Moral Hazard

so far

Probability and magnitude of loss are exogenously given.

now one change

Insuree can influence either the probability or the magnitude of a loss by exerting effort.

Moral Hazard is an old term in the insurance literature. It refers to the situation in which the purchase of insurance changes the behaviour of the buyer, usually with results unfavourable to the insurance seller.

A first Distinction

Ex ante moral hazard: Loss prevention or loss reduction is undertaken before the loss event occurs, and the existence of insurance affects the incentive to do this. In fact all loss prevention is obviously of this type, some loss reduction may be but some may not.

Ex–post moral hazard: The existence of insurance creates the incentive to make choices that increase the level of the loss, after the event has taken place. This is held to be a particular problem with health insurance: if the patient is insured and so pays only a small proportion of the cost of treatment, his demand for treatment will increase.

We will concentrate on the ex ante MH case and treat the ex–post case only briefly in the end.

Loss reduction

The individual may carry out some costly action to reduce the amount of loss, when the event occurs.

Examples:

- Automatic water sprinklers reduce the damage done by a fire, once it has started.
- Buildings can be located or constructed so as to reduce earthquake, hurricane or flood damage.
- One can treat sickness with more or less expensive medications and procedures, i.e. the patient can demand varying amounts of medical care.
- etc.

Loss reduction – Caveat

In the loss reduction case the first best may become attainable. Assume the loss distribution consists of the four possible losses $0, L_1, L_2, L_3$ if e is incurred and $0, L_1, L_2, L_3'$ if e is not incurred, with $L_3 < L_3'$. This fits what we would call loss reduction.

But this means that there is at least one loss level, L_3', which cannot result if e is incurred, and which can result if e is not incurred. In that case **observation of this loss level ex post tells the insurer that e certainly was not incurred**. We call this a “fully revealing signal”.

Given its probability, the insurer could then write a contract ex ante specifying a large enough punishment if this loss would be observed, to induce the insurance buyer to incur e. Such a solution would only be ruled out if there is some limit on the punishment that could be imposed (**limited liability**).

Loss Prevention

The individual may carry out some costly action to reduce the probability of the loss occurring.

Examples:

- Driving more carefully reduces the probability of an accident, but costs more time and effort.
- Giving up smoking reduces the probability of a number of unpleasant illnesses, but involves fighting with addiction.
- Installing burglar alarms, locks, ferocious dogs reduces the risk of burglary, but at a cost.
- Smoke detectors lead to earlier detection of a fire and therefore lower risk of it taking hold, but at a cost.
- etc.

A (too) simple solution

In reality, the problem of moral hazard may be solved by specifying in the insurance contract certain loss prevention or loss reduction measures, and then checking, in the event of a claim, if these have been met. If not, compensation may be reduced or denied altogether.

E.g., flood damage insurance may specify that a house must not be built on a flood plain. Car accident insurance may not be payable if the claimant is shown to have driven negligently.

We are interested in the case in which such clauses would be too costly to incorporate into the contract, perhaps because it would be too costly to verify (prove before a law court) that the appropriate measures had not been taken.

We want to study how contracts should be designed in the face of the moral hazard problem.

Basic Model

- An individual faces the risk of losing wealth L.
- If she spends an amount e on loss prevention the probability of a loss is $\underline{p}$.
- If she does not spend e the probability of loss is $\overline{p} > \underline{p}$.
- She can also buy insurance against the loss.

We are interested in the relation between the insurance contract and her decision whether or not to spend e. Note that this decision will be taken after insurance is bought, but before the loss may occur.

We assume throughout that the insurance market is perfectly competitive and that there are no insurance costs. This implies two things:

- The premium in equilibrium will always be fair. (Thus we do not have to impose an explicit zero profit condition for sellers.)
- The equilibrium contract will maximise the buyer's expected utility subject to whatever constraints may have to be imposed.

Result under symmetric information

The insurer can observe (and verify !) whether e has been spent or not. In this case he can offer two alternative contracts.

Contract 0: Coverage q is offered at premium $\overline{p}q$. We know that if the buyer chooses this type of contract, she will choose $q = L$, and her utility will therefore be $u(y - \overline{p}L)$ for sure.

Contract 1: Coverage q is offered at premium $\underline{p}q$, and the contract contains a clause that says no compensation will be paid if e is not spent. We know that if the buyer chooses this type of contract, she will choose $q = L$, and her utility will therefore be $u(y - e - \underline{p}L)$ for sure. Then, contract 1 will be chosen if and only if

$$u(y - e - \underline{p}L) \geq u(y - \overline{p}L).$$

I.e., the fall in probability of loss, and therefore fair premium, must be enough to compensate for e. We assume that this is the case.

