Information Economics notes for Econ 8106

Dimitrios Diamantaras
Department of Economics
Temple University
Philadelphia, PA 19122
dimitris [at] temple.edu

2009-11-16

1 Adverse Selection

This is a very brief note on JR chapter 8. We start with Adverse Selection, relying on these authors' presentation (there is very little that is original here). We concentrate on an auto insurance example, but there are other great examples of the applicability of the tools we are talking about here, for instance in the labor market (see MWG chapter 13).

There are m identical consumers, but they have different probabilities of getting in an accident. Probability for consumer i is π_i. Accidents occur independently across consumers. Each consumer has wealth w and suffers loss L if an accident occurs. Each has utility function u, continuous, strictly increasing, vnM, strictly concave.

Insurance companies are identical. They only sell full insurance. You can only buy or not buy this insurance; they do not sell it in fractional amounts. If bought at price p, a single insurance policy gives expected profit of $p - \pi_i L$ to the issuing insurance company.

1.1 Symmetric Information

Here we assume that the insurance companies can costlessly see each consumer's true accident probability.

The price of a commodity depends on the state of the world. Policies in this model pay out in different states, so they are different commodities and receive different prices. Denote by p_i the price of the policy to the ith consumer. We want to find the competitive equilibrium price p_i^* for each consumer.

Consider the supply of policy i. For a price less than $\pi_i L$, quantity supplied is 0. For a price more than $\pi_i L$, quantity supplied is infinite. It follows that we can only find a competitive equilibrium at the price $p_i^* = \pi_i L$, as at this price the insurance companies are ready to supply any number of insurance policies to person i.

Consider now the demand. Consumers are risk averse because of the strictly concave utility function u assumption. So at any price less than or equal to $\pi_i L$ (the actuarially fair price), consumer i wants to fully insure. There will be a price above $\pi_i L$ at which, and at prices above it, the consumer will no longer want a single policy.

Putting demand and supply together we see that $p_i^* = \pi_i L$ is indeed the equilibrium price for consumer i.

Claim 1.1. This competitive outcome (under all the strong assumptions we have made) is Pareto efficient.

We skip the proof as it is well done in JR and we need the time for other models.

1.2 Asymmetric Information and Adverse Selection

Now remove the assumption that insurance companies know the π_i values. We will assume that they know the distribution of accident probabilities. Also, we will assume that the accident probabilities all come from some non-trivial interval $[\underline{\pi}, \overline{\pi}]$. The cumulative is F, so $F(\pi)$ is the fraction of consumers with accident probability less than or equal to π. This is the only change of assumptions.

But this change makes a huge difference:

Claim 1.2. There can no longer be different prices for different consumers at a competitive equilibrium.

We prove this by contradiction: if there were two consumers purchasing policies at different prices, then insurance companies must be making no losses at either policy, so must make a positive profit at one policy, the one with the higher price, and then there would be infinite quantity supplied of that policy at that higher price, thus precluding equilibrium.

So we can only have one price, say p^*, at equilibrium. What can it be?

Claim 1.3. It cannot be the natural candidate, $E(\pi)L$, where $E(\pi) = \int_{\underline{\pi}}^{\overline{\pi}} \pi dF(\pi)$.

Indeed, if it were, the estimate of the probability as the unconditional mean would be too low, leading to losses for the insurance companies. For the reason, we look at the decision of a consumer on buying insurance.

Take a generic consumer with accident probability π. This consumer will buy insurance at price p if and only if

$$u(w-p) \geq \pi u(w-L) + (1-\pi)u(w).$$

Rearranging leads to

$$\pi \geq \frac{u(w) - u(w-p)}{u(w) - u(w-L)}.$$

We'll need a shortcut for the function of p on the right hand side, so call it $h(p)$.

Definition 1.4. We call p^* a competitive equilibrium price under asymmetric information if it satisfies the condition

$$p^* = E(\pi|\pi \geq h(p^*))L, \tag{1}$$

where $E(\pi|\pi \geq h(p^*)) = \int_{h(p^*)}^{\overline{\pi}} \pi dF(\pi)$ is the expected accident probability conditional on the condition that ensures that people will buy the policy.

