MORAL HAZARD

Moral hazard has somewhat different meanings in insurance industry parlance and in micro economics. In either case, the problem arises because the insured typically has a degree of control over the probability that an insured event will occur, or the size of the loss if it does arise. The driver can drive slower, or with more care. The property owner can install smoke alarms and can make sure inflammable wastes are cleared up. The firm buying worker's compensation insurance can choose to educate its workers in safety techniques. The same firm can minimize worker's compensation claims after the fact by occupational therapy. The insured who has suffered a fire can often minimize the loss by removing and protecting undamaged goods. I will refer to such actions, or inactions, of the insured as the chosen "level of care" or the chosen "investment in safety".

In insurance, the term refers to the tendency for those with insurance to relax their level of care, or their investment in safety and loss prevention. It is sometimes used with ethical overtones. Some people are given to take advantage of their insurance protection by failing to take care. Others, of more upright character, behave impeccably (from the insurer's viewpoint) despite their insurance coverage. Thus, moral hazard describes a behavioral characteristic of the insured which is used in the underwriting decision.

In micro economics, the focus is on information, incentives and rational behavior. Insurance, and other hedges do relax incentives to take care or invest in safety. The expected loss is not fixed. Rather, the expected loss can be expected to increase as a result of the insurance protection. This would present little problem to the insurer, if the behavior of the insured could be monitored perfectly; the insurer could simply make coverage, or price, conditional on the level of care or safety chosen by the insured. The problem arises when the level of care is hidden to the insurer. Since pre-arranged contract conditions, nor a pre-arranged price, can be made conditional on the subsequent behavior of the insured, the latter may be tempted to choose a low investment in safety once the insurance is in place. It would seem that the best the insurer can now do is to anticipate that the insured will reduce his chosen level of safety and take this into account in the price. This will make insurance more expensive since the insured will be required to pre-pay for the anticipated increase in expected loss. The insured, is hoisted on his own petard.

The moral hazard problem is not unique to insurance. It is an illustration of a wider set of *principal - agent* problems. A principle agent problem arises when one party, the principal, employs another, the agent, to perform a task on the principal's behalf. Unless the principal can monitor all actions of the agent (in which case doesn't the principal perform the task herself?), the agent may be tempted to act in his own interest. For example, a lawyer acting on behalf of the client may favor the publicity and challenge of a trial whereas the plaintiff, the agent, could be better off with a settlement. Another common example, is the relationship between shareholders and managers of a firm. The managers are there to act in the interest of shareholders which implies maximizing the value of the shareholders' investment. However, the manager might be tempted to work with less vigor, to consume perquisites, to choose project which minimize the risk of job loss, and to engage in relationships which maximize the manager's marketability for a new job. The more difficult it is to monitor these actions, the more likely that managers will undertake such actions. The problem is that one party chooses a set of actions, but another party gets the rewards of those actions, or picks up the pieces. One party chooses the meal, but someone else pays the bill. There is a separation of ownership and control. Similarly in insurance, the insurer pays for the protection, but the insurer pays the losses if the protection fails.

Moral hazard can be divided into *ex ante* and *ex post.* If wealth is hedged against some event, then the party buying the hedge will have little incentive to undertake actions *before the event* that will reduce the expected cost, nor will he have incentive to take actions *after the event* that can minimize its impact, and therefore the payout under the hedging instrument. For example, the firm may be little inclined to spend a lot on product safety if it is insured against liability; nor will it have a great incentive to deny liability if considers that a legal fight would have an adverse effect on product demand and the insurer was paying for the settlement anyway. These are *ex ante* moral hazard and *ex post* moral hazard. We will first look at *ex ante* moral hazard.

EX ANTE **MORAL HAZARD**

Optimal Safety for the Insured.

Insurance gives rise to a classic principle agent problem. The insured chooses safety and thereby determines the expected loss; but the insurer pays the loss. To see the basis of the insurance problem we will first consider that an individual must choose some level of care or safety. For example, the person must choose how much to invest in sprinklers, smoke detectors, fire doors and other forms of fire safety. The insured faces the cost, c(s) which says that the cost is a function of the level of safety chosen. One would expect this cost to increase the higher the chosen level of safety, i.e., c is an increasing function of s. On the other hand, the individual receives a benefit. Since safety reduces the expected value of loss (either by reducing the probability or the size), then the expected cost of fires will fall. Let us consider the effect is a fall in probability. Thus the probability of a loss of size L, depends on the chosen level of s, i.e., probability is p(s) where p falls as s increases. Consider first that the person is indifferent to risk. The level of safety which will maximize expected utility is that which equates the marginal cost and marginal benefit of safety. In Figure ***, the rising curve is the marginal cost of safety, c(s). The marginal benefit, assuming no insurance, is shown as the steepest falling line denoted "no hedge". The optimal choice is where the cost of any incremental improvement (marginal cost) just matches the incremental reduction in expected losses (marginal benefit). This is shown as position A.