Result under asymmetric information

The insurer cannot observe (verify) whether e has been spent or not. It is therefore no use including a clause specifying nonpayment of compensation if e is not spent. Instead, the insurer must offer an **incentive compatible contract**.

If the contract assumes that e has been spent, it must provide the incentive to ensure that this will in fact happen. The only instrument the insurer has to do this is the amount of cover, q, that is offered. This must be chosen to ensure incentive compatibility. More precisely, to induce the insurance buyer to spend e, the cover q must be chosen to satisfy the incentive compatibility constraint

$$(1-\underline{p})u(y-e-\underline{p}q)+\underline{p}u(y-e-L+(1-\underline{p})q)$$
$$\geq$$
$$(1-\overline{p})u(y-\underline{p}q)+\overline{p}u(y-L+(1-\underline{p})q)$$

I.e., the insuree must be at least as well off with the insurance contract giving cover q at premium $\underline{p}q$ when spending e, as she would be with the same contract and not spending e.

This must involve less than full cover.

Suppose $q = L$:

The inequality becomes

$$u(y - e - \underline{p}q) \geq u(y - \underline{p}q)$$

which cannot be true as long as $e > 0$.

The competitive market assumption implies that the equilibrium contract is found by maximising the buyer's utility subject to the incentive compatibilty constraint. It is clear that this latter constraint must be binding. For if not, the solution would imply that cover simply maximises the buyer's expected utility, i.e. must be full cover, but we have already seen this cannot satisfy the incentive compatibility constraint.

Thus the optimal $q^* < L$ will satisfy

$$(1-\underline{p})u(y-e-\underline{p}q^*)+\underline{p}u(y-e-L+(1-\underline{p})q^*)$$

$$=$$

$$(1-\overline{p})u(y-\underline{p}q^*)+\overline{p}u(y-L+(1-\underline{p})q^*)$$

Another Basic Model

The above model assumed that the variable e was a monetary cost. But we could well argue that for monetary expenditures the MH problem is less severe. It is easy (or at least easier) to verify whether some investments have been made to reduce risk (like installing smoke detectors or burglar alarms).

Whereas the MH problem gets more severe (and plausible) if we are talking about non–monetary effort costs that have to be born in order to reduce risk (like giving up smoking). We write this cost in terms of utility as $c(e)$, with $c'(e) > 0$, $c''(e) > 0$.

It is convenient to assume the buyer's utility function takes the additively separable form $u(\cdot) - c(e)$, where $u(\cdot)$ is the (concave) utility of income.

The **symmetric information** case goes through just as before. [check this]

For the **asymmetric information** case, the incentive compatibility constraint now takes the form

$$(1-\underline{p})u(y-\underline{p}q^*)+\underline{p}u(y-L+(1-\underline{p})q^*)-c(a)$$
$$=$$
$$(1-\overline{p})u(y-\underline{p}q^*)+\overline{p}u(y-L+(1-\underline{p})q^*)$$

The same argument as before shows that this constraint must be binding at the optimum.

Let $y_1 \equiv y - pq^*$ and $y_2 \equiv y - L + (1-p)q^*$.

Then we can rearrange this constraint to obtain

$$u(y_1^*) - u(y_2^*) = \frac{c(e)}{\overline{p}-\underline{p}} > 0$$

From this simple condition we know that

- $q^* < L$ (since this is the only way that we can have $y_2 < y_1$).
- this difference must be greater

– the higher $c(e)$, the utility cost of the loss prevention activity, is.

– the lower $\overline{p} - \underline{p}$, the reduction in loss probability brought about by the loss prevention activity, is.

The intuition is straightforward. The greater the difference in income between the two states, the lower must be the cover q. The higher the utility cost of loss prevention, the lower must be the cover, to provide suffcient incentive to the buyer to reduce the risk of loss to herself by undertaking loss prevention. We can think of $\overline{p} - \underline{p}$ as the effectiveness of the loss prevention activity. The lower is this effectiveness, the lower must be the cover, again in order to provide the incentive to undertake loss prevention.

Contractual structure

The previous basic models confirmed the intuitive result:

The reponse to the moral hazard problem is partial cover.

Open question: What kind of partial cover? A simple deductible or something more complicated?

