MEASUREMENT DISTORTION AND MISSING CONTINGENCIES IN OPTIMAL CONTRACTS

by

Franklin Allen
Douglas Gale

26-90

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367

MEASUREMENT DISTORTION AND MISSING CONTINGENCIES IN OPTIMAL CONTRACTS

Franklin Allen

University of Pennsylvania

Douglas Gale

University of Pittsburgh

and

Boston University

Revised September 1990.

We have benefitted from conversations with Oliver Hart, Michael Riordan and Jean Tirole. Financial support from the National Science Foundation under grants SES-8813719 and SES-8720589 is gratefully acknowledged.

ABSTRACT

Standard contract theory suggests that in optimal contracts payments should be contingent on many events, but in practice this rarely happens. For example, financial securities typically do not make payments contingent on accounting information. This paper develops a theory to explain these missing contingencies. The first important element of the theory is that contracts are based on signals produced by measurement systems which are manipulable. The second is that the contracting parties have incomplete information about each other's type. Given these two assumptions, it is shown that in equilibrium a non-contingent contract is optimal.

I. INTRODUCTION

Contract theory suggests that optimal contracts are often extremely complicated (Hart and Holmstrom, 1987). In particular, Holmstrom (1979) and Shavell (1979) have shown that in the presence of moral hazard, optimal contracts will be contingent on all relevant information. Even if we ignore agency problems, simple risk-sharing arguments suggest that optimal contracts should include many contingencies. But whatever the implications of contract *theory*, it is a commonplace that *in practice* contracts are much simpler. In particular, they leave out contingencies that have the potential to improve risk-sharing. A few examples will make clear the sort of contingencies we have in mind.

Consider first the standard debt contract. Under this contract, the borrower is committed to a fixed schedule of repayments of principal and interest, independently of the state of nature. Optimal risk-sharing requires that borrower and lender have equal marginal rates of substitution between states. But under a debt contract, this condition will clearly be violated, except in very special cases. In addition, the borrower may declare bankruptcy in some states because he cannot meet the non-contingent payment. In that case, both parties may incur substantial deadweight costs.

Some writers have argued that the standard debt contract is an optimal incentive-compatible contract when there is incomplete information (see Townsend, 1979; Diamond, 1984; Gale-Hellwig, 1985). But this argument does not explain why the debt contract does not make repayments contingent on readily available information, such as company profits. It is the absence of this sort of contingency that requires explanation.

A practical example of how such contingencies might be used to improve the risk-sharing is provided by the income bond (Dewing, 1955; McConnell and

Schlarbaum, 1987; De and Kale, 1990). The income bond combines the advantages of debt and preferred stock. Income bonds are a contractual obligation of the issuer, so the interest payments are tax deductible. On the other hand, interest is paid only if earned, which makes the borrower's repayment contingent on accounting profits. Income bonds have been issued in certain periods, but their use has been extremely limited. Before the First World War, income bonds were used with some success by railway companies during reorganizations (McConnell and Schlarbaum, 1987). Income bonds were used in the nineteen-thirties, by companies restructuring their debt, and in the nineteen-forties, mainly by railways undergoing reorganization. A few strong companies have issued income bonds in the ensuing decades, but despite the availability of the income bond "technology" and the theoretical advantages it offers, it has fallen into disuse.

Another example of missing contingencies is provided by interest and currency swaps (Antl, 1986). One of the advantages of the swap is its flexibility. Because the swap is a contract written between two parties, the contract can, in principle, be tailored to their individual needs. Despite this fact and despite the fact that swaps are typically not "traded", the form of the contract remains very simple and is fairly standard. For example, in a fixed to floating swap, one counterparty is exposed to all the risk of fluctuations in the floating rate. Optimal risk-sharing would seem to require some intermediate allocation in which both counterparties shared part of the risk. Although swaps are an important tool for improving risk-sharing, they seem to fall short of the goal of complete, optimal risk-sharing.

Missing contingencies are not limited to financial contracts. For example, risk-sharing is a major concern of contractors developing and producing weapons for the U. S. Defense Department. In his classic study of the weapons acquisition process, Scherer (1964) describes the very limited number of types of contract used by the Defense Department. Most are variants on the polar extremes of the Cost Plus Fixed Fee (CPFF) contract and the Firm Fixed Price (FFP) contract. The contractors' desire to avoid the risk of cost overruns explains the use of the CPFF contract, despite the fact that the CPFF appears to offer no incentive to the contractor to minimize costs. The FFP contract, on the other hand, offers strong incentives but no risk-sharing. Scherer's study describes a number of other types of contracts that combine the features of these two polar cases; but the interesting fact is that the CPFF and FFP contracts are used at all. It might be thought that the FFP, which admits no contingencies, will be used only when the uncertainty about cost is minimal. But the fact that realized profit rates are substantially higher for these contracts suggests that, on the contrary, there is a significant risk premium built in, presumably because the contractor is bearing substantial risk. In any case, the contracts are simpler than one would expect on the basis of optimal contract theory.

In this paper, we develop a theory to explain the absence of contingencies that would allow better risk-sharing. A central role in the theory is played by *measurement systems*. We begin with the idea that it is impossible to write contracts that are contingent on "states of nature". The contingencies incorporated in actual contracts depend on signals produced by measurement systems, which reveal the state of nature only indirectly, if at all. For instance, if bond repayments are contingent on accounting profits, there has to be an accounting system to produce the profit information. Similarly, the cost reimbursement contracts used in the procurement process require accounting systems to specify costs.

A measurment system does not necessarily reveal the true state of nature. In the first place, the system may be "noisy", so that it produces at best imperfect information about the variables that really matter to the contracting parties. Second, and more importantly, the system may be *manipulable*, that is, it can be distorted by one party to the contract.

The fact that measurement systems are manipulable is an important part of our explanation of why some contingencies are not more widely used. But it is not the whole story. Even if the information provided by measurement systems can be distorted, it does not become completely useless. There are costs of distorting the information produced by the measurement system and this will limit the extent of distortion. Even if there is little information in the distorted signal, it should be used in the contract, according to the Holmstrom-Shavell theorem. In order to explain the complete absence of a contingency based on manipulable data, something more is needed.

When two parties to a contract first meet, they have *incomplete information* about each other's type. The proposal to include a certain kind of contingency may be interpreted by one party as a signal about the other's type. If the proposal of the contingency is a bad signal, it may be excluded altogether in equilibrium. The clearest example of this kind of signalling effect seems to arise in the case of income bonds. The evident resistance of the financial community to the issue of income bonds is conventionally attributed to their association with the reorganization of bankrupt or struggling companies. They are said to have "the smell of death about them". (See McConnell and Schlarbaum, 1987 and the quotations therein). In other words, the attempt to issue these securities is interpreted as a bad signal about the issuer.

The interplay between the measurement distortion and incomplete information is crucial in obtaining these results. As we pointed out above, the mere possibility of distorting the signal on which a contingency is based will not eliminate that contingency from the equilibrium contract.

Similarly, it is well known that adverse selection may reduce the amount of insurance contained in equilibrium contracts, but it will not lead to missing contingencies in our sense. Our results require that different types of supplier pool in equilibrium and that they pool at the non-contingent contract. What makes the non-contingent contract special in our theory is that this is the only contract at which no type has an incentive to distort. At a non-contingent contract, on the other hand, different types have positive incentives to distort and furthermore, their marginal incentives are different. This is what prevents pooling. So it is the interaction of adverse selection with measurement distortion that produces a novel result.

There are undoubtedly many reasons why contracts have missing contingencies. In this paper we have only offered one kind of explanation. A closely related approach is found in Holmstrom and Milgrom (1990), in the context of principal-agent problems. They argue that if the agent engages in several tasks, some of which can be monitored more accurately than others, the use of high-powered incentives may distract the agent from productive, poorly monitored tasks to spend more time on less productive, better monitored tasks. The moral hazard problems they consider are analogous to the measurement distortion in our theory, but there is no role for adverse selection. They show that in some cases it may be optimal to eliminate incentives altogether, but this does not seem to be the typical outcome.

There are also a number of papers that have tried to use pure adverse selection arguments to show why contracts have missing contingencies. We review these theories in Section II before setting out a model of contracting with measurement systems in Section III.

We begin by describing a simple contracting problem to serve as a paradigm. This example was chosen mainly on the grounds of tractability. But we are confident that the ideas developed here can be applied to other situations. To simplify the discussion, we restrict attention in Sections II and III to problems with two dimensional contracts. This allows us to describe many of the results diagrammatically. A general analysis is presented in Section IV.

II. CONTRACTING PROBLEMS WITH PURE ADVERSE SELECTION

In this section we study the contracting problem under the assumption that there is no measurement distortion.

There are two agents who want to write a contract for the supply of an indivisible good. The agents are referred to in the sequel as the *purchaser* and the *supplier*, respectively. The supplier can produce at most one unit of the good. The cost of production is uncertain and takes on two values, high (C_H) and low (C_L), with $0 < C_L < C_H$. There are two types of supplier, a bad type (B) and a good type (G). The supplier knows his type; the purchaser does not. The probability of having high costs is π_B for the bad type and π_G for the good type, with $0 < \pi_G < \pi_B < 1$.

The purchaser is assumed to be risk neutral. He places a value of $v > 0$ on one unit of the good. If he obtains the good at a price p his utility is $v - p$. If there is no trade his utility is zero.

The supplier, on the other hand, is risk averse. His preferences are represented by a von Neumann - Morgenstern utility function $U(\cdot)$ with the usual properties. If he produces the good at a cost of c and sells it at a price of p, then his utility is $U(p - c)$. If there is no trade, his utility is $U(0) = 0$.