To see what is going on here, let us first consider that the person's expected wealth, W, is some intial level, W_0, minus the cost of safety, c(s), minus the cost of paying for accidents, p(s)L:

$$W \quad = \quad W_0 - c(s) - \quad \Sigma\, p_i(s)\, L$$

Let us simplify a little to start and assume the person is indifferent to risk. The task is to choose the level of s that maximizes expected wealth. This can be done by taking the derivative of wealth with respect to s, and setting this derivative to zero to derive the maximum.

$$\frac{\partial W}{\partial s} \quad = \quad - \frac{\partial c}{\partial s} \quad - \quad \sum \frac{\partial p_i}{\partial s} L \quad = \quad 0$$

This leads to the straightforward interpretation, marginal costs, $\partial c/\partial s$, equals marginal benefit, $-\sum(\partial p_i /\partial s)L$.

If the person is averse to risk (why after all would she wish to buy insurance?), the analysis can be adapted by looking at the utility of wealth, U(W), and taking the derivative with respect to s. Formally this will look as follows:

$$U(W) \quad = \quad \Sigma\, p_i(s)\, U(\, W_0 - c(s) - p_i(s)L_i\,)$$

and the derivative turns out to be a little messier and you can ignore this step and move straight to the interpretation at the end of the paragraph.

$$\frac{\partial U(W_i)}{\partial s} \quad = \quad - \Sigma_i \left(\frac{\partial U}{\partial W_i} \left[\frac{\partial c}{\partial s} \quad - \quad \frac{\partial p_i}{\partial s} L \right] \quad + \quad \frac{\partial p_i}{\partial s} U(W_i) \right) \quad = \quad 0$$

This is still equating marginal costs and marginal benefits of safety, but now everything is stated in utility terms. The term involving $\partial c/\partial s$ is the marginal cost and all the terms involving $\partial p_i/\partial s$ comprise the marginal benefit. The marginal cost is the utility loss from paying for safety. And the marginal benefit is the utility gain from reducing the probability of paying any potential future loss. While the actual level of safety chosen will be depend on the shape of the utility function (and will not generally be the same as the case of risk neutrality), the risk averse person will still be willing to invest more in safety, the greater the saving in potential future losses.

Now suppose that insurance is purchased. The marginal benefit is zero if the loss is fully insured; the insurer pays for all damage. Thus, the marginal benefit to the insured is zero, regardless of the level of safety. This is shown as the line that coincides with the horizontal axis, labeled "100% hedge". The point where marginal cost and marginal benefit are now zero is where s=0. The interpretation is straightforward. The insured pays the costs of safety but gets none of the benefit. Consequently, the insured invests nothing in safety. An intermediate case is also shown where 50% of the loss is insured. Here the insured pays all the costs but gets half the benefit (the full benefit being shared with the insurer)[1]. The chosen level of safety is shown at the intermediate level "B".

The problem with insurance can be shown as follows. Some proportion of the loss, α, is insured and the remaining portion, $1-\alpha$, is paid by the policyholder. For this insurance there is a premium, αP which is scaled according to the proportion of the risk insured. Thus, wealth is now the starting level, W_0, minus the cost of safety, c(s), minus that portion of the expected loss that is not insured, $(1-\alpha)p(s)L$, minus the premium, αP:

$$W = W_0 - c(s) - (1-\alpha)p(s)L - \alpha P$$

The task is still to choose the level of s that maximizes expected wealth. Setting this derivative to zero to derive the maximum:[2]

$$\frac{\partial W}{\partial s} = -\frac{\partial c}{\partial s} - (1-\alpha)\ \Sigma\frac{\partial p_i}{\partial s} L = 0$$

Now the marginal cost is set equal to $(1-\alpha)$ times the marginal benefit. The effect is to scale down the level of safety chosen by the policyholder.[3]

The Optimal Insurance Contract, Given that "s" is Not Observed by the Insurer.

Now consider the problem from the insurer's point of view. The insurer has to design and sell an insurance policy knowing that the policyholder will select s (or at least can change s) after the policy has been sold. One way the insurer can go is to scale back on the

---

[1]The utility benefit may not be exactly half that for no insurance. However, the marginal utility benefit will lie between the full insurance and no insurance cases.

[2]Or, with utility:

$$\frac{\partial U(W_i)}{\partial s} = -\Sigma_i \left(\frac{\partial U}{\partial W_i} \left[\frac{\partial c}{\partial s} - (1-\alpha)\frac{\partial p_i}{\partial s} L \right] + \frac{\partial p_i}{\partial s} U(W_i) \right) = 0$$

[3]Notice that the insurer could not monitor the chosen level of s. If this was not the case, the insurer could have set the premium relative to s and there would have been no problem.

level of coverage. This will at least give the policyholder some stake in future losses and will encourage him to take an interest in preventing or containing them.

Suppose that the expected loss were independent of the level of coverage; there was no moral hazard issue. Now imagine that the insured can choose some proportion, α, of the loss to insure. Apart from any fixed transaction costs, it seems plausible that the premium would be proportional to α. Figure *** shows such a premium schedule as the straight upwards sloping line labeled “price w/o moral hazard”. Since the vertical axis measures price (which is undesirable to the insured) and the, and the horizontal axis shows coverage (which is desirable), the policyholder’s utility will be higher the further to the southeast. Thus, indifference curves for the policyholder are show with I(1) having higher utility than I(2) since it is further to the southeast. The individual chooses the level of coverage which attains the highest indifference curve, I(1). This coverage is shown as “A” in the diagram.