Does such an equilibrium exist? The answer is affirmative. But there can be multiple equilibria, and equilibria do not have to be Pareto efficient. The book has an example of how extremely inefficient such an equilibrium can be, in which case we should be talking about complete market failure. (Akerlof's famous "lemons" paper of 1970 was the first to show this possibility.) This is an example of **adverse selection**.

Conclusion: when adverse selection exists, competitive equilibrium can be very inefficient, missing out on many mutually beneficial trades.

1.3 Signaling

Consumers can try to avoid the adverse selection problem by signaling their accident probability, with those with lower probabilities presumably more eager to signal this fact. But are the signals to be believed?

They may, if they involve offering to buy different types of policies, and if the bad risks would not want to buy the policies that the good risks would buy.

Here is a simplified model. There are now only two accident probabilities: $0 < \underline{\pi} < \overline{\pi} < 1$. Fraction $\alpha \in (0,1)$ of consumers have the low probability, $\underline{\pi}$. We model the situation with a stylized game.

- Nature moves first. It determines which consumer will make a proposal to the insurance company. A low-risk consumer is chosen with probability α.
- The chosen consumer chooses a policy (B, p), where $B \geq 0$ is the benefit and p, with $0 \leq p \leq w$ is the premium.
- The insurance company moves last without knowing the type of the consumer, and either accepts or rejects the proposal.

When thinking about this game, remember that the insurance company is one of many competing ones, and that the consumer is a randomly chosen one from the overall population of consumers.

Pure strategies for the two kinds of consumers: $\psi_l = (B_l, p_l)$, $\psi_h = (B_h, p_h)$. Pure strategy for the insurance company: specify A (accept) or R (reject) for each potential policy proposed. Notation: $\sigma(B, p) \in \{A, R\}$ for each (B, p). Note the argument: it is the proposal, not the consumer's type. If at equilibrium the two types make different proposals, then the company can use the proposals to tell which consumer type it is facing.

Beliefs: $\beta(B, p)$ is the probability that the company believes that the proposal (B, p) came from a low type.

A continuum of strategies for the consumer exists, so we have a technical issue with sequential equilibrium, as it was only defined for finite games. This is neatly sidestepped in the book, so we don't have to worry. We thus get to the following.

Definition 1.5. [Signaling Game Pure Strategy Sequential Equilibrium]

The assessment $(\psi_l, \psi_h, \sigma(\cdot), \beta(\cdot))$ is a pure strategy sequential equilibrium of the insurance signaling game if

(1) Given the Insurance company's strategy $\sigma(\cdot)$, proposing ψ_l maximizes the low-risk consumer's utility and proposing ψ_h maximizes the high-risk consumer's utility;

(2) The Insurance company's beliefs satisfy Bayes' rule, which here means three things: (a) for all policies $\psi = (B, p)$, $\beta(\psi) \in [0,1]$, (b) if $\psi_l \neq \psi_h$, then $\beta(\psi_l) = 1$ and $\beta(\psi_h) = 0$, and (c) if $\psi_l = \psi_h$, then $\beta(\psi_l) = \beta(\psi_h) = \alpha$;

(3) For every policy $\psi = (B, p)$, $\sigma(\psi)$ maximizes the company's expected profit given the belief $\beta(B, p)$.

Conditions (1) and (3) impose sequential rationality, and (2) Bayes' rule. Bayes' rule is here very simple because we are only considering pure strategies.

It is not obvious that a low-risk consumer can distinguish herself by choosing a different policy proposal than the high-risk one (and here is where the insurance example is better than the education example used in MWG). This is because the signals are unproductive. Offering to buy a lot or a little insurance does not change the probability of having an accident.

Analysis of the Signaling Game

Notation: $u_l(B,p) = \underline{\pi} u(w-L+B-p) + (1-\underline{\pi})u(w-p)$, and similarly for $u_h(B,p)$. We have the following from our basic assumptions, provable without much effort. One piece of notation before we start: $MRS_i(B,p)$ stands for the marginal rate of substitution of a customer with type $i = l$ or $i = h$, and the marginal rate of substitution is in the form $dB/dp = D_p u_i(B,p)/D_B u_i(B,p)$.

1. $u_l(B,p)$ and $u_h(B,p)$ are continuous, differentiable, strictly concave in (B,p), strictly increasing in B, and strictly decreasing in p.