A contract specifies the payment made by the purchaser to the supplier in the event that the good is exchanged. The payment may be contingent on the cost of production, which is assumed to be observable to both parties.

If there is no trade, there is no payment. Thus, a contract can be represented by an ordered pair (P_H, P_L), where P_i denotes the payment by the purchaser to the supplier when trade occurs and the cost of production turns out to be C_i $(i = H,L)$.

The contracting process is represented by a simple game. The supplier proposes a contract P to the purchaser. The purchaser accepts or rejects the contract. If the purchaser rejects the contract, the game ends: there is no trade and no transfer. If the purchaser accepts the contract, the supplier delivers the good and the purchaser makes the appropriate payment, contingent on the observed cost of production.

Under complete information, there is a unique, efficient equilibrium of this game. This equilibrium is illustrated in Figure 1(a). Since the purchaser is risk neutral, he bears all the risk. The supplier receives a constant net income, independently of the state.

When information is incomplete, the contracting procedure described above defines a standard signalling game. It is well known that such games have many equilibria. We shall focus on the Riley equilibrium (Cho-Kreps, 1987). In this equilibrium, the two types separate. The bad type chooses the most preferred contract that is consistent with the purchaser's individual rationality constraint. The good type chooses his most preferred contract, subject to the purchaser's individual rationality constraint and the constraint that the contract is not preferred by the bad type. The equilibrium is illustrated in Figure 1(b).

Comparing the complete information equilibrium illustrated in Figure 1(a) with the incomplete information equilibrium illustrated in Figure 1(b), we notice that in both cases the bad type gets full insurance. The good type, on the other hand gets strictly less insurance under incomplete information. In this sense, adverse selection reduces the importance of contingencies in the optimal contract. However, there is always some insurance in the good type's equilibrium contract. We never observe a completely non-contingent contract.

Contracting Costs and Non-Contingent Contracts

To address this last point, Spier (1989) has used a similar setup, but assumed that writing contingent contract involved positive fixed costs. By contrast, a non-contingent contract such as (v,v) is costless to write. When there are fixed costs of writing a contingent contract it may be optimal, even in the first best, to choose a completely non-contingent contract. Whether a contingent contract is used depends on all the parameters of the model. For example, if the contracting costs are sufficiently large or if the riskiness of production costs is sufficiently small, it will not pay to use a contingent contract. Similar considerations apply in the analysis of the signalling game.

Comparing the unique equilibrium under complete information with the Riley equilibrium under incomplete information, we find once again that adverse selection reduces the use of contingencies in equilibrium contracts. But the result is now sharper. Spier shows that for some parameter values contingent contracts will be used when there is complete information but not when there is incomplete information. Conversely, if a contingent contract is not used in the complete information equilibrium, it will not be used in the incomplete information equilibrium. In this sense, adverse selection together with fixed costs can be said to cause incomplete contracts.

The explanation for this result is simple. As we saw before, if contingent contracts are chosen, the amount of insurance obtained by the good type is lower under incomplete information than it would be under complete information. The incentive to write a contingent contract is correspondingly less and so the good type is less likely to propose a contingent contract under incomplete information.

In this model, the non-contingent contract (v,v) represents a lower bound on the equilibrium outcome that both types of supplier can receive. Since the two types differ only in the probability distribution of costs, the purchaser is indifferent about the supplier's type when he is offered the non-contingent contract. Adverse selection can force the good type's contract toward the non-contingent contract, but not beyond it.

This property disappears if the two types differ in some other dimension. For example, suppose that they produce different qualities of the good. Let v_G and v_B denote the values of the goods produced by the good and bad types, respectively, and assume that $0 < v_B < v_G$. The Riley equilibrium under incomplete information is illustrated in Figure 2. The good type can no longer be certain of getting the non-contingent contract (v_G, v_G). If he is confused with the bad type, he may be forced to offer a much worse contract, even if it does not contain any contingencies. In order to separate himself from the bad type, the good type may be willing to accept negative insurance, i.e., contingencies that increase his risk, as shown in Figure 2.

Now we can see that, by comparison with the equilibrium under complete information, adverse selection may actually increase the amount of insurance in the equilibrium contract. Suppose for example that the value of insurance is not very great. (This will be the case if the risk is small or the supplier's risk aversion is low). Then even under complete information, the good type may not choose a contingent contract. However, under incomplete information, he may choose a contingent contract in order to separate himself from the bad type. Thus, Spier's result may be reversed when the purchaser cares which type of supplier offers a non-contingent contract.

In contrast to Spier's model, we do not assume that there are fixed costs of writing contingent contracts. Nonetheless, we can show that in a stable equilibrium, agents will choose completely non-contingent contracts. Also, whereas by modifying the Spier model we have shown that adverse selection can increase the use of contingencies, in our model the effect always is to reduce them.

"Corner Solutions" and Non-Contingent Contracts

In work on a related theme, Nachman and Noe (1989) offer another explanation of the use of non-contingent contracts. Their theory explains the use of debt as an optimal contract for financing investment in a firm. Suppose there is a single firm and a single investor. Both agents are risk neutral. There are assumed to be two types of firms, high risk (B) and low risk (G), say. Both types need to raise a fixed amount of capital K in order to finance a risky venture. A contract specifies how much the firm must repay to the investor as a function of the return to the venture. Suppose for simplicity that the risky venture has three possible outcomes, 0, R_L and R_H, where $0 < R_L < R_H$. Since the firm has no resources other than the return to its risky venture, it cannot repay the loan when the return is zero. If a debt contract is used, the debt must be risky. The two types of firm are defined by their probability distributions over returns, which are denoted by $p_B(R)$ and $p_G(R)$. Assume that $p_G(\cdot)$ stochastically dominates $p_B(\cdot)$.

Formally, a contract is a schedule of repayments S(R). The contract is feasible if $S(R) \leq R$ for all R. One additional condition is imposed on the contract: it must be *monotonic*. That is, $S(\cdot)$ must be a non-decreasing function.

The contracting process is represented by the usual game form. The firm offers a contract, which the investor accepts or rejects. If the contract is accepted, the investor pays the cost of the project and the return is divided according to the agreed payment schedule. There are many equilibria of this game. Nachman-Noe propose a refinement which selects a unique equilibrium.

In this equilibrium both types of firm pool at a risky debt contract. How can we explain this result? Note first of all that there cannot be a separating equilibrium. If the two types offered monotonic contracts which had the same expected value (evaluated using the probability distribution of the type offering the contract), the bad type would always prefer the good type's contract to his own. Second, at any pooling equilibrium, the good type will try to separate from the bad type. He can do this by offering a less monotonic contract, which is better for him but worse for the bad type. The only time he cannot do this is when he is already offering the least monotonic contract, that is, the debt contract.

Without the assumption that contracts must be monotonic, the pooling equilibrium will not be stable. Instead, the refinement proposed by Nachman-Noe would select the Riley equilibrium, in which the good firm signals its type by offering a non-monotonic contract. In fact, the good type of firm is accepting negative insurance (a payoff that is negatively correlated with the project return). As in the case of Spier's model, a small change in the model has led to the result that adverse selection actually *increases* the use of contingencies in optimal contracts.

Note also that the Nachman-Noe result is dependent on the assumption of risk neutrality. The choice of optimal contract has no welfare implications in this case.

It is well known that the single-crossing property loses its force on the boundary of the contract space. This seems to be crucial to the Nachman-Noe analysis. The assumption that contracts are monotonic establishes a convenient boundary at the point where we want a stable pooling equilibrium. De and Kale (1990) exploit a similar idea in their analysis of debt packaging. They assume that the firm seeking finance has a binary choice. It can either choose Fixed-Periodic-Obligation Debt (FPOD) contracts or No-Periodic-Obligation Debt (NPOD) contracts. Having offered the market its choice, the market determines the coupons the firm must pay in each period. Under certain conditions, it can be shown that the unique equilibrium that is stable in the sense of Universal Divinity (Banks-Sobel, 1987) is one in which FPOD contracts, i.e, non-contingent contracts, are used. Although we have not had a chance to study their results thoroughly, it appears that the explanation of these results, like those of Nachman-Noe, turns on the fact that the equilibrium contract is on the boundary of the contract space, in this case because of the firm's restriction to a binary choice of contract forms.

Stiglitz-Weiss (1981) and DeMeza-Webb (1987) have used a similar idea to obtain pooling equilibria with credit rationing. In both cases, the power of the single-crossing property is vitiated by adopting assumptions that ensure the equilibrium occurs on the boundary of the contract space. Like the Nachman-Noe and De-Kale constructions, these equilibria seem a little fragile, since small changes in the model could ensure that the equilibrium contract was located in the interior of the contract space. In the case of Stiglitz-Weiss, this possibility has been studied by Bester (198*).

To sum up, these papers have made important contributions to our understanding of the ways in which incomplete information limits the use of contingencies in optimal contracts. However, they may be fragile solutions to the extent that they rely, in different ways, on the use of "corner solutions". The present theory, in contrast, supports a robust equilibrium in which agents pool at a non-contingent contract in the interior of the contract space.

III. A CONTRACTING PROBLEM WITH MEASUREMENT DISTORTION

In Section II, we saw that adverse selection could increase or decrease the completeness of contracts. In this section, we consider what happens when manipulable measurement systems are introduced. In order to introduce the basic ideas with a minimum of technical detail, we again restrict attention to the case of two-dimensional contracts.

As before, the purchaser wants to obtain one unit of an indivisible good from the supplier. The supplier can produce one unit of the good at a cost of $C(\theta)$, where θ is an exogenous random variable. If the good is not produced, there is no cost.