With moral hazard, the policyholder’s chosen level of “s” will likely decline, the more coverage is purchased. Thus, the expected loss, p(s)L will increase as more coverage is purchased. This means that the price charged by the insurer, just to breakeven will not be proportional to α, but will increase in greater proportion than α. After all, more insurance means that the policyholder is getting a larger proportion of an increasing expected loss. This double whammy makes the premium rise steeply to cover the increasing moral hazard effect. The new price line is that marked “”price w/ moral hazard”. With the increased cost from the moral hazard, the original level of utility, I(1) is no longer attainable. The highest utility that can be achieved by the policyholder is shown by the indifference curve I(2). The result shows that less insurance is purchased.[4]

The optimal level of insurance can be solved with moral hazard along the lines shown in the previous paragraphs. But that is not necessarily the optimal insurance contract that can be arranged between the policyholder and insurer. Solving the optimal *level* of insurance assumes the structure of contract is fixed. For example, one can assume the form of the insurance coverage is such that a proportion of the loss is covered by the insurer and the choice is “what proportion?”. Or one could assume that the structure is one with a deductible and the choice is “what level of deductible?”. There are many other potential structures for insurance coverage. The question are, “what is the optimal structure?” and, given that structure, “what are the optimal parameters?”.

While the formal solution of this problem is a little messy, the solution is quite accessible. Consider another agency problem with similar structure. The board of directors

---

[4]In these models, it is usually the case that moral hazard reduces the level of insurance that the individual would choose. However, a formal proof does require certain assumptions. Winter (   ) has a nice review of the economics of moral hazard.

of a firm must decide how to structure the compensation of the CEO. The issue is not simply the level, but on what factors will compensation depend? There are two sets of issues that clearly are relevant. One is risk sharing. Shareholders can diversify risk at low cost in how they form portfolios of assets. Managers cannot diversify the risk surrounding their employment since it typically represents a large portion of their wealth. Accordingly, it would appear to make sense that any risk attached to the firm's earnings or value, is best allocated to shareholders who have a comparative advantage in bearing that risk. If one were to require the manager to take part of that performance risk (say by rewarding them with profit related bonuses) the manager would require a risk premium. This can be avoided by paying flat compensation (unrelated to profit).

Now the other factor is incentive. Managers, as argued before, have an incentive to work more on their own behalf if they cannot be monitored. Accordingly, one would like to make compensation related to the quantity and quality of effort exercised on behalf of the firm's owners; the greater the effort, the higher the pay. In this way, the interests of the managers would be aligned with those of the shareholders. Accordingly, such contracts are often referred to as *incentive compatible*. But managers cannot observe effort; therefore they cannot make a feasible contract conditional on effort. If they tried to do so, they would have to rely on the manager's own statements on effort and, since these could not be independently verified, managers could not be penalized for lying. The answer relies on there being a relationship between effort and something the managers can observe. If managerial inputs cannot be observed, the fruits of those inputs can. Managers can observe the price at which shares are traded. They can, and do observe periodic earnings. Even though profit reports rely on information provided by management, this is independently audited. So, if pay is related to such value or profit measures, managers will have an incentive to take those actions (supply more effort) that are likely to increase profit.

Notice that the effectiveness of this compensation scheme in increasing shareholder wealth, does depend on there being a positive relationship between effort and share value or earnings. But it does not rely on this relationship being deterministic. It may be objected that earnings are not determined solely by managerial decisions and effort. The best laid plans can go astray. Thus, profits will be influenced by factors outside the manager's control, such as the state of the economy, interest rates, changes in exchange rates, or technical problems of production that could not have been foreseen. Thus, profits will partly reflect managerial talent and effort and partly reflect chance. It is sometimes said that there is "signal and noise". If managers are given incentive compensation, then there is a problem they could be rewarded, or punished, due to factors outside their control. This is simply a problem of being unable to monitor effort. The point is that, prospectively, managers will tend to supply greater effort is it increases the *likelihood* that there bonuses will increase.

With these two ideas in conflict (optimal risk sharing would favor flat compensation whereas incentive compatibility favors output related compensation) we have a trade off. As in many economic dilemmas, there is a trade off between risk sharing and efficiency. The optimal arrangement is often one which has a little of both

Now consider insurance. Insurers wish policyholders to take those actions that reduce expected loss, but the actions cannot be observed. But actual losses are observed; indeed the insurer has to pay them. And there is a stochastic relationship between care/safety and losses (just as there is a stochastic relationship between managerial inputs and profit). Thus, the insurance contract will have its price related to actual loss experience. This does court the risk that the insured can be penalized for bad luck. For example, an automobile policy might increase premiums after two accidents within a space of three years. There is a chance that even the best driver could have had two accidents simply due to ill fortune. But, if you do not know whether the driver is good or bad, the fact of two accidents may appear to lend more support to the view that the driver is bad (and the record is representative of a bad driver) than that the driver is really good (but that the experience is not representative). There is scene in Oscar Wilde's "The Importance of Being Ernest", where a young man is trying to explain that he is an orphan, only to be told that losing one parent is bad luck but losing two is somewhat careless.