To address this point we now make a small extension to our last model with non–monetary effort costs. We will show that when loss prevention changes the probability that a loss will occur, without however changing the probabilities of specific loss levels conditional on a loss having occurred, then the correct form of partial cover is a deductible. If however the loss probability distribution is changed more generally, then partial cover will optimally take more complicated forms.

- There are three loss states, with losses $L_1 < L_2 < L_3$.
- Let $\underline{p}$ be the probability that a loss occurs if loss prevention activity e is undertaken, and $\overline{p} > \underline{p}$ that if it is not.
- Conditional on there being a loss, π_s, $s = 1, 2, 3$ is the probability of loss of size L_s, with $\Sigma\pi_s = 1$.
- Thus the loss values $0, L_1, L_2, L_3$ have respectively the probabilities $1-\underline{p}, \underline{p}\pi_1, \underline{p}\pi_2, \underline{p}\pi_3$, if loss prevention is undertaken, and $1-\overline{p}, \overline{p}\pi_1, \overline{p}\pi_2, \overline{p}\pi_3$ if not.
- Recall that the competitive market assumption (with no insurance costs) implies that the equilibrium contract must **maximise the buyer's expected utility**, and **have a fair premium**. It must also **satisfy the incentive compatibility constraint**.

The problem takes the form

$$\begin{aligned}
\max \quad & (1-\underline{p})u(y_0) + \underline{p}\Sigma\pi_s u(y_s) - c(e) &=& \ EU \\
\text{(s.t. PC)} \quad & P - \underline{p}\Sigma\pi_s q_s &=& \ 0 \\
\text{(s.t. IC)} \quad & (1-\underline{p})u(y_0) + \underline{p}\Sigma\pi_s u(y_s) - c(e) &\geq& \\
& (1-\overline{p})u(y_0) + \overline{p}\Sigma\pi_s u(y_s) & &
\end{aligned}$$

where $y_0 = y - P$, $y_s = y - P - L_s + q_s$.

The Lagrange function for this problem is

$$L = EU + \lambda(P - \underline{p}\Sigma\pi_s q_s) + \mu\left[EU - (1-\overline{p})u(y_0) - \overline{p}\Sigma\pi_s u(y_s)\right].$$

For simplicity we assume all coverage values q_s are positive in the optimum (you can check the alternative case as an exercise). The FOCs are then given by

$$\frac{\partial L}{\partial q_s} = \underline{p}\pi_s u'(y_s^*) - \lambda^*\underline{p}\pi_s - \mu^*\pi_s u'(y_s^*)[\overline{p} - \underline{p}] = 0$$

and

$$\begin{aligned}
\frac{\partial L}{\partial P} &= \lambda^* - (1-\underline{p})u'(y_0^*) + \underline{p}\Sigma\pi_s u'(y_s^*) \\
&+ \mu^*(\underline{p} - \overline{p})\Sigma\pi_s u'(y_s^*) = 0
\end{aligned}$$

together with the constraints. From $\frac{\partial L}{\partial q_s}$ we obtain $u'(y_s^*) = \frac{\lambda^*}{1-\mu^*[\overline{p}-\underline{p}]/\underline{p}}$

Note that the right hand side is independent of s. In other words we have

$$u'(y_1^*) = u'(y_2^*) = u'(y_3^*)$$

implying

$$y - P - L_1 + q_1^* = y - P - L_2 + q_2^* = y - P - L_3 + q_3^*.$$

Since we know full cover cannot satisfy the incentive compatibility constraint, we must have

$$L_1 - q_1^* = L_2 - q_2^* = L_3 - q_3^* = D > 0.$$

Thus the optimal contract has cover in each state equal to loss minus a deductible $q_s^* = L_s - D$.

The reason is that this allows marginal utilities across loss states to be equalised, while meeting the condition that there should only be partial cover.

The Sufficient Statistics Result

Holmström (1982) has shown the following:

The optimal contract should condition on those and only those variables / signals that are informative with respect to the agent's effort choice (or more generally the action the principal is interested in).

So **all** informative signals and **no** uninformative one should be included in the optimal contract.

Examples:
manager & firm : $\Rightarrow$ profit
insuree & insurance: $\Rightarrow$ accident/loss or not; magnitude of loss (?)
...