The supplier can observe the cost shock θ directly, but the purchaser cannot. As a result, the contract cannot be made contingent on the cost of production. In order to incorporate risk-sharing in the contract, they make use of a *measurement system.* The measurement system is represented by an exogenous signal S. The signal is publicly observed and is correlated with the supplier's cost. However, the signal can be distorted by the supplier. Let d denote the amount of distortion caused by the supplier and let θ be the true value of the cost shock. Then the value of the signal is assumed to be a function of $\theta + d$.

If the supplier could distort the measurement system without cost, then obviously there would be no possibility of writing a useful contingent contract. But distortion is assumed to take effort and effort is costly. We assume that the supplier's effort is proportional to the amount of distortion and that his utility depends on this effort.

In the *following* section, the general model is analyzed. But first, to introduce the essential ideas, we discuss an illustrative example. There are assumed to be two types of supplier, good (G) and bad (B). The two types are equally likely. The supplier's utility is assumed to be a function U(c,d) of his consumption c and his distortion d. The utility function is assumed to be additively separable:

$$U_k(c,d) = U(c) - B_k|d|, \quad (k = B,G),$$

where $B_G > B_B > 0$ and U has the usual properties. If there is no trade, his utility is U(0) = 0. Note that the two types have the same preferences over consumption. In Section V we indicate how this can be relaxed.

The cost shock θ is a real random variable. For simplicity, assume that θ is uniformly distributed on the interval [a,b]. As in the earlier example, the cost of production is either high or low:

$$C(\theta) = \begin{cases} C_H \text{ if } \theta \geq 0; \\ C_L \text{ if } \theta < 0, \end{cases}$$

where $0 < C_L < C_H$. Likewise, the signal S takes on the two values

$$S(\theta) = \begin{cases} H \text{ if } \theta + d \geq 0; \\ L \text{ if } \theta + d < 0. \end{cases}$$

In the absence of distortion, the signal reveals the true cost of production. Otherwise, it gives imperfect information about costs.

The payment to the supplier cannot be made contingent on costs since costs are not directly observable. Instead, the payment must be made contingent on the signal produced by the measurement system. Because the signal can take on only two values, a contract is represented by an ordered pair $P = (P_H, P_L)$, where P_i represents the payment contingent on the value of the signal i = H,L.

The purchaser is assumed to be risk-neutral. One unit of the good purchased from a supplier of type k has a utility v_k for the purchaser. If the purchaser buys the good from a type k supplier at a price p, his utility is $v_k - p$. If there is no trade, his utility is zero. We assume that

$$0 < v_B < v_G$$

so that the bad type produces an inferior good and has lower costs of distortion than the good type.

The assumption that the supplier with the lower costs of distortion produces the lower quality good plays a crucial role in the analysis to follow. The sort of situation we have in mind to motivate this assumption is the following. Company A is a high quality company that has invested in information systems and a highly professional management over a number of years. As a result, it has good costs and quality control which allow it to develop a high quality good. At the same time, the institutional structure of the firm (the existence of high quality data and the professionalism of the staff) makes it costly to distort the true costs of production. Company B, on the other hand, is a fly-by-night outfit. Its accounting and information systems are chaotic or non-existent and management is either incompetent or dishonest. These characteristics mean that their product is likely to be inferior and, for the same reason, it is relatively easy for them to manipulate accounting data.

Of course, one could easily tell a story where distortion costs and quality were positively correlated. The results then would be quite different. We do not want to claim that negative correlation is the only case worth considering, only that it is not implausible and produces interesting results.

The trading process is represented by a simple game. The supplier proposes a contract which the purchaser can accept or reject. If the contract is rejected, the game ends. Otherwise, the good is produced, the supplier observes the cost shock and decides how much effort to exert to distort the signal.[1] Finally, the signal is observed, the good is exchanged and payment is made.

This game can be reduced to a standard signalling game, once we note that the supplier's optimal distortion is uniquely determined as a function of θ and P. Suppose that a contract P has been accepted, the good has been produced and the supplier observes that the cost shock is θ. He has to decide whether and by how much to distort the signal. Suppose that $P_H > P_L$. Obviously, if the supplier chooses to distort the signal, he will use the smallest effort needed to send the high signal. Thus, the distortion will be $d = -\theta$ or zero. He will never distort the signal when $\theta \geq 0$ since this involves positive effort and gains nothing. On the other hand, it is worthwhile taking effort if $\theta_k < \theta < 0$, where $\theta_k = -(P_H - P_L)/B_k$, since the cost of distortion is less than the gain. For $\theta < \theta_k$, no effort is taken since the cost is greater than the gain. In each case, there is a unique choice of distortion and a unique payoff associated with every value of θ. The supplier's income under the optimal policy is given by

$$W(\theta) = \begin{cases} P_L - C_L & \text{if } \theta < \theta_k \\ P_H - C_L & \text{if } \theta_k < \theta < 0 \\ P_H - C_H & \text{if } \theta > 0. \end{cases}$$

If $P_H < P_L$, the situation is similar except that the discontinuity of the signal seems to prevent the existence of an optimal distortion. There are various ways this could be dealt with; we shall simply ignore it and assume the set of distortions that achieves a particular value of the signal is closed.[2] In that case, there is no distortion if $\theta > \theta_k = -(P_H - P_L)/B_k$ whereas the optimal distortion is $d = -\theta$ if $\theta > \theta_k$. The supplier's income under the optimal policy is

$$W(\theta) = \begin{cases} P_L - C_L & \text{if } \theta < 0 \\ P_L - C_H & \text{if } 0 < \theta < \theta_k. \\ P_H - C_H & \text{if } \theta > \theta_k \end{cases}$$

In either case, i.e., whether $P_H > P_L$ or $P_H < P_L$, the maximized expected utility when a contract P is accepted is given by the formula

$$U^*(P,k) = \text{Prob}\{\theta < \theta_k\}U(P_L - C_L) + \text{Prob}\{\theta_k < \theta < 0\}U(P_H - C_L)$$
$$+ \text{Prob }\{0 < \theta < \theta_k\}U(P_L - C_H) + \text{Prob}\{\theta > \theta_k\}U(P_H - C_H)$$
$$- \int_{\theta_k}^{0} B_k\theta d\nu.$$

By similar reasoning, there is a unique expected utility for the purchaser associated with each contract P and type k:

$$V^*(P,k) = v_k - \text{Prob}\{\theta < \theta_k\}P_L - \text{Prob}\{\theta > \theta_k\}P_H.$$

In what follows, we can use these "payoff functions" to describe a reduced form signalling game which is equivalent to the original game. In this game, the supplier proposes a contract P, which the purchaser accepts or rejects. If the contract is accepted, the payoffs are $U^*(P,k)$ and $V^*(P,k)$. Otherwise, both agents receive zero.

In this framework we can find equilibria in which only non-contingent contracts are used. Unlike the examples considered in Section 2, these equilibria are quite robust. We begin by examining a *pooling equilibrium.* Let $P^* = (p,p)$ denote a non-contingent contract. Certain conditions must be satisfied if both types of supplier are to propose the same non-contingent contract P^* in equilibrium. First, the contract must be acceptable to the purchaser, so we assume the non-contingent payment is no greater than the expected value of the good:

$$p \le (v_H + v_L)/2.$$

Second, the bad type must prefer the pooling contract to the best contract that he could get if he were to reveal his type:

$$\mathbb{E}\{U(p - C(\theta))\} \geq \sup \{U^*(P,B): V^*(P,B) \geq 0\} > 0.$$

P^* cannot be an equilibrium contract unless this condition is satisfied because, in a sequential equilibrium, the purchaser must accept any contract that satisfies $V^*(P,B) > 0$.

Now we are ready to define equilibrium strategies. The strategy for both types of supplier is to propose the contract P^*. The purchaser's strategy is to accept any contract P such that $V^*(P,B) \geq 0$ or $P = P^*$ and reject the rest. Under the maintained assumptions, this is clearly an equilibrium. First of all, the purchaser cannot do better than to accept the contract that is offered, since by assumption it is individually rational. Second, both types of supplier prefer the contract P^* to any other contract that the purchaser will accept. Since the equilibrium contract gives them a positive payoff, it must be optimal for them to propose it.

It can also be shown that the equilibrium is sequentially rational. Suppose the purchaser believes that any contract $P \neq P^*$ is offered by the bad type. Then his decision to reject any contract $P \neq P^*$ such that $V^*(P,B) < 0$ is optimal relative to his beliefs.

There are many sequentially rational equilibria, each of them supported by more or less arbitrary off-the-equilibrium-path beliefs, so it is not surprising that non-contingent contracts are used in some equilibria. However, the pooling equilibrium is in fact robust to precisely the sort of arguments that are often used to eliminate equilibria supported by "unreasonable" off-the-equilibrium-path beliefs. The reason can perhaps be grasped intuitively from the diagram in Figure 3. At a non-contingent contract, there is no incentive to distort the measurement system. Since the preferences of good and bad suppliers differ only with respect to the disutility of distortion, their indifference curves will be tangent at a non-contingent contract. Moreover, the bad type's indifference curve must lie below the indifference of the good type. This follows from a revealed preference argument. Notice first that both types obtain the same expected utility from a non-contingent contract. This follows from the fact that there is no distortion and their utility of consumption is the same. At a contingent contract, on the other hand, the bad type will enjoy a higher expected utility than the good type, simply because his disutility of distortion is lower. Any contract which is weakly preferred to the equilibrium contract by the good type will be strictly preferred by the bad type. Thus it is impossible for the good supplier to signal his type, as he could if the single-crossing property were satisfied. In fact, nestedness of the indifference curves here is almost equivalent to Universal Divinity.