There are two ways in which premiums can be related to losses, one is an experience related premium and the other is a retrospective premium. Experience rated premiums are calculated with reference to the loss experience of prior periods. If I know my premiums in the future are to be experience rated, and I anticipate buying insurance in the future, then I will encourage to take care even though the insurer will pay for losses. The second method is to treat each year's premium as conditional on the losses that occur in the year. Usually, this is done in two stages. A provisional premium is paid up front. At the end of the insurance period (often with some delay since claims take some time to evaluate) the provisional premium will be adjusted to reflect the loss experience. At the extreme, the adjustment could be the difference between the provisional premium and the actual losses. But this would undo the insurance protection. The insured would end up paying for his own losses and would receive effectively, a short term loan to pay for those losses. Short of this extreme, the adjustment would normally be some portion of the difference between the provisional premium and actual losses. Thus, there is part risk sharing, and part an incentive compatible contract.

ADVERSE SELECTION

One of Groucho Marx's often attributed remarks is that he would not like to belong to any club that would have him as a member. Marx captured the essence of adverse selection. Adverse selection can arise in any market place where people trade with different information. If one party has better information about some relevant features of the trade, then he may use that knowledge to his advantage and, often to the disadvantage of the trading partner. The way the partner reacts to this potential disadvantage can act as a barrier to trading and can lead to market inefficiencies. While moral hazard has been described as a problem of hidden actions, adverse selection has been defined as a problem of hidden information; one party has access to material information that is denied to the other. It is a problem of information asymmetry.

An appealing illustration of this concept was used some years ago by George Ackeloff and, since this was the market for second hand cars, adverse selection is often referred to as a "lemons" problem. The sellers of secondhand cars have better information about their quality than the buyers. They know of any prior accidents, they know the service and repair record, and they know of any chronic problems with the vehicle. The buyer may be able to uncover some, but not all, of this information. Thus, there is an information asymmetry with (in the language of political correctness) the seller being informationally advantaged and the buyer informationally challenged. Now cars coming onto the secondhand market will vary in quality, some being better than average and some worse than average. The seller will know the quality but the best the buyer will know is the average quality of cars of this type. Not being able to verify individual quality, the best the buyer will be willing to offer is the price suitable for a car of average quality. If the seller knows his car to be better than average, this price is unlikely to be adequate and he may well withdraw his vehicle from the market. But for the seller with a worse than average car, a price based on average quality is more than it is worth and she will gladly sell. This trade in this market will comprise mostly worse than average cars.

But the problem gets worse. Buyers can now anticipate that mostly poor quality cars will be offered for sale and this will lower further the price they are willing to pay. This will make any remaining sellers with high quality cars even less willing to sell. And this dynamic process continues we will be left with few, if any, good second hand cars for sale. Good cars are "driven out" of the market.

Ah ! but why does the seller of a high quality car simply say, "my car is better than average and therefore should command a high price". The problem is that, since the buyer cannot verify this statement, then the sellers of lemons will also straighten their faces and say that there cars also are wonderful and so buyers will be denied meaningful information. Cheap talk like this is simply not convincing. We will see later that there may be mechanisms to separate the good and bad cars, but they will need to be more subtle. For these mechanisms to work, the seller of the high quality vehicle must be able to send some signal of quality, which will not be replicated by the seller of the low quality car.

To see how adverse selection can arise in insurance markets consider the case of automobile insurance. To keep things simple, drivers can be grouped into two classes: those with a high expected value of loss and those with low loss expectancy. If the insurer can distinguish drivers according to their respective loss characteristics, each can be charged a premium that reflects his or her expected value of loss. Thus, insurers may use observable characteristics, such as automobile type, location, or age, to distinguish different risk groups of automobile policies. In this way, the insurer will have many categories of policyholders; each category containing drivers that are similar in observable features. But, even after classifying in this way, there may still be considerable variation within each class; not all 22 year old males driving sedans in Philadelphia have the same loss potential. The skill levels and behavioral characteristics can vary substantially. So there will be an effective subsidy from low-risk drivers to high-risk drivers within each class. This subsidy can destabilize the insurance market and reduce insurance supply, as shown below.

In Figure *** we consider a category of drivers who are similarly classified on the basis of observable characteristics.. However, there is hidden variation in loss potential; some being worse than average and others better than average. We will simply cal these high and low risk drivers. Each insured starts with a wealth level of $125, but a loss can reduce the wealth to $25. The groups differ in the probability of loss. For the high-risk group, the probability of loss is 0.75, resulting in an expected loss of $75. For the low-risk group, the probability loss is 0.25, resulting in an expected loss of $25. If the insurer can distinguish between the two groups, respective competitive premiums of $75 and $25 can be charged (ignoring transaction costs). With premiums set at the expected value of loss for each insured, the Bernoulli principle asserts that each would fully insure. For the low risk group, the utility of insuring and having wealth of $100 with certainty, (that is, U($100)), is higher than the expected utility of not insuring EU_L. Thus,

$$U(\$100) > EU_L = (0.75)U(\$125) + (0.25)U(\$25)$$

and, for the high-risk group,

$$U(\$50) > EU_H = (0.25)U(\$125) + (0.75)U(\$25)$$

The respective positions are shown on the vertical axis of Figure ***.

