) characterize the optimal values of e_H and e_L, respectively. The optimal values for w_L and w_H are then determined from constraints (i) and (iii), which we have seen hold with equality at the solution.

An alternative approach to solving problem (14.BB.1) that avoids this somewhat cumbersome argument involves the following "trick": Solve problem (14.BB.1) ignoring constraint (iv). Then show that the solution derived in this way also satisfies constraint (iv). If so, this must be a solution to the (more constrained) problem (14.BB.1). (Exercise 14.BB.1 asks you to try this approach.)

REFERENCES

Allen, F. (1985). Repeated principal-agent relationships with lending and borrowing. *Economic Letters* **17**: 27–31.

Baron, D., and R. Myerson. (1982). Regulating a monopolist with unknown costs. *Econometrica* **50**: 911–30.

Bernheim, B. D., and M. D. Whinston. (1986). Common agency. *Econometrica* **54**: 923–42.

Dasgupta, P., P. Hammond, and E. Maskin. (1979). The implementation of social choice rules: Some results on incentive compatibility. *Review of Economic Studies* **46**: 185–216.

Dye, R. (1986). Optimal monitoring policies in agencies. *Rand Journal of Economics* **17**: 339–50.

Fudenberg, D., B. Holmstrom, and P. Milgrom. (1990). Short-term contracts and long-term agency relationships. *Journal of Economic Theory* **52**: 194–206.

Green, J., and N. Stokey. (1983). A comparison of tournaments and contests. *Journal of Political Economy* **91**: 349–64.

Grossman, S. J., and O. D. Hart. (1983). An analysis of the principal-agent problem. *Econometrica* **51**: 7–45.

Hart, O. D., and B. Holmstrom. (1987). The theory of contracts. In *Advances in Economic Theory, Fifth World Congress*, edited by T. Bewley. New York: Cambridge University Press.

Holmstrom, B. (1979). Moral hazard and observability. *Bell Journal of Economics* **10**: 74–91.

Holmstrom, B. (1982). Moral hazard in teams. *Bell Journal of Economics* **13**: 324–40.

Holmstrom, B., and P. Milgrom. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica* **55**: 303–28.

Holmstrom, B., and P. Milgrom. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organizations* **7**: 24–52.

Maskin, E., and J. Riley. (1984a). Optimal auctions with risk averse buyers. *Econometrica* **52**: 1473–1518.

Maskin, E., and J. Riley. (1984b). Monopoly with incomplete information. *Rand Journal of Economics* **15**: 171–96.

Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics* **12**: 380–91.

Myerson, R. (1979). Incentive compatibility and the bargaining problem. *Econometrica* **47**: 61–74.

Myerson, R. (1982). Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics* **10**: 67–81.

Nalebuff, B., and J. E. Stiglitz. (1983). Prizes and incentives: Towards a general theory of compensation and competition. *Bell Journal of Economics* **13**: 21–43.

Rogerson, W. (1985a). Repeated moral hazard. *Econometrica* **53**: 69–76.

Rogerson, W. (1985b). The first-order approach to principal-agent problems. *Econometrica* **53**: 1357–68.

EXERCISES

14.B.1[B] Consider the two-effort-level hidden action model discussed in Section 14.B with the general utility function $u(w, e)$ for the agent. Must the reservation utility constraint be binding in an optimal contract?

14.B.2[B] Derive the first-order condition characterizing the optimal compensation scheme for the two-effort-level hidden action model studied in Section 14.B when the principal is strictly risk averse.

14.B.3[B] Consider a hidden action model in which the owner is risk neutral while the manager has preferences defined over the mean and the variance of his income w and his effort level e as follows: Expected utility $= E[w] - \phi \operatorname{Var}(w) - g(e)$, where $g'(0) = 0$, $(g'(e), g''(e), g'''(e)) \gg 0$ for $e > 0$, and $\operatorname{Lim}_{e \to \infty} g'(e) = \infty$. Possible effort choices are $e \in \mathbb{R}_+$. Conditional on effort level e, the realization of profit is normally distributed with mean e and variance σ^2.