In addition to the pooling equilibrium, there is a unique *separating equilibrium* that satisfies the usual criteria of strategic stability. This equilibrium is illustrated in Figure 4. Two properties of this equilibrium are noteworthy. *First*, one of the types (the good type in this case) chooses a completely non-contingent contract. Thus, even if separation occurs, there will still be some use of non-contingent contracts. *Second*, both types are worse off than they would be in the pooling equilibrium. We may therefore argue that the separating equilibrium is less plausible than the pooling equilibrium. In order to support the separating equilibrium, the purchaser must reject non-contingent contracts that are preferred by both types. This is rational if his beliefs are sufficiently negative, i.e., put sufficiently high weight on the bad type. But when faced with a non-contingent contract that both types prefer, why should he assume that this is a deviation by the bad type. It might be more reasonable to argue as follows: "Every type prefers the pooling equilibrium to the separating equilibrium. Therefore, if I observe the pooling contract, it is probably because all types are coordinating on the pooling equilibrium. If so, my beliefs ought to be determined by my prior probabilities and I should accept the contract." This sort of reflection will cause the pooling equilibrium to be observed.

Since the supplier's preferences do not satisfy the single-crossing property, it might be thought that there will be other stable equilibria, but in fact there are not. In particular, it can be shown that in any equilibrium satisfying the Cho-Kreps Intuitive Criterion, *there will never be pooling at a contingent contract*. The reason is simply that the tangency of indifference curves illustrated in Figure 3 occurs only on the diagonal, i.e., only on the set of non-contingent contracts. Elsewhere, a version of the single-crossing property will be satisfied. More precisely, above the diagonal the bad type's indifference curve will be flatter than the good type's and below the diagonal it will be steeper. This is precisely what is needed to "destabilize" pooling equilibria.

The crucial element here is, of course, the ability of the supplier to distort the signal on which contingencies are based. The bad type has a stronger preference for the inclusion of contingencies in contracts simply because he has a greater ability to exploit any contingency for his own advantage. Since this effect is symmetric with respect to the inclusion of positive or negative insurance, it explains the stability of an *interior* pooling equilibrium.

Note that there is no requirement that distortion costs be low. In this particular model, one can make the costs of distortion arbitrarily large.[3] The amount of distortion that will occur in equilibrium will then be very small, but the results described above will still hold. It is only the *difference* in distortion costs between types that matters.

IV. THE GENERAL CONTRACTING PROBLEM

In this section we consider the general version of the contracting problem discussed above. As before, there are two agents, a supplier and a purchaser. The supplier can produce one unit of an indivisible good which he exchanges with the purchaser for money. Now, however, there are many possible types of supplier and many possible observations of the measurement system. The contracting process is represented by the usual game.

In this framework, we derive the following results. We define a non-contingent contract to be one in which the delivery price is independent of information that is publicly available. First, we show that there exists an equilibrium in which a completely non-contingent contract is always chosen. Furthermore, this equilibrium is stable in the (quite strong) sense of Universal Divinity. Thus, the failure to include potentially useful contingencies in the contract is a robust phenomenon. Second, in any equilibrium that satisfies the (relatively weak) Intuitive Criterion of Cho-Kreps, the best types must always choose a completely non-contingent contract. More precisely, we can show that there is an interval of types, including the best type, who choose the same, non-contingent contract in equilibrium. This further strengthens the argument that the use of non-contingent contracts is a robust phenomenon, but it also shows that where contingent contracts are used, they will be used by unattractive firms. We can also show that when some types of supplier use contingent contracts, all types of supplier are worse off than they would be in the completely non-contingent equilibrium. It can be argued, by a kind of "forward induction", that the completely non-contingent equilibrium is more likely to be observed.

The Model

There is a finite number of different types of the supplier indexed by $k = 1,\ldots,\ell$. We let K stand for the set of types. For any finite set S, $\Delta(S)$ denotes the set of probability distributions on S. The prior probability distribution of the supplier's type is denoted by $\nu \in \Delta(K)$.

The cost function is denoted by $C(\theta)$, where θ is a real random variable, continuously distributed with support [a,b]. For simplicity, we assume that $C(\cdot)$ takes on a finite number of values indexed by $c_1,\ldots,c_i,\ldots,c_n$. Let $E_i = C^{-1}(c_i) \equiv \{\theta|C(\theta) = c_i\}$ denote the set of states θ in which the cost is equal to c_i. Again for simplicity, we assume that E_i is an interval for each $i = 1,\ldots,n$.

The signal S, representing the measurement system, is assumed to be a function of the distorted shock $\theta + d$. In the absence of distortion, the measurement system is perfectly informative, that is, $S(\theta) \equiv C(\theta)$.[4]

The supplier's preferences are described by a von Neumman-Morgenstern utility function that is assumed to be additively separable:

$$U_k(c,d) = U(c) + T_k(d).$$

The purchaser places a value v_k on the good produced by a supplier of type k. If he purchases the good from supplier k at a price of p, then his utility is $v_k - p$. Otherwise his utility is zero.

The following regularity assumptions are maintained throughout:

(A.1) $v_k < v_{k+1}$ for $k = 1,\ldots,\ell - 1$.

(A.2) (i) $T_k(0) = 0$ and $T_k(d) < 0$ for any $d \neq 0$ and $k = 1,\ldots,\ell - 1$,

(ii) For any d, d′ such that $|d'| < |d|$,

$$T_k(d') - T_k(d) > T_{k+1}(d') - T_{k+1}(d).$$

(A.3) $U(\cdot)$ and $T_k(\cdot)$ are C^1 and strictly concave and $U(\cdot)$ is increasing.

The meaning of (A.1) is obvious. (A.2i) says that distortion is costly and, together with (A.3), implies that the costs of distortion are strictly increasing in the absolute value of the distortion. (A.2ii) says, roughly, that the marginal costs of distortion are strictly increasing in type. But note that we only compare absolute values. Again, (A.3) is self-explanatory.

We also assume that mutually advantageous trade is possible between the purchaser and every type of supplier. The most that the purchaser is willing to pay for a unit of the good from a supplier of type k is v_k. Then trade is always possible if

(A.4) $\mathbb{E}\{U(v_k - C(\theta))\} > 0, \forall k \in K.$

where $\mathbb{E}\{\cdot\}$ denotes the expectation operator.

A *contract* specifies the payment for the good, contingent on the observed value of the signal, assuming that trade takes place. The contract can be identified with a vector $P = (P_1,\ldots,P_n)$, where P_i is the payment required when $S = c_i$. For some purposes, it is convenient to treat the contract as a function of the distorted shock $\theta + d$. In that case, we write the payment as $P(\theta + d)$, where of course $P(\theta + d) = P_i$ for all $\theta + d \in E_i$.

The Reduced Form Game

We use the extensive form introduced in Section III. The supplier proposes a contract, which the purchaser accepts or rejects. If the contract is accepted, the supplier produces the good, observes the cost shock and distorts the signal. After the signal is observed, the good is delivered and the appropriate contingent payment is made.

We analyze this game using the concept of Perfect Bayesian Equilibrium. As applied to the complete extensive form game, Perfect Bayesian Equilibrium requires that four conditions be satisfied. *First*, given that a contract P has been accepted and a cost shock θ has been observed, the supplier must choose a distortion d optimally, as a function of P and θ. *Second*, once the supplier has proposed a contract P, the purchaser's decision to accept or reject the contract must be made optimally, given the purchaser's beliefs about the supplier's type and his anticipation of the supplier's distortion of the signal. *Third*, the supplier must propose an optimal contract, given the optimal response of the purchaser and the anticipated optimal distortion of the supplier himself. *Finally*, the purchaser's beliefs must be consistent with the equilibrium strategies of the supplier.

Rather than deal with the complete extensive form game, we use this informal description of an equilibrium of the entire game to eliminate the last two stages of the game, thus reducing it to a standard signalling game. This reduced form game is then analyzed in the sequel.

For analytical purposes, it is useful to represent the supplier's choice variable as the value of the signal S rather than the amount of distortion d. Consider the decision of the supplier once the contract P has been accepted and the cost shock θ has been observed. The supplier can choose among n different payments by choosing an appropriate distortion d. The minimum cost of obtaining P_i is

$$T_{ki}(\theta) = \sup \{T_k(d) | \theta + d \in E_i\}.$$

Then the supplier will choose the payment P_i that maximizes $U(P_i - C(\theta)) + T_{ki}(\theta)$.

To reduce the extensive form game to a two-stage game, we must show that the supplier's optimal distortion decision is essentially unique.

LEMMA: Suppose that a contract P has been accepted. Then for all but a finite number of values of θ, there is a unique value of i that maximizes $U(P_i - C(\theta)) + T_{ki}(\theta)$.

PROOF: To see this, we need to consider a number of cases. Let i, j and h ($i \neq j$) be arbitrarily given and suppose that θ belongs to E_h. Let $\bar{E}$ denote the closure of E. Let x_i and x_j be the points in $\bar{E}_i$ and $\bar{E}_j$, respectively, that minimize the distance to θ. Then the supplier will be indifferent between P_i and P_j if and only if

$$U(P_i - C_h) + T_{ki}(\theta) = U(P_j - C_h) + T_{kj}(\theta).$$

Case (i). Suppose first of all that i = h. (The case j = h is exactly similar). Then

$$U(P_h - C_h) = U(P_j - C_h) + T_{kj}(\theta).$$

Since $T_k(d)$ decreases when the absolute value of d increases, it is clear that there can be at most one value of θ for which this condition is satisfied.