Now suppose that the insurer is unable to distinguish between high and low-risk drivers. If there are equal numbers in each group, the break-even premium will be $50. However, at this premium the low risk group will not insure because the utility of not insuring, EU_L , is greater than the utility of insuring and having a wealth level of $75 for certain, that is:

$$EU_L > U(\$75)$$

Conversely, the high-risk group will find insurance to be a bargain and will choose to insure, that is,

$$EU_H < U(\$75)$$

Consequently, the portfolio composition will change as low risk drivers cancel their policies, leaving a portfolio of high-risk drivers (each having an expected cost of $75) and an inadequate premium.

In practice, the process will be somewhat smoother. There may be several risk groups, and coverage may be arranged on a partial basis. The insurer that averages premiums over a number of risk groups will find that it tends to lose the good risks as they cancel or reduce their coverage. The resulting change in the composition of the portfolio will cause the average premium to be inadequate, forcing the insurer to raise the premium. This aggravates the flight of low-risk drivers, and so the process continues until only bad risks are left. This process is usually attributed to information deficiency on the part of the insurer, but the same effect can result from regulation designed to prevent insurers from using classification variables that are politically sensitive, such as sex or race. (see Dahlby, 1983, Crocker and Snow, 1986 and Hoy 1989).

Competition between insurers may help reduce problems of adverse selection. Information on loss expectancies of individual drivers is of economic value to an insurer. Armed with such information, an insurer can selectively attract low risk drivers from a rival that is unable to discriminate simply by offering a lower price and admitting only low-risk drivers. Thus, competition will induce insurers to seek and compile information that will enable them to use premium structures that discriminate between risk groups. Of course, information will never be perfect, and adverse selection will never disappear. But in an actively competitive market, adverse selection will be reduced to a level that reflects the cost of information.

The Rothschild Stiglitz Model

While competition may stimulate insurers to use information that is observable to classify insureds, there will always be some information asymmetry between insured and insurers. Insofar as insured are better informed, the problems of adverse selection remain. But there are ways in which insureds can be induced to reveal private information in a way which is credible. The principle works as follows. Suppose I offer a pair of contract to an individual without knowing whether that person is high or low risk. The person receiving this offer however does know their risk type. Now the contracts are designed such that one contract would only be appealing to those who knew they were low risk; and the other contract would be appealing to those who knew themselves to be high risk. If the person receiving such an offer will select according to their private information about their loss type. Accordingly, if he or she chooses the contract that would only be appealing to high (low) risks, I can safely conclude that he/she is indeed a high (low) risk. By making their contract choice, the person reveals their hidden identity. It works because of the clever design of the two contracts. This is essentially the Rothschild Stiglitz model. This concept is known as *self selection*. By their own choices, the different types will separate into their respective groups.

Consider Figure 4. The axes show the wealth of the individual under two different circumstances; first if she does not have a loss and second if she has a loss. A "loss" in this context means some event that would affect their net worth. The 45° line traces out equal values on both axes. On this line, the person's wealth is unaffected by whether a loss occurs, it can only mean that they have made some arrangement to protect themselves from the financial impact of the event. In other words, they are fully insured, or fully hedged. Call the 45° line the "full insurance line". The point A is the starting position. This is the wealth combination if no insurance is purchased so that, at A, wealth is much higher if she has no loss than if a loss is suffered. The difference in wealth on the two axes is the size of the loss. Now look at line A-P_L. This line shows the opportunities for buying insurance at a premium rate which is equal to the expected value of loss for the low risk type. If we ignore transaction costs, and insurers could sell only to low risks at this price, the insurers would break even. For example, if the low risk has a probability "p_L" of suffering a loss of size "D" and insures for a percentage "α" of the loss. The premium would be α p_L D. The line show the wealth combination after buying different levels of insurance. Thus if α is one, there is full insurance and the person is one the full insurance line. The reduction in wealth in the no loss state is simply the premium paid. And the increase in wealth in the loss state is the insurance payout net of the premium. Similarly, the line A-P_H shows the insurance possibilities based on a high risk premium. This line is much shallower signifying that a higher premium must be paid to approach the 45° line.

We know from chapter 2, that if anyone is offered an insurance contract in which the premium equals the expected loss, they would fully insure. If high risk types could be identified and offered insurance at premium rate A-P_H they would fully insure and buy contract labeled "H". This is shown by the point of tangency of the price line and the high risk indifference curve I_{H1}. Similarly, if low risks could be identified and offered a policy priced according to their risk level, they would buy a contract at the point where the 45° line intersects A-P_L. Of course, these contracts cannot be offered, since the types cannot be identified.