A more general model

It is clearly quite special to assume that the loss prevention activity affects only the probability of having a loss or not, and not the probability of a loss conditional on there being one. So now we generalize and assume that, given the 4 possible loss levels $\{0, L_1, L_2, L_3\}$, the respective probabilities are $\underline{p}_s$ if effort e is incurred and $\overline{p}_s$ if not, with $s = 0, .., 3$ and $\Sigma \underline{p}_s = \Sigma \overline{p}_s = 1$.

We expect that loss prevention would lead to an improvement, in some sense, in the loss distribution. A general formulation would be to say that the distribution of income with e would stochastically dominate the distribution without e to the first or even second order.

A problem with the analysis is that this leaves open a large number of possibilities for the changes in probabilities.

To make things concrete, let us assume

$$\underline{p}_0 > \overline{p}_0,\ \underline{p}_1 > \overline{p}_1,\ \underline{p}_2 < \overline{p}_2,\ \underline{p}_3 < \overline{p}_3$$

in other words, the probabilities of the lower loss levels are increased and those of the higher loss levels are reduced by incurring e. The problem of finding the optimal contract is now written as

$$\begin{array}{lrcl} \max & \Sigma \underline{p}_s u(y-s) - c(e) & = & EU \\ \text{(s.t. PC)} & P - \Sigma \underline{p}_s q_s & = & 0 \\ \text{(s.t. IC)} & \Sigma \underline{p}_s u(y_s) - c(e) \geq \Sigma \overline{p}_s u(y_s) & & \end{array}$$

The Lagrange function is now

$$L = EU + \lambda(P - \Sigma \underline{p}_s q_s) + \mu \left[EU - \Sigma \overline{p}_s u(y_s) \right]$$

and the FOCs are then given by

$$\frac{\partial L}{\partial q_s} = \underline{p}_s u'(y_s^*) - \lambda^* \underline{p}_s - \mu^* u'(y_s^*) \left[\overline{p}_s - \underline{p}_s \right] = 0 \text{ with } s = 1, 2, 3$$

and

$$\frac{\partial L}{\partial P} = \lambda^* - \left[\Sigma \underline{p}_s u'(y_s^*) + \mu^* (\underline{p}_s - \overline{p}_s) \Sigma u'(y_s^*) \right] = 0.$$

together with the constraints (as equalities). We can write the first conditions for $s = 1, 2$ as

$$\begin{aligned} u'(y_1^*) &= \frac{\lambda^*}{1 + \mu^*[\underline{p}_1 - \overline{p}_1]/\underline{p}_1} \\ u'(y_2^*) &= \frac{\lambda^*}{1 + \mu^*[\underline{p}_2 - \overline{p}_2]/\underline{p}_2}. \end{aligned}$$

Recall that $\underline{p}_1 > \overline{p}_1, \underline{p}_2 < \overline{p}_2$, by assumption. It then follows that

$$u'(y_1^*) < u'(y_2^*)$$

or

$$L_1 - q_1^* < L_2 - q_2^*$$

Thus in this case **we cannot have a constant deductible**, but rather **the difference between loss and cover increases** with the loss.

It would be tempting to conclude in the same way that

$$L_2 - q_2^* < L_3 - q_3^*$$

so that we could talk perhaps of coinsurance in this case. Note however that **this is not implied by the assumptions** we have made so far, even though this assumption was pretty special. For this we require

$$u'(y_2^*) < u'(y_3^*)$$

.

implying

$$\frac{\lambda^*}{1 + \mu^*[\underline{p}_2 - \overline{p}_2]/\underline{p}_2} < \frac{\lambda^*}{1 + \mu^*[\underline{p}_3 - \overline{p}_3]/\underline{p}_3}$$

in turn implying

$$[\underline{p}_2 - \overline{p}_2]/\underline{p}_2 > [\underline{p}_3 - \overline{p}_3]/\underline{p}_3$$

which is called a **monotone likelihood ratio condition** and tells us something about the informativeness of a particular outcome w.r.t. the effort choice.

If we do not assume this then the gap between loss and cover may not be increasing monotonically with the loss. This kind of issue is familiar from general Principal Agent Theory.

Continuous effort levels

We can model the MH problem (more elegantly) in a continuous way. The insuree now can choose his level of care (effort) continuously, thus influencing the loss probability continuously.