Case (ii). In what follows, we can suppose without loss of generality that $i \neq h \neq j$. Suppose that θ lies between x_i and x_j, say $x_i < \theta < x_j$. Then

$$U(P_i - C_h) - U(P_j - C_h) = T_k(x_j - \theta) - T_k(\theta - x_i).$$

It is clear from the fact that the distortion cost is increasing in absolute distortion that an increase in θ will raise $T_k(x_j - \theta)$ and reduce $T_k(\theta - x_i)$ and that a decrease in θ will have the opposite effects. There is at most a single value of θ for which this equation can be satisfied.

Case (iii). Now suppose that both x_i and x_j lie to the right of θ. (The case where both lie to the left is exactly similar). Then

$$U(P_i - C_h) - U(P_j - C_h) = T_k(\theta - x_j) - T_k(\theta - x_i).$$

The strict convexity of the distortion cost ensures that there is at most one value of θ in E_h that can satisfy the equation. If $x_i < x_j$, say, then an increase in θ will increase both $T_k(\theta - x_j)$ and $T_k(\theta - x_i)$, but the increase in $T_k(\theta - x_j)$ will be greater than the increase in $T_k(\theta - x_i)$. Exactly the same argument works if θ decreases or if $x_j < x_i$.

We have shown that for any choice of i, j and h there is at most one value of θ in E_h at which the supplier is indifferent between P_i and P_j. There is a finite number of choices of i, j and h. Thus, under the maintained assumptions, there is a unique optimal distortion for the supplier for all but a finite number of values of θ.∎

Let H denote the set of functions from $\mathbb{R}\times\mathbb{R}^n$ to $\{e_1,\dots,e_n\}$ where $e_i = (0,\dots,0,1,0,\dots,0)$ is an n-vector with a one in the i-th place and zeroes elsewhere. If $h \in H$ then we interpret $h(\theta,P)$ as the choice of signal to send conditional on the contract P and the cost shock θ.

Since we have shown that the selection function h is essentially unique for any contract P, we can express *both* players' equilibrium payoffs as functions of the equilibrium contract P.[5] Let $U^*(P,k)$ denote the expected utility of a supplier of type k if a contract P is accepted by the purchaser. Let $V^*(P,k)$ denote the expected utility of the purchaser who has accepted a contract P from a supplier of type k. Then we can define the (reduced form) *payoff functions* U^* and V^* by putting

$$U^*(P,k) = \mathbb{E}\Big[\max_i\{U(P_i - C(\theta)) + T_{ki}(\theta)\}\Big]$$

and

$$V^*(P,k) = \mathbb{E}\Big[v_k - \textstyle\sum_{i=1}^{\ell} h_i(\theta,P)P_i\Big],$$

where h is the unique best response for a supplier of type k.

Equilibrium

From now on, we use these reduced form payoff functions to describe the game. A *pure strategy* for the supplier, in the reduced form game, is a choice of contract $P \in \mathbb{R}^n$. A *pure strategy* for the purchaser, in the reduced form game, is a decision rule $\alpha:\mathbb{R}^\ell \to [0,1]$ with the usual interpretation. In what follows, we only consider pure strategy equilibria.

Let $\Delta(K)$ denote the set of probability distributions on the set K. The purchaser's beliefs about the supplier's type are represented by a function $\mu:\mathbb{R}^n \to \Delta(K)$, with the interpretation that $\mu(P)$ is the purchaser's probability assessment if he observes the contract P. Let $m \in \Delta(K)$ denote the purchaser's beliefs about the supplier's type. Extend the purchaser's payoff function to $\mathbb{R}^\ell \times \Delta(K)$ by putting

$$V^*(P,m) = \sum_{k\in K} V^*(P,k)m(k),$$

where $V^*(P,m)$ denotes expected utility when beliefs are given by m. A *Perfect Bayesian Equilibrium* is defined to be a strategy profile $(\{P_k\},\alpha)$ together with an assignment of beliefs μ such that the following conditions are satisfied:

(E.1) $\quad P_k \in \operatorname{Arg\,Max}_P \alpha(P)U^*(P,k), \ \forall k \in K;$

(E.2) $\quad \alpha(P) \in \operatorname{Arg\,Max}_a aV^*(P,\mu(P)), \ \forall P \in \mathbb{R}^n;$

(E.4) $\quad \forall k \in K, \ \forall P \in \mathbb{R}^n$, if $P_k = P$ then

$$\mu(P)(k) = \frac{\nu_k}{\sum_{\{j:P_j=P\}} \nu_j}.$$

The first condition simply requires the supplier to choose his contract optimally. The second condition requires the purchaser's acceptance decision to be optimal, given his beliefs, for every contract, not just those that are proposed in equilibrium. The third condition requires beliefs to be consistent with equilibrium strategies.

Refinements

In analyzing the stability of equilibrium, two refinements are used. The first is the *Intuitive Criterion* of Cho-Kreps. Let $BR_0(P,m)$ denote the set of *best responses* for the purchaser to the contract P when beliefs are given by m. That is,

$$BR_0(P,m) = \underset{\alpha}{\text{Arg Max}}\ \alpha(P)V^*(P,m).$$

Let $\Delta(K')$ denote the set of probability distributions in $\Delta(K)$ concentrated on the set $K' \subset K$. For any $K' \subseteq K$, let

$$BR_0(P,K') = \cup\{a \in BR_0(P,m) \mid m \in \Delta(K')\}$$

denote the set of best responses of the purchaser for some beliefs in $\Delta(K')$. An equilibrium $(\{(P_k\},\alpha)$ *fails to satisfy the Intuitive Criterion* if and only if there exist a non-empty set K′ in K and a contract P such that:

(i) $\forall k \in K \backslash K'$, $aU^*(P,k) < U_k^*$, for any $a \in BR_0(P,K)$;

(ii) $\exists k \in K'$, $aU^*(P,k) > U_k^*$ for any $a \in BR_0(P,K')$.

Note that we are assuming here that in all equilibria, the contract is accepted with probability one. This will be true for the equilibria considered in the sequel.

Next we define the notion of Universal Divinity, a notion that is weaker than stability in the sense of Kohlberg-Mertens (1986) but stronger than the Intuitive Criterion. Choose some fixed but arbitrary equilibrium $(\{P_k\},\alpha,\mu)$. For any $k = 1,\ldots,\ell$ and any $P \in \mathbb{R}^n$ let

$$R(k|P) = \{a \in BR(P,K) | U^*(P,k) > U_k^*\}$$

and

$$R^0(k|P) = \{a \in BR(P,K) | U^*(P,k) = U_k^*\}.$$

Universal Divinity eliminates a strategy P for the supplier of type k, i.e., requires that $\mu(P)(k) = 0$, if

$$R(k|P) \cup R^0(k|P) \subset \cup_{k' \neq k} R(k'|P).$$

Universal Divinity provides a criterion for deciding which type is most likely to make a particular deviation from equilibrium, and concentrates the purchaser's beliefs on that type. Type k is assumed to be "unlikely" to choose a non-equilibrium contract P if, whenever type k weakly (strictly) prefers that contract for some best response from the purchaser, there is some type $k' \neq k$ that strictly (strictly) prefers it. This is a very strong refinement of equilibrium. It eliminates a strategy for a type as long as there is some combination of types that in aggregate seem more likely to use the strategy.

Pooling Equilibrium

An equilibrium $(\{P_k\},\alpha)$ is called a *pooling equilibrium* if $P_k = P$ for every k in K. Note that in a pooling equilibrium, the purchaser may receive some positive surplus, that is, we may have $V^*(P,\nu) > 0$.

A contract P is said to be *non-contingent* if and only if $P_i = P_j$, for all i and j. Otherwise it is said to be *contingent*. Note that the definition says nothing about whether the payments actually made in equilibrium are contingent. For example, it is conceivable that, by distorting the signal, the supplier can ensure that he gets the same payment in each state of nature, even though the contract is contingent in the sense that it specifies different payments in different states. However, this cannot happen in "stable" equilibria. In fact, the payments made in equilibrium will be contingent if and only if the contract is contingent in the sense of the definition. Obviously, if the contract is non-contingent the payment made in equilibrium is non-contingent. On the other hand, if the contract is contingent, the equilibrium payments can be non-contingent only if there is distortion. But in that case, the supplier would be better off proposing an equivalent non-contingent contract and saving the distortion cost. Furthermore, the purchaser should accept the alternative non-contingent contract since his payments would be the same and the contract would not attract types that are worse than the one offering the original contract.[6]

The first result is that there exists a Universally Divine equilibrium in which only non-contingent contracts are used. Let $\tilde{U}_k$ denote the maximum expected utility a supplier of type k can obtain by offering an acceptable contract to the purchaser when the purchaser knows his type. That is,

$$\tilde{U}_k = \sup\left\{U(P,k) \mid V^*(P,k) \geq 0\right\}.$$

PROPOSITION 1: Let $(\{P_k\},\alpha,\mu)$ be a pooling equilibrium in which all types choose a non-contingent contract P^*. Suppose that for any contract $P \neq P^*$ the purchaser's beliefs are concentrated on the worst type $k = 1$. Suppose further that $\tilde{U}_1 < U_1^*$. Then the equilibrium is Universally Divine.

PROOF: Note first of all that if a non-contingent contract is chosen, the optimal distortion is zero for all k and θ. Hence, for any k,

$$U_k^* = \mathbb{E}\{U(P_k^*(\theta + \delta_k(\theta)) - C(\theta)) + T_k(\delta_k(\theta))\} = \mathbb{E}\{U(p - C(\theta))\},$$

where p is the payment specified by the non-contingent contract. The important point is that all types of supplier receive the same payoff from the non-contingent contract.