Suppose the insurer offers everyone a choice between contract "H" and "L". Notice H is full insurance at the high risk price and L is partial insurance at the low risk price. Who would buy which contract? The high risk type is indifferent; both contracts are on the same high risk indifference curve, I_{H1}. If contract L is defined as marginally below this indifference curve, then the high risk will prefer contract H. Now the low risk type will have indifference curves of different slope. Since the low risk knows her risk type, and therefore knows that there is a smaller probability of having a loss than the high risk type, she will be less willing to give up money in the no loss state, to gain compensation in the loss state.

This implies that, at any point in the diagram, the low risk indifference curves will be steeper. I_L is a los risk indifference curve through the point L which is steeper than I_{H1} at point L. You can see that the curve I_L passes to the right of contract H implying that the low risk strictly prefers contract L to contract H. her is the separation. The high risk will choose contract H and the low risk will choose contract L. They reveal themselves by their contract choices. This is a *separating equilibrium* in which there are no cross subsidies between the risk types.

Despite the separating equilibrium shown, adverse selection is not resolved without a cost. To derive the separation, low risks must choose less than full insurance to signal that they are not high risk types. Low risks must therefore sacrifice some desired insurance protection in order to avoid being pooled with, or mis-classified as, high risk types. The separation works because the cost of giving up some insurance protection is lower for low risk since they know they are less likely to need the insurance protection than the high risks.

Unfortunately, the R/S model does not always produce a separating equilibrium. This problem can be shown in Figure 5 which reproduces some of the essentials from Figure 4 but also has a "pooled price" line. Suppose that a contract could be offered to everyone (high and low risk) which would just break even for the insurer. The pooled price would obviously be somewhere between the high and los price lines, just where in this space would depend on the relative numbers of high and low risk. For example, with equal numbers of high and low risks, the pooled price would be halfway between the high and low price lines. In this way, the deficit to the insurer on each high risk policy would be covered by an equal subsidy from the low risk policy. With more high than low risks, the pooled price would have to extract large subsidies from each of the few low risks, to cover the small deficits to the insurer on each of the many high risk policies . Thus, the pooled price would lie closer to the high risk price line.

Figure 5 depicts a situation where there is a sufficiently large proportion of low risks that the low risk indifference curve, I_L , intersects the pooled price line. In this situation, another insurer could come along and offer a policy such as A which would be preferred by both types to a choice from the menu of H and L; high risk prefer A to H and low risks prefer A to L. So the menu of H and L would not survive in a competitive market. It would seem that the market would settle down to a *pooled equilibrium* in which both types bought a contract such as A. But this situation is not stable. Imagine another insurer now offers a single contract such as B. This is situated between The indifference curves I_{H2} and I_{L2}. Notice that low risks will prefer B to A but that high risks will prefer A to B. Thus all low risks would cancel their A policy and buy B from the rival insurer. A is profitable since it attract low only risks but is priced at a price higher than the fair low risk price P_L. But the insurer stuck with B is now in trouble. The contract was priced to breakeven if were bought by both high and low types, but only highs are left and the insurer will lose money and withdraw the coverage. Without the choice to buy A, the low risks will now flock to B which will become unprofitable if bought by both types since it lies above the pooled price line (indicating that it is priced below the breakeven price. So *there is no equilibrium*.

Let us summarize the messages coming out of this analysis. Whether a separating equilibrium exists or not depends on the relative numbers of high and low risks. Unless there are too man low risks, a separating equilibrium can exist in which low risks select partial coverage at a low price and high risks types take full coverage at a high price. Low risk signal their type by accepting a policy with partial coverage. Since they know themselves to be high risks, this type is not willing to sacrifice some insurance protection to obtain the lower price since they know there is a high probability they will have a loss and need the coverage.

Other Adverse selection Models

The main insight of Rothschild and Stiglitz was that hidden information can be signaled by offering a contract choice in which the contracts will have different appeal to those with different (but hidden) characteristics. Later models have refined the nature of the equilibrium or have examined different signals. In the 1980's attention was focused on different types of equilibrium. Perhaps the most influential approach was that coming from the separate works of Wilson, Spence and Miyazaki. Riley formulated an equilibrium concept in which only contracts which would at least break even, after other firms' anticipated reactions to the introduction of these contracts, would remain. With this equilibrium, Spence and Miyazaki showed that a separating equilibrium always existed but it might require that some subsidy be paid from low to high risks. Figure 6 illustrated their model. The policies H and L are the R/S equilibrium. These policies separately break even for an insurer, thus an insurer offering both policies would obviously break even. Policies H* and L* are another pair, with H* offering full insurance and L* offering partial coverage but over which high risks are indifferent. This pair differs from H and L in that H* is cheaper than H (makes a loss for the insurer) but L* is more expensive than L (and makes a profit for the insurer). However the combination of H* and L* breaks even. The subsidy on L* just covers the deficit on H*. The line LZ denotes pairs of contract matching these criteria. Thus, for any point on LZ there is a policy L', with a corresponding full insurance policy H', such that high risks are indifferent between H' and L' and the pair of policies breaks even for the insurer. H* and L* is the pair of such policies which offers highest utility to low risks which satisfies the self selection constraint. In this equilibrium, both types might be better off than under R/S, but at least we are assured that an equilibrium exists. Notice that the low risk types are signaling their type now by their willingness to accept partial insurance and to pay a subsidy to the high risks types. There is still a price to resolve adverse selection and this price is borne by the low risks.