(a) Restrict attention to linear compensation schemes $w(\pi) = \alpha + \beta\pi$. Show that the manager's expected utility given $w(\pi)$, e, and σ^2 is given by $\alpha + \beta e - \phi\beta^2\sigma^2 - g(e)$.

(b) Derive the optimal contract when e is observable.

(c) Derive the optimal linear compensation scheme when e is not observable. What effects do changes in β and σ^2 have?

14.B.4[B] Consider the following hidden action model with three possible actions $E = \{e_1, e_2, e_3\}$. There are two possible profit outcomes: $\pi_H = 10$ and $\pi_L = 0$. The probabilities of π_H conditional on the three effort levels are $f(\pi_H \mid e_1) = \frac{2}{3}$, $f(\pi_H \mid e_2) = \frac{1}{2}$, and $f(\pi_H \mid e_3) = \frac{1}{3}$. The agent's effort cost function has $g(e_1) = \frac{5}{3}$, $g(e_2) = \frac{8}{5}$, $g(e_3) = \frac{4}{3}$. Finally, $v(w) = \sqrt{w}$, and the manager's reservation utility is $\bar{u} = 0$.

(a) What is the optimal contract when effort is observable?

(b) Show that if effort is not observable, then e_2 is not implementable. For what levels of $g(e_2)$ would e_2 be implementable? [*Hint*: Focus on the utility levels the manager will get for the two outcomes, v_1 and v_2, rather than on the wage payments themselves.]

(c) What is the optimal contract when effort is not observable?

(d) Suppose, instead, that $g(e_1) = \sqrt{8}$, and let $f(\pi_H \mid e_1) = x \in (0, 1)$. What is the optimal contract if effort is observable as x approaches 1? What is the optimal contract as x approaches 1 if it is not observable? As x approaches 1, is the level of effort implemented higher or lower when effort is not observable than when it is observable?

14.B.5[B] Suppose that in the hidden action model explored in Section 14.B the manager can not only choose how much effort to exert but can also, after observing the realization of the firm's profits π, unobservably reduce them in a way that is of no direct benefit to him (e.g., he can voluntarily offer to pay more for his inputs). Show that in this case there is always an optimal incentive scheme that is nondecreasing in observed profits.

14.B.6[B] Amend the two-effort-level model studied in Section 14.B as follows: Suppose now that effort has distinct effects on revenues R and costs C, where $\pi = R - C$. Let $f_R(R \mid e)$ and $f_C(C \mid e)$ denote the density functions of R and C conditional on e, and assume that, conditional on e, R and C are independently distributed. Assume $R \in [\underline{R}, \bar{R}]$, $C \in [\underline{C}, \bar{C}]$, and that for all e, $f_R(R \mid e) > 0$ for all $R \in [\underline{R}, \bar{R}]$ and $f_C(C \mid e) > 0$ for all $C \in [\underline{C}, \bar{C}]$.

The two effort choices are now $\{e_R, e_C\}$, where e_R is an effort choice that devotes more time to revenue enhancement and less to cost reduction, and the opposite is true for e_C. In particular, assume that $F_R(R \mid e_R) < F_R(R \mid e_C)$ for all $R \in (\underline{R}, \bar{R})$ and that $F_C(C \mid e_C) > F_C(C \mid e_R)$ for all $C \in (\underline{C}, \bar{C})$. Moreover, assume that the monotone likelihood ratio property holds for each of these variables in the following form: $[f_R(R \mid e_R)/f_R(R \mid e_C)]$ is increasing in R, and $[f_C(C \mid e_R)/f_C(C \mid e_C)]$ is increasing in C. Finally, the manager prefers revenue enhancement over cost reduction: that is, $g(e_C) > g(e_R)$.

(a) Suppose that the owner wants to implement effort choice e_C and that both R and C are observable. Derive the first-order condition for the optimal compensation scheme $w(R, C)$. How does it depend on R and C?

(b) How would your answer to **(a)** change if the manager could always unobservably reduce the revenues of the firm (in a way that is of no direct benefit to him)?

(c) What if, in addition, costs are now unobservable by a court (so that compensation can be made contingent only on revenues)?