Consider some alternative contract P. If P is a non-contingent contract, all types receive the same payoff from P, so the Universal Divinity criterion is automatically satisfied. Now suppose that P is contingent. As noted previously, there must be positive distortion for some values of θ. It is easy to see that if a type k weakly prefers P to P^* then any type $k' < k$ must strictly prefer it since k' has lower costs of distortion that k. In other words, $U^*(P,k) < U^*(P,k')$.

From these facts it is easy to see that the pooling equilibrium is Universally Divine. For any P, $R(1|P) = \phi$ implies that $R(k|P) \cup R^0(k|P) = \phi$ for any $k = 2,\ldots,\ell$. On the other hand, if $R(1|P) \neq \phi$, then the preceding argument shows that for some $a \in BR(P,K)$, $aU^*(P,1) = U_1^*$ and $aU^*(P,k) < U_k^*$ for any $k = 2,\ldots,\ell$. Then we cannot have

$$R(1|P) \cup R^0(1|P) \subset \cup_{k' \neq 1} R(k'|P).$$

So the pooling equilibrium satisfies the definition of Universal Divinity.∎

The next step is to show that other pooling equilibria do not satisfy reasonable stability properties. A crucial step in this argument is an envelope theorem. Recall that the supplier's decision problem at the distortion stage can be represented as the choice of an optimal selection function h:

$$\max_{h \in H} \mathbb{E}\{\textstyle\sum_{i=1}^n h_i(\theta,P)[U(P_i - C(\theta)) - T_{ki}(\theta)]\}.$$

As we saw earlier, there is an essentially unique optimal value of h for each contract P.

PROPOSITION 2: $U^*(\cdot,k)$ is a continuously differentiable function for any value of k. The derivative is given by the formula:

$$\partial U^*(P,k)(z) = \mathbb{E}\{\textstyle\sum_{i=1}^n h_i(\theta)U'(P_i - C(\theta))z_i\},\ \forall z \in \mathbb{R}^n,$$

where h is the optimal selection function.

PROOF: Without ambiguity we can suppress the reference to k in what follows. Let P be a fixed but arbitrary contract and let δP be some variation in P. Let $h(\theta,P)$ and $h(\theta,P + \delta P)$ denote the optimal selections for the contracts P and $\delta P + P$. Let $W_i(\theta,P) \equiv U(P_i - C(\theta)) + T_i(\theta)$ and let $W(\theta,P) = (W_1(\theta,P),\ldots,W_n(\theta,P))$. Then

$$U^*(p + \delta P) - U^*(P) = \mathbb{E}\{h(P + \delta P)\cdot W(P + \delta P) - h(P)\cdot W(P)\}$$

$$= \mathbb{E}\{(h(P + \delta P) - h(P))\cdot W(P + \delta P) +$$

$$h(P)\cdot(W(P + \delta P) - W(P)\}.$$

Consider the first term on the right.

$$0 \le \mathbb{E}\{(h(P + \delta P) - h(P))\cdot W(P + \delta P)$$

$$= \mathbb{E}\{(h(P + \delta P) - h(P))\cdot(W(P) + \partial W(P)\cdot\delta P + o(\delta P))\}$$

$$\le \mathbb{E}\{(h(P + \delta P) - h(P))\cdot(\partial W(P)\cdot\delta P + o(\delta P))\}$$

since

$$0 \ge \mathbb{E}\{(h(P + \delta P) - h(P))\cdot W(P)\}.$$

Suppose that $h(P + \delta P)$ converges weakly to $h(P)$ as δP converges to zero. Then the first term is $o(\delta P)$. Hence,

$$U^*(p + \delta P) - U^*(P) = \mathbb{E}\{h(P)\cdot\partial W(P)\delta P\} + o(\delta P).$$

It remains to show that $h(P + \delta P)$ converges weakly to $h(P)$ as δP converges to zero. If not, select a weakly convergent subsequence (this exists by a standard result). It is straightforward to show that the limit of the sequence is optimal in the limit. Since the optimal selection is unique, we have a contradiction, as required.∎

Using this envelope theorem, we can show that equilibria in which there is pooling at a contingent contract fails the Intuitive Criterion.

PROPOSITION 3: Suppose there is an equilibrium in which different types of supplier pool at a contingent contract. Then the equilibrium fails to satisfy the Intuitive Criterion.

PROOF: The proof is by contradiction. We assume that, contrary to what is to be proved, two or more types of supplier offer the same, contingent contract. Let P^C denote the contract chosen by these types and let k^C denote the highest type that chooses the contract. Then two further sub-cases must be considered according to whether (a) every type $k > k^C$ strictly prefers his equilibrium contract to P^C or (b) some type $k > k^C$ is indifferent between his equilibrium contract and P^C.

Case (a). In this case, we can find a contract P' very close to P^C which is preferred by the supplier of type k^C and by the purchaser. To see this, note that any contract sufficiently close to P^C which is offered only by type k^C must be strictly preferred by the purchaser, since the higher type distorts less than the lower types and produces a good of higher value. Then it is only necessary to show that the contract is preferred by k^C and strictly inferior to the equilibrium contracts for $k < k^C$. Let m and M denote the values of i for which P_i is respectively a minimum and a maximum. Define P' by putting

$$P'_i = \begin{cases} P^C_i \text{ for } 1 < i < n; \\ P^C_M - \varepsilon_M \text{ for } i = M; \\ P^C_m + \varepsilon_m \text{ for } i = m. \end{cases}$$

For any type k,

$$U^*(P',k) - U^*(P^C,k) = \mathbb{E}\{h_{km}(\theta)U'(P_i - C(\theta))\varepsilon_m - h_{kM}(\theta)U'(P_i - C(\theta))\varepsilon_M\} + o(\varepsilon)$$

from the envelope theorem. We claim that for some ε_m, $\varepsilon_M > 0$ sufficiently small, $U^*(P',k) - U^*(P^C,k) > 0$ for $k \geq k^C$ and $U^*(P',k) - U^*(P^C,k) < 0$ for $k < k^C$. To prove this, we make the following observations about the selections h_k.

Claim 1: For any $k = 1,\ldots,K-1$, $h_{k+1m} > h_{km}$.[7] Suppose that $h_{k+1m}(\theta) = 1$ for some k and θ. Then it must be the case that θ belongs to E_m and $d = 0$. Optimality requires that for any $d \neq 0$,

$$U(P^C(\theta + d) - C(\theta)) + T_{k+1}(d) \leq U(P^C(\theta))$$

which in turn implies that

$$U(P^C(\theta + d) - C(\theta)) + T_k(d) < U(P^C(\theta)).$$

Thus

the first inequality holds as an equation (and there must be one such) then there is a non-negligible set of θ for which $h_{k+1}(\theta) = 1$ and $h_{km}(\theta) = 0$.

Claim 2: For any $k = 1,\dots,K - 1$, $h_{k+1M} < h_{kM}$. Obviously, $d = 0$ is optimal for any θ in E_M and any k. So we only need to consider cases where $h_{k+1M}(\theta) = 1$ and $d = \delta_k(\theta) \neq 0$. Obviously, it cannot be optimal for k to choose d′ with $|d'| > |d|$; conversely, for any $|d'| < |d|$ we have

$$T_k(d) - T_k(d') > T_{k+1}(d) - T_{k+1}(d')$$

so that

$$U(P_M - C(\theta)) + T_k(d) > U(P(\theta + d') - C(\theta)) + T_k(d').$$

There must also be a non-negligible set of θ such that $h_{kM}(\theta) = 1$ and $h_{k+1M}(\theta) = 0$ for the usual reasons.

Returning to the comparison of P′ and P^C we see that

$$\mathbb{E}\{h_{km}(\theta)U'(P_m - C(\theta))\varepsilon_m - h_{kM}(\theta)U'(P_M - C(\theta))\varepsilon_M\}$$
$$< \mathbb{E}\{h_{k+1m}(\theta)U'(P_m - C(\theta))\varepsilon_m - h_{k+1M}(\theta)U'(P_M - C(\theta))\varepsilon_M\},$$

for every $k = 1,\dots,K - 1$. Thus, by an apppropriate choice of ε_m and ε_M we can ensure that type k^C prefers P′ and all types $k < k^C$ prefer P^C.

Case (b). Now suppose that $k > k^C$ is indifferent between P^C and some other contract. Without loss of generality, we can choose k′ to be the largest such type. By the previous argument, with k′ in the place of k^C, we can show that there is a contract P′ close to P^C that is strictly preferred to P^C by the purchaser and strictly preferred to P^C by k′ but not by $k \le k^C$. This again leads to a contradiction.∎

Separating equilibria

Although there are no stable equilibria that involve pooling at contingent contracts, there will exist some stable equilibria in which contingent contracts are chosen. In these equilibria, the choice of a contingent contract reveals the supplier's type. The amount of separation in these equilibria varies. What we can show is that, even in a separating equilibrium, at least one type chooses a non-contingent contract and possibly several types pool at a non-contingent contract.

PROPOSITION 4: Suppose that an arbitrary equilibrium satisfies the Intuitive Criterion. Then at least one type of supplier (the "best" type) must choose a non-contingent contract.

PROOF: Once again, the proof is by contradiction. Suppose that the highest type $k = K$ has chosen a contingent contract P^C in equilibrium. Let P^{NC} denote the best non-contingent contract that is not strictly preferred by type K to his equilibrium contract. Then every type $k < K$ must strictly prefer his equilibrium contract to P^{NC}. This follows since each type $k < K$ weakly prefers his equilibrium contract to P^C and if K weakly prefers P^{NC} to P^C, then every type $k < K$ must strictly prefer P^{NC} to P^C.