Multi Periods Contracts and Adverse Selection

The models discussed so far are single period models. The problem is set up as though insurance is traded just once and then the parties go their separate ways. This may be suitable for some markets, such as life insurance, but for other situations the parties can do repeated business over many years. This multi period setting offers other possibilities for resolving adverse selection. We will not go into the various models in detail, but will draw out one or two of the interesting possibilities.

In a series of papers, inlcuding Kunreuther and Pauly (1985), Cooper and Hayes (1987), Hosios and Peters (1989), Doherty and Dionne (1994), Nilssen (1990) have examined the such multi period insurance situations. Perhaps the most interesting thing to emerge is the ability to lower the costs of adverse selection by conditioning the terms of contracts in one period (the price and level of insurance coverage) on the loss experience in prior periods. This, *experience rating*,

Figure 1. Value Maximizing Safety and Insurance

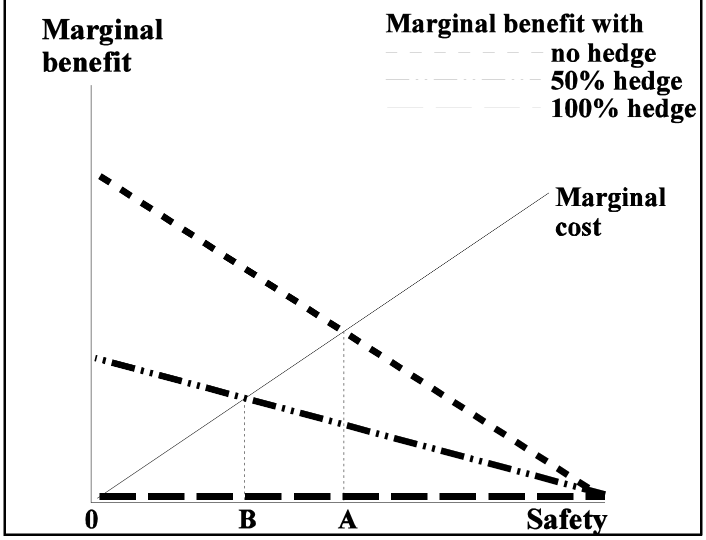

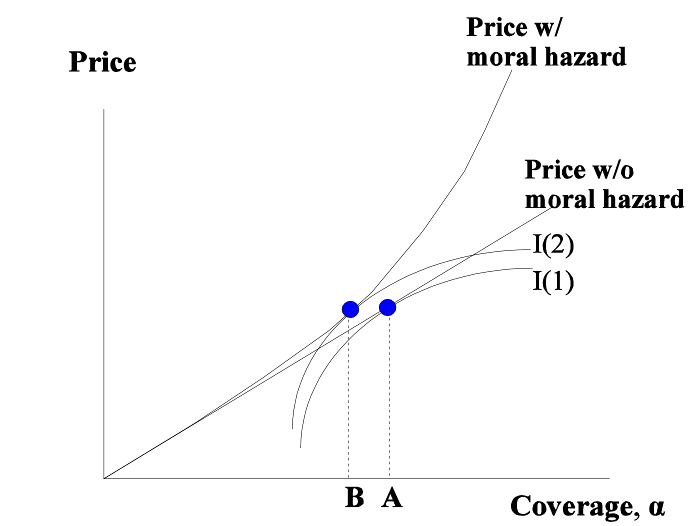