14.B.7[C] Consider a two-period model that involves two repetitions of the two-effort-level hidden action model studied in Section 14.B. There is no discounting by either the firm or the manager. The manager's expected utility over the two periods is the sum of his two single-period expected utilities $E[v(w) - g(e)]$, where $v'(\cdot) > 0$ and $v''(\cdot) < 0$.

Suppose that a contract can be signed ex ante that gives payoffs in each period as a function of performance up until then. Will period 2 wages depend on period 1 profits in the optimal contract?

14.B.8[C] Amend the two-effort-choice hidden action model discussed in Section 14.B as follows: Suppose the principal can, for a cost of c, observe an extra signal $\tilde{y}$ of the agent's effort. Profits π and the signal y have a joint distribution $f(\pi, y \mid e)$ conditional on e. The decision to investigate the value of y can be made after observing π.

A contract now specifies a wage schedule $w(\pi)$ in the event of no investigation, a wage schedule $w(\pi, y)$ if an investigation occurs, and a probability $p(\pi)$ of investigation conditional on π. Characterize the optimal contract for implementing effort level e_H.

14.C.1[C] Analyze the extension of the hidden information model discussed in Section 14.C where there are an arbitrary finite number of states $(\theta_1, \ldots, \theta_N)$ where $\theta_{i+1} > \theta_i$ for all i.

14.C.2[B] Consider the hidden information model in Section 14.C, but now let the manager be risk neutral with utility function $v(w) = w$. Show that the owner can do as well when θ is unobservable as when it is observable. In particular, show that he can accomplish this with a contract that offers the manager a compensation scheme of the form $w(\pi) = \pi - \alpha$ and allows him to choose any effort level he wants. Graph this function and the manager's choices in (w, e)-space. What revelation mechanism would give this same outcome?

14.C.3[B] Suppose that in the two-state hidden information model examined in Section 14.C, $u(w, e, \theta) = v(w) - g(e, \theta)$.

(a) Characterize the optimal contract under full observability.

(b) Is this contract feasible when the state θ is not observable?

14.C.4[C] Characterize the solution to the two-state principal-agent model with hidden information when the manager is risk averse, but not infinitely so.

14.C.5[B] Confirm that the analysis in Section 14.C could not change if the owner's profits depended on the state and were not publicly observable and if, letting $\pi_i(e)$ denote the profits in state θ_i for $i = L, H$, $\pi'_H(e) \geq \pi'_L(e) > 0$ for all $e \geq 0$. What happens if this condition is not satisfied?

14.C.6[C] Reconsider the labor market screening model in Exercise 13.D.1, but now suppose that there is a single employer. Characterize the solution to this firm's screening problem (assume that both types of workers have a reservation utility level of 0). Compare the task levels in this solution with those in the equilibrium of the competitive screening model (assuming an equilibrium exists) that you derived in Exercise 13.D.1.

14.C.7[B] (J. Tirole) Assume that there are two types of consumers for a firm's product, θ_H and θ_L. The proportion of type θ_L consumers is λ. A type θ's utility when consuming amount x of the good and paying a total of T for it is $u(x, T) = \theta v(x) - T$, where

$$v(x) = \frac{1 - (1 - x)^2}{2}.$$

The firm is the sole producer of this good, and its cost of production per unit is $c > 0$.

(a) Consider a nondiscriminating monopolist. Derive his optimal pricing policy. Show that he serves both classes of consumers if either θ_L or λ is "large enough."

(b) Consider a monopolist who can distinguish the two types (by some characteristic) but can only charge a simple price p_i to each type θ_i. Characterize his optimal prices.

(c) Suppose the monopolist cannot distinguish the types. Derive the optimal two-part tariff (a pricing policy consisting of a lump-sum charge F plus a linear price per unit purchased of p) under the assumption that the monopolist serves both types. Interpret. When will the monopolist serve both types?

(d) Compute the fully optimal nonlinear tariff. How do the quantities purchased by the two types compare with the levels in **(a)** to **(c)**?

14.C.8[B] Air Shangri-la is the only airline allowed to fly between the islands of Shangri-la and Nirvana. There are two types of passengers, tourist and business. Business travelers are willing to pay more than tourists. The airline, however, cannot tell directly whether a ticket purchaser is a tourist or a business traveler. The two types do differ, though, in how much they are willing to pay to avoid having to purchase their tickets in advance. (Passengers do not like to commit themselves in advance to traveling at a particular time.)