Since there is no pooling at contingent contracts, the worst type 1 reveals his type by proposing the contract P^1 in equilibrium. The assumptions of Proposition 1 tell us that type 1 is worse off at P^1 than at the non-contingent contract P^* chosen in the pooling equilibrium. Then, P^{NC} must be strictly worse than P^* for the supplier and therefore strictly better than nothing for the purchaser. Hence, the contract P^{NC} must be accepted by the purchaser if he believes it is offered by type K, it will never be offered by types $k < K$, whatever the response of the purchaser, and it is at least as good as the equilibrium contract P^C for the supplier of type K, if it is accepted. We have not shown that the supplier will be strictly better off at P^{NC}, but a small increase in the contract price will make type K strictly better off without changing any of the other conditions. Thus, the Intuitive Criterion is violated, a contradiction.∎

PROPOSITION 5: Suppose that an equilibrium satisfies the Intuitive Criterion and the conditions of Proposition 1 are satisfied. In the given equilibrium every type of supplier is worse off than he would be in the pooling equilibrium described in Proposition 1.

PROOF: From Propositions 3 and 4 we know that, in the given equilibrium, at least one type of supplier chooses a non-contingent contract P^{NC} and there is no pooling at contingent contracts. In particular, the worst type, k = 1, is revealed and must get the best contract that is acceptable to the purchaser, conditional on his type. That is, the supplier of type 1 gets a payoff of $\bar{U}_1$.

In the pooling equilibrium described in Proposition 1, every type of supplier receives the same payoff, call it U^*. By hypothesis, $U^* > \bar{U}_1$. Therefore at least one type is better off in the pooling equilibrium. To see that the other types are also better off in the pooling equilibrium, note first that the non-contingent contract P^{NC} must be strictly worse P^*; otherwise, the worst type, k = 1, would choose it. So any type that chooses P^{NC} is strictly worse off than in the pooling equilibrium. Second, suppose there is some type k > 1 who chooses a contingent contract P^k in the given equilibrium and that, contrary to what is to be proved, type k weakly prefers P^k to P^*. Then by a familiar argument, the worst type must strictly prefer P^k to P^*. But since the equilibrium conditions require that P^1 is weakly preferred to P^k by type 1, P^1 must be strictly preferred to P^* by type 1, contradicting our hypothesis. This completes the proof.█

It would be nice to obtain results that characterize the non-pooling equilibria in more detail or that say more about the welfare properties of the equilibria that satisfy the Intuitive Criterion. For example, are they

Pareto-ranked from the suppliers' point of view? Proving such results is difficult because the multi-dimensional character of the problem means that there is very little evident structure. The following result is all that I have found.

PROPOSITION 6: Suppose that an equilibrium satisfies the Intuitive Criterion. Then the set of types pooling at a non-contingent contract is an interval of the form $\{k^{NC},\ldots,K\}$. Suppose that I^k is the non-contingent contract that makes type k indifferent with his equilibrium contract. Then if I^k lies strictly within the acceptance region for type k, i.e., $V^*(I^k,k) > 0$, the type k supplier chooses the non-contingent contract in equilibrium, i.e., $k \geq k^{NC}$.

PROOF: The first claim follows from the fact that if type k weakly prefers a non-contingent contract P^{NC} to a contingent contract P^C, then any type $k' > k$ strictly prefers P^{NC} to P^C.

To establish the second claim, suppose that type k chooses a contingent contract P^k in equilibrium and suppose, contrary to what we want to prove, that $V^*(I^k,k) > 0$. By a previous argument, every type $k' < k$ will strictly prefer P^k to I^k and so strictly prefer their equilibrium contracts to I^k, even if we improve it a little. Hence, there exists a contract P′ that is strictly preferred to P^k by k, strictly inferior to the equilibrium contract for every $k' < k$ and satisfies $V^*(P',k') > 0$ for every $k' \geq k$. Then the Intuitive Criterion is violated.∎

As a corollary of the second claim, if $V^*(I_1,k) \geq 0$ for some $k > 1$, then type k chooses the non-contingent contract in equilibrium. Thus, if type 1 is "bad" enough relative to the rest, all types $k > 1$ will pool at the non-contingent contract.

V. MORE GENERAL PREFERENCES

One of the assumptions used in the model of Sections III and IV is that different types of supplier have the same preferences for consumption. More precisely, the additively separable utility function $U(c) + T_k(d)$ assumes that the utility of consumption is independent of k. We make essential use of this assumption in the analysis, where it ensures that the indifference curves of different types of supplier are tangent at a non-contingent contract. In order to extend the analysis to situations where this condition does not hold, we allow suppliers to differ in their costs but assume that types occur in matched pairs. For each distribution of production costs there are (at least) two types, one with high distortion costs producing a high value product and one with low distortion costs producing a low value product. To illustrate, we start with the two-type example in Section 3. A natural way to extend this example is to use the utility function to $U(c) + T_k(d)$ $(k = B,G)$ but assume that the distribution of θ is different for different values of k. But this will not work, since it clearly makes the pooling equilibrium unstable under standard conditions. Instead, we assume there are four types, $k = G_1$, G_2, B_1 and B_2, with preferences given by

$$U_k(c,d) = \begin{cases} U(c) + T_G(d) \text{ if } k = G_1,G_2 \\ U(c) + T_B(d) \text{ if } k = B_1,B_2 \end{cases}$$

and costs given by

$$C_k(\theta) = \begin{cases} C(\theta) & \text{if } k = G_1,B_1 \\ C(\theta + \psi) & \text{if } k = G_2,B_2 \end{cases}$$

for some $\psi > 0$. In other words, for each distribution of costs there is a "good" type and a "bad" type. This pair of types will function like the pair in the original example of Section 3. However, between the two pairs there is a difference in costs. The difference in costs generates an additional source of differences in the types' preferences over contracts, which may serve as a additional source of signalling.

The pooling equilibrium will continue to have the stability properties described in Section III. The reason is that each of the good types is "shadowed" by its own bad type: any contract that is weakly preferred by a good type is strictly preferred by its corresponding bad type. Thus, Universal Divinity continues to hold. This kind of equilibrium is illustrated in Figure 3(a).

At contracts off the diagonal, however, things are different. It is now possible to have pooling, as illustrated in Figure 3(b). The reason is that the two bad types may prevent one of the good types from signalling. Although the slopes of the indifference curves are different, the fact that the lower envelope of the two bad types' indifference curves lies below the good type's indifference curve is sufficient to make this equilibrium stable. But this can only occur if the contract is not "too far" from the diagonal.

Note that this equilibrium has a single non-contingent contract. This is a robust feature. Any stable equilibrium must have at least one type choosing a non-contingent contract. Consider an equilibrium in which no type chooses a non-contingent contract. Whichever side of the diagonal a good type finds himself on, his indifference curve tilts toward the diagonal more than the corresponding bad type. Then at least one of the good types must have an indifference curve the tilts further toward the diagonal than any other type. This property is all that is needed to make the equilibrium unstable. The distinguished type can increase his utility by moving toward the diagonal, without attracting any other type.

REFERENCES

Antl, B., ed., (1986) *Swap Finance*. London: Euromoney Publications Ltd.

Bester, H. (1985) "Screening vs. Rationing in Credit Markets with Imperfect Information," *American Economic Review* **57**, 850-855.

Cho, I.-K. and J. Sobel (1990) "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory* **50**, 381-414.

Cho, I.-K. and D. Kreps (1987) "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics* **102**, 179-222.

De, S. and J. Kale (1990) "Security Design and Equilibrium Refinements: Implications for Debt Packaging," GSB, University of Wisconsin-Madison (mimeo).

De Meza, D. and D. Webb (1987) "Too Much Investment: a Problem of Asymmetric Information," *Quarterly Journal of Economics* **102**, 281-292.

Dewing, A. (1911) "The Position of Income Bonds, as Illustrated by Those of the Central of Georgia Railway," *Quarterly Journal of Economics* **15**, 396.

Dewing, A. (1955) *The Financial Policy of Corporations* (Fifth Edition). New York: The Ronald Press Company.

Diamond, D. (1984) "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies* **51**, 393-414.

Gale, D. and M. Hellwig (1985) "Incentive-Compatible Debt Contracts: The One-Period Case," *Review of Economics Studies* **52**, 647-663.

Hart, O. and B. Holmstrom (1987) "Theory of Contracts," in Truman Bewley (ed.) *Advances in Economic Theory*, Econometric Society Monographs, Fifth World Congress. New York: Cambridge University Press.

Holmstrom, B. and P. Milgrom (1990) "Multi-Task Principal Agent Problems: Incentive Contracts, Asset Ownership and Job Design," SOM, Yale University (mimeo).

Holmstrom, B. (1979) "Moral Hazard and Observability," *Bell Journal of Economics* 10, 74-91.

McConnell, J. and G. Schlarbaum (1987) "The Income Bond Puzzle," in J. Stern and D. Chen (eds.) *The Revolution in Corporate Finance*. New York: Basil Blackwell.

Nachman, D. and T. Noe (1989) "Design of Securities under Asymmetric Information," Georgia Institute of Technology (mimeo).

Shavell, S. (1979) "Risk Sharing and Incentives in the Principal Agent Relationship," *Bell Journal of Economics* 10, 55-73.

Scherer. F. (1964) *The Weapons Acquisition Process: Economic Incentives*. Cambridge MA: GSBA, Harvard University.

Spier, C. (1989) "Incomplete Contracts in a Model with Adverse Selection and Exogenous Costs of Enforcement," M.I.T. (mimeo).