Figure 2. Moral Hazard & Optimal Hedging

Figure 3. Adverse Selection

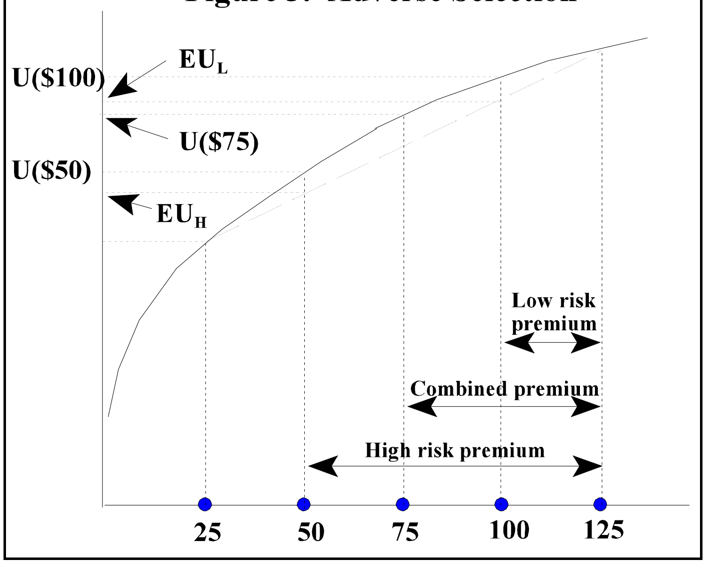

Figure 4. Rothschild Stiglitz Separating Equilibrium

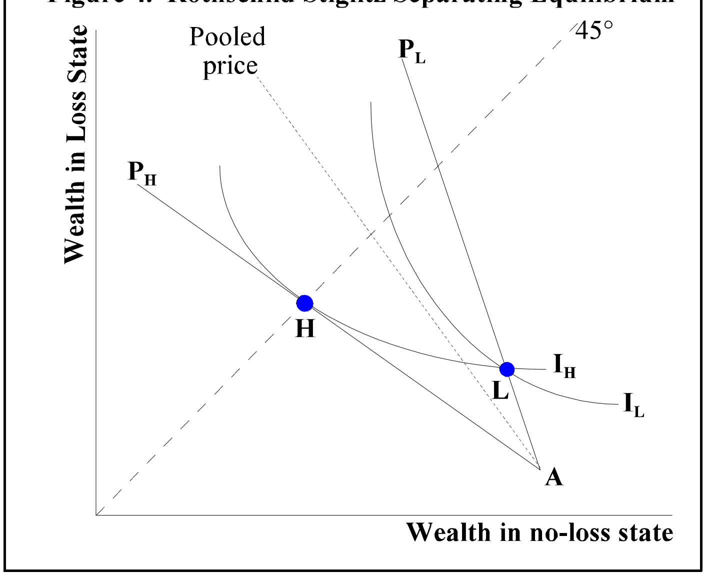

Figure 5. Rothschild Stiglitz - No Equilibrium

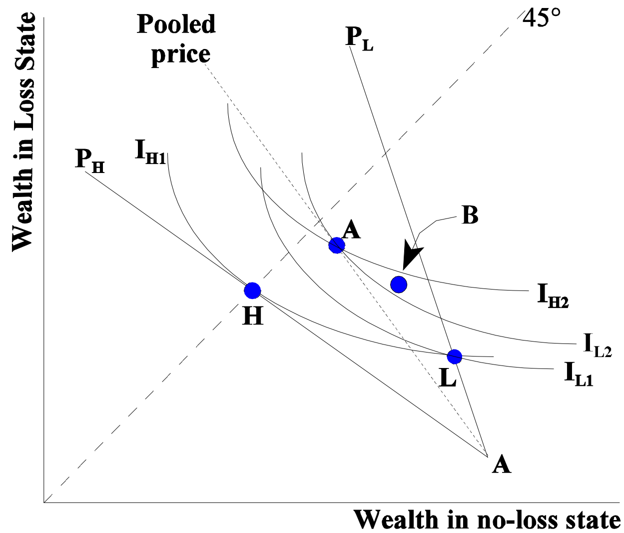

Figure 6. Wilson-Miyazaki-Spence Equilibrium

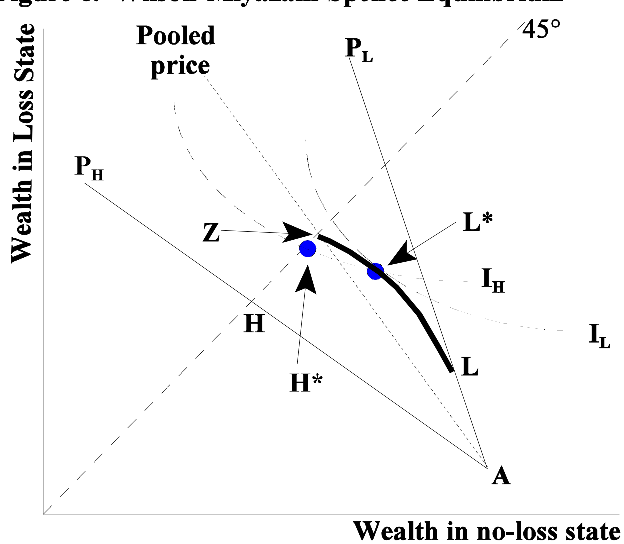

Ackerlof, George, 1970, "The Market for Lemons: Quality Uncertainty and the Market Mechanism", Quarterly Journal of Economics, 84, 488-500

Crocker, Keith and Arthur Snow, 1986, "The Efficiency Effects of Categorical Discrimination in the Insurance Industry", Journal of Political Economy, 94, 321-344.

Dionne, Georges and Neil Doherty, 1982, "Adverse Selection in Insurance Markets: A Selective Survey", in Contributions to Insurance Economics, (ed. Georges Dionne), Kluwer Academic Publishers, Boston, Mass.

Hoy Michael, 1989, "The Value of Screening Mechanisms Under Alternative Insurance Possibilities", Journal of Public Economics, 39, 177-206

Winter, Ralph. 1992, "Moral Hazard and Insurance Contracts", in Contributions to Insurance Economics, (ed. Georges Dionne), Kluwer Academic Publishers, Boston, Mass.