More specifically, the utility levels of each of the two types net of the price of the ticket, P, for any given amount of time W prior to the flight that the ticket is purchased are given by

$$\begin{aligned} &\textit{Business:} \quad v - \theta_B P - W, \\ &\textit{Tourist:} \quad v - \theta_T P - W, \end{aligned}$$

where $0 < \theta_B < \theta_T$. (Note that for any given level of W, the business traveler is willing to pay more for his ticket. Also, the business traveler is willing to pay more for any given reduction in W.)

The proportion of travelers who are tourists is λ. Assume that the cost of transporting a passenger is c.

Assume in **(a)** to **(d)** that Air Shangri-la wants to carry both types of passengers.

(a) Draw the indifference curves of the two types in (P, W)-space. Draw the airline's isoprofit curves. Now formulate the optimal (profit-maximizing) price discrimination problem mathematically that Air Shangri-la would want to solve. [*Hint*: Impose nonnegativity of prices as a constraint since, if it charged a negative price, it would sell an infinite number of tickets at this price.]

(b) Show that in the optimal solution, tourists are indifferent between buying a ticket and not going at all.

(c) Show that in the optimal solution, business travelers never buy their ticket prior to the flight and are just indifferent between doing this and buying when tourists buy.

(d) Describe fully the optimal price discrimination scheme under the assumption that they sell to both types. How does it depend on the underlying parameters λ, θ_B, θ_T, and c?

(e) Under what circumstances will Air Shangri-la choose to serve only business travelers?

14.C.9^C Consider a risk-averse individual who is an expected utility maximizer with a Bernoulli utility function over wealth $u(\cdot)$. The individual has initial wealth W and faces a probability θ of suffering a loss of size L, where $W > L > 0$.

An insurance contract may be described by a pair (c_1, c_2), where c_1 is the amount of wealth the individual has in the event of no loss and c_2 is the amount the individual has if a loss is suffered. That is, in the event no loss occurs the individual pays the insurance company an amount $(W - c_1)$, whereas if a loss occurs the individual receives a payment $[c_2 - (W - L)]$ from the company.

(a) Suppose that the individual's only source of insurance is a risk-neutral monopolist (i.e., the monopolist seeks to maximize its expected profits). Characterize the contract the monopolist will offer the individual in the case in which the individual's probability of loss, θ, is observable.

(b) Suppose, instead, that θ is not observable by the insurance company (the individual knows θ). The parameter θ can take one of two values $\{\theta_L, \theta_H\}$, where $\theta_H > \theta_L > 0$ and $\text{Prob}(\theta_L) = \lambda$. Characterize the optimal contract offers of the monopolist. Can one speak of one type of insured individual being "rationed" in his purchases of insurance (i.e., he would want to purchase more insurance if allowed to at fair odds)? Intuitively, why does this rationing occur? [*Hint*: It might be helpful to draw a picture in (c_1, c_2)-space. To do so, start by locating the individual's endowment point, that is, what he gets if he does not purchase any insurance.]

(c) Compare your solution in **(b)** with your answer to Exercise 13.D.2.

14.AA.1^B Show that $[f_e(\pi \mid e)/f(\pi \mid e)]$ is increasing in π for all $e \in [a, b] \subset \mathbb{R}$ if and only if for any $e', e'' \in [a, b]$, with $e'' > e'$, $[f(\pi \mid e'')/f(\pi \mid e')]$ is increasing in π.

14.AA.2^B Consider a hidden action model with $e \in [0, \bar{e}]$ and two outcomes π_H and π_L, with $\pi_H > \pi_L$. The probability of π_H given effort level e is $f(\pi_H \mid e)$. Give sufficient conditions for the first-order approach to be valid. Characterize the optimal contract when these conditions are satisfied.

14.BB.1^B Try solving problem (14.BB.1) by first solving it while ignoring constraint (iv) and then arguing that the solution you derive to this "relaxed" problem is actually the solution to problem (14.BB.1).