The problem takes then the following form:

$max_{P,q,e}$

$(1-\pi(e))\,u(Y-P)+\pi(e)u(Y-L-P+q)-c(e)$

s.t. PC

$$(1-\pi(e))\,P-\pi(e)(q-P)\geq 0$$

s.t. IC

$e\in\arg\max$

$[(1-\pi(\tilde{e}))\,u(Y-P)+\pi(\tilde{e})u(Y-L-P+q)-c(\tilde{e})]$

We are confronted with two intertwined optimization problems. On the one hand the principal tries to maximize the agent's expected profit by choosing an appropriate insurance scheme – subject to the agent's effort choice. The latter in turn is an optimal reaction to this insurance scheme chosen by the principal. Sir James Mirrlees came up with an idea how to get rid of the tricky max term in the IC. He replaced the IC by its FOC w.r.t. e.

What does that mean?

In optimum, if the agent has chosen e correctly this condition must hold as then the marginal utility of exerting a marginally higher effort level equals 0.

The new, more easily to handle, problem now reads:

$max_{P,q,e}$

$$(1-\pi(e))\,u(Y-P)+\pi(e)u(Y-L-P+q)-c(e)$$

s.t. PC

$$(1-\pi(e))\,P-\pi(e)(q-P)\geq 0$$

s.t. IC'

$$[-\pi'(\tilde{e})u(Y-P)+\pi'(\tilde{e})u(Y-L-P+q)-c'(\tilde{e})]=0$$

Unfortunately the *First Order Approach* is not always applicable. To be sure we have to restrict to special distribution functions. These ensure that the agent's problem – given the optimal wage scheme – is concave in effort.

Otherwise we have to use the FOA, solve for the optimal wage scheme and then check whether the agent's problem indeed is concave in e, i.e. whether we used the FOA justly.

Dynamic Properties of MH contracts

The intuition for a repeated MH situation is similar to the one in the AS context. We should observe punishment or reward for past performance in later periods of the relationship.

However clearcut predictions of theses types of models are hard to derive and depend delicately on the assumptions we make concerning the access to financial markets, i.e. whether we allow the agent to smooth his income over time and states herself by borrowing and saving.

Ex–post Moral Hazard

The basic problem

Consider the case of health insurance. The agent is confronted with the risk of getting ill / having an accident.

Now the magnitude of the loss depends on the agent's actions after the risk has realized, i.e. the agent has fallen ill. The agent is (roughly speaking) the one to decide how much treatment to "consume".

If he is fully insured, i.e. the insurance has to cover all the treatment costs, the agent will tend to consume to much treatment as he does not have to bear its costs.

Solution

$\Rightarrow$ The agent has to bear the cost partly himself.

$\Rightarrow$ partial insurance (deductible or coinsurance ?)

A simple model

- two possible states i(ll) and h(ealthy)
- probability of illness π is exogenously given.
- two goods: treatment x and consumption good y with prices equal to 1; income Y

Now consider an insurance.

- Premium P
- Insurance sets a coinsurance rate c (cf. Sufficient Statistics result)
- Net income
  $Y_i = Y - P - cx$ and $Y_h = Y - P$

Expected utility

$$EU = \pi \left[u(Y_i) + v(x)\right] + (1 - \pi)u(Y_h)$$

with $u'(\cdot) > 0$ and $u''(\cdot) < 0$ and where v(x) is the utility from the consumption of treatment. Again $v'(\cdot) > 0$ and $v''(\cdot) < 0$ is assumed.

Agent maximizes his utility over the choiice of x.

$$[FOC] \qquad \frac{v'(x)}{u'(Y_i)} = c$$

As always, the MRS has to equal the price ratio. But note that the socially correct price of x is not c but 1. Thus we have overconsumption of x whenever $c < 1$, i.e. whenever the agent is insured.

The overconsumption decreases in c.

Now the task is to find an optimal tradeoff between provision of insurance (consumption smoothing) and providing incentives to avoid treatment overconsumption.

The optimal level of c will depend on the price elasticity of the demand for x and the degree of risk aversion.

Insurance Fraud

Costly state verification

The insurance company can only by a costly Gutachten control whether there really was an accident.

The optimal contract (assuming CARA prefernces) entails auditing by the insurance company.

- If fraud is detected that causes maximal punishment.
- Auditing is random. The auditing probability increases in the magnitude of the loss.
- Contracts have a deductible. Audited losses have a lower deductible.

Problem:

If auditing rules out fraud in the first place, why audit after all?

But this would be anticipated by the insurees ...

Costly state falsification

The magnitude of the loss can be manipulated (at a cost) by the insuree.

- The optimal contract entails partial insurance at the margin.
- The marginal coverage equals the costs of falsifying the claim.
- Small losses are overinsured, large losses underinsured.