2. $MRS_l(B,p)$ is greater than, equal to, or less than $\underline{\pi}$ as B is less than, equal to, or greater than L. For $MRS_h(B,p)$, the same statements hold, with $\overline{\pi}$.

3. $MRS_l(B,p) < MRS_h(B,p)$ for all (B,p).

The last one is the single-crossing property. It means that the indifference curves of the two types intersect exactly once. Further, these indifference curves have different MRSs at the same policy. The next graph illustrates items 1 and 3.

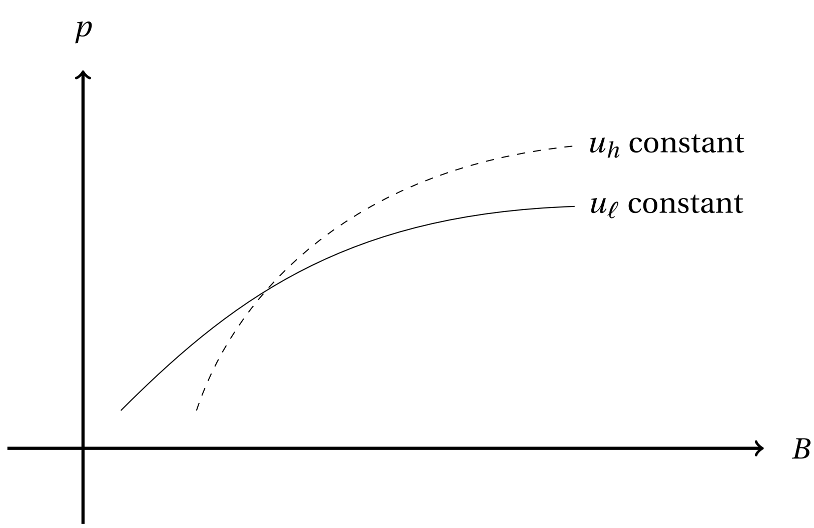

The insurance company will accept a policy with $p \geq \pi B$, when it is confident that the individual proposing the policy has accident probability π. When it believes that the consumer is a low-risk individual, it uses $\pi = \underline{\pi}$; a high-risk one, $\pi = \overline{\pi}$; when unsure, $\pi = \alpha\underline{\pi} + (1-\alpha)\overline{\pi}$. Its break-even lines are shown below.

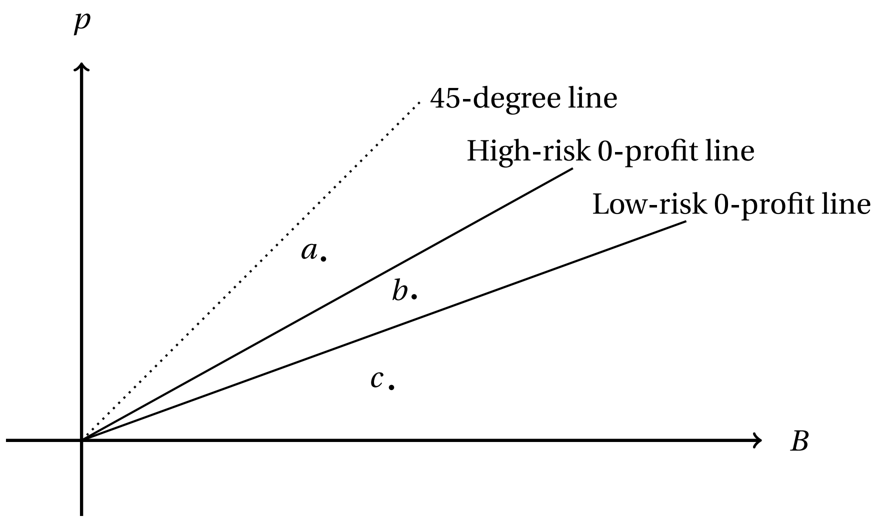

The policy at point a is profitable for both consumer types; b is profitable for the low-risk type but unprofitable for the high-risk type; c is unprofitable for all types. The high-risk zero-profit line has slope $\overline{\pi}$ and the low-risk one has slope $\underline{\pi}$.

The competitive outcome in case the insurance company can tell the types apart by costless observation is shown below.