Stiglitz, J. and A. Weiss (1981) "Credit Rationing in Markets with Imperfect Information," *American Economic Review* 71, 393-410.

Townsend, R. (1979) "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory* 22, 265-293.

NOTES

1 This is not the only game form that could be used. One could assume that the distortion decision is made before the cost shock is known. The results would be much the same, but it would be necessary to assume that $C'(0) = 0$ in order to ensure that some distortion is always optimal in contingent contracts. In general, information arrives and distortion decisions are made over time and one would need a much more complex dynamic game to capture this time structure.

2 One way of dealing with the discontinuity is to define an equilibrium distortion to be optimal for a supplier if it is the limit of distortions that yield maximum utility asymptotically. Alternatively, one could treat the signal S as a correspondence that is bi-valued at 0. These all amount to the same thing and add nothing to the theory.

3 One qualification must be noted. Because the distortion decision is made ex post, after the cost shock has been observed, even if the costs of distortion are large there will always be some values of θ for which distortion is optimal. If the distortion decision were made ex ante, before the observation of the cost shock, it might be optimal to choose no distortion because of the costs. In that case, one would obtain a degenerate case in which there would effectively be no role for distortion. This last possibility is avoided, even in the case of an ex ante decision, by assuming that $C'(0) = 0$.

4 The crucial assumption is that the signal produced by the measurement system takes on a finite number of values. From this is follows that a contract can be described by a finite number of parameters so the supplier's strategy set and the domain of the purchaser's strategies are both contained in a finite dimensional space. Obviously, the extension to an infinite dimensional contract space would pose formidable technical difficulties. Even under the assumption of risk neutrality, the analysis of infinite dimensional contracts is not easy. Cf. Nachman-Noe (1989).

5 The maximum payoff the supplier can obtain from a contract P is obviously well defined, whether or not his optimal distortion is uniquely determined by (θ,P). On the other hand, the payoff received by the purchaser may not be. If the supplier is indifferent between two different actions which have different payoffs for the purchaser, we cannot represent the payoff to the purchaser as a function of the contract. Even if we were to make an arbitrary selection of the supplier's optimal action, the payoff would not be a continuous function of the contract. To reduce the game to a well behaved signalling game, it is crucial that the supplier's optimal distortion be unique.

6 The argument is somewhat more involved than this suggests: we use the fact that there is no pooling at contingent contracts. See Proposition 3.

7 We write $h' > h$ if $h'(\theta) \geq h(\theta)$ for all θ and the inequality is strict on a set of positive measure.

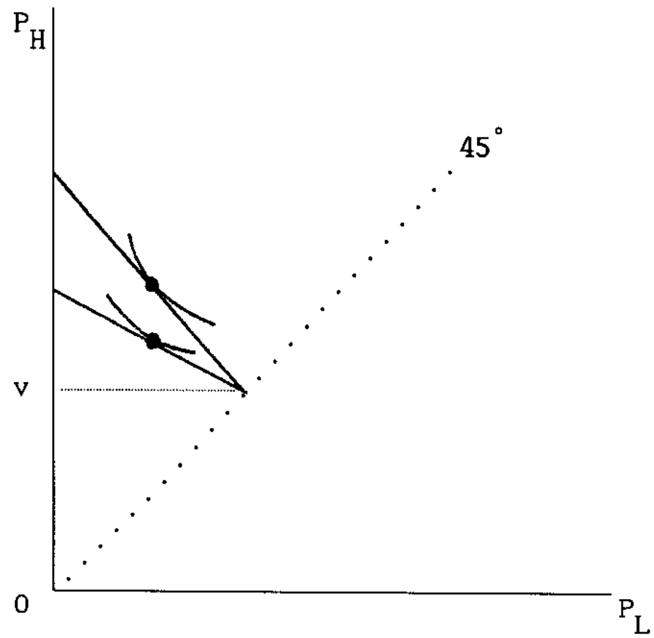

Figure 1(a) Complete Information Equilibrium

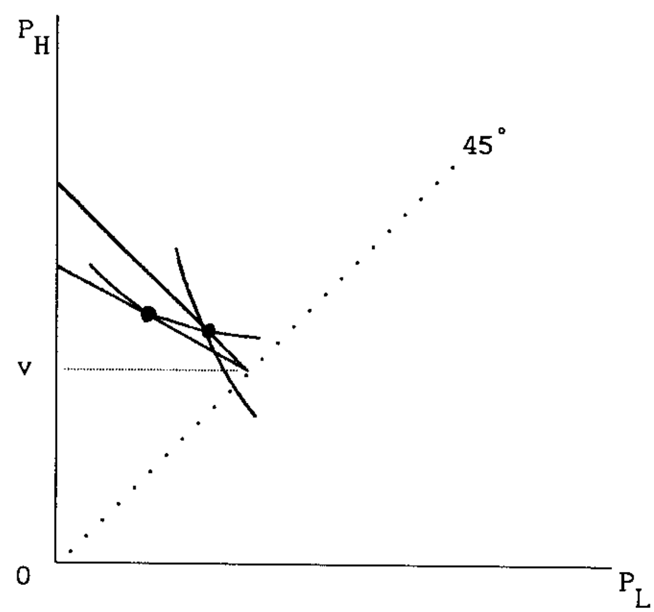

Figure 1(b) Incomplete Informatio Equilibrium

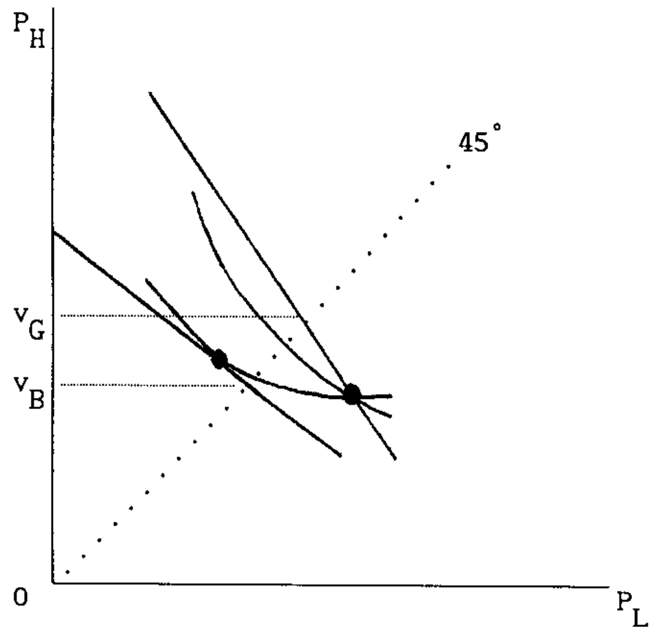

Figure 2 Incomplete Information Equilibrium when Different Types Produces Goods of Different Value

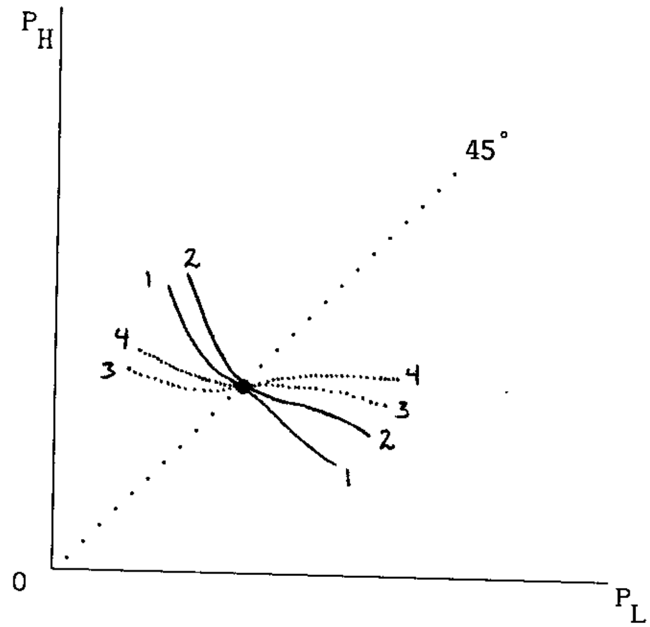

Figure 3(a) Stable Equilibrium with Pooling at a Non-Contingent Contract

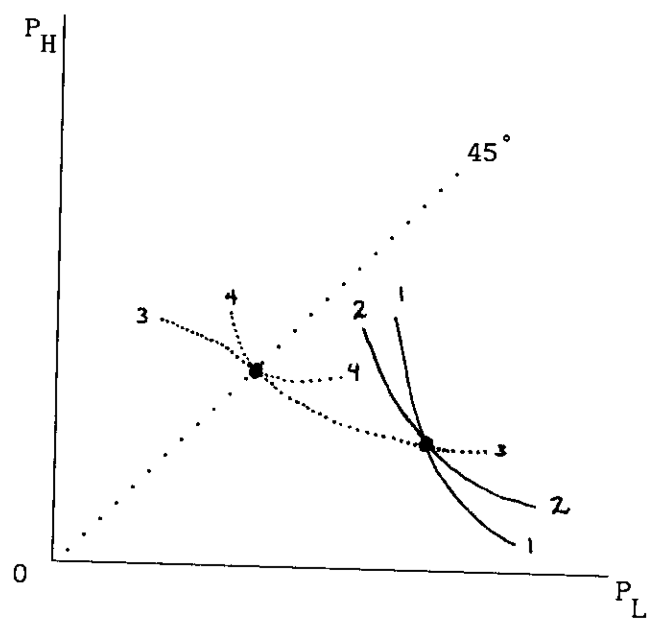

Figure 3(a) Stable Equilibrium with Pooling at a Contingent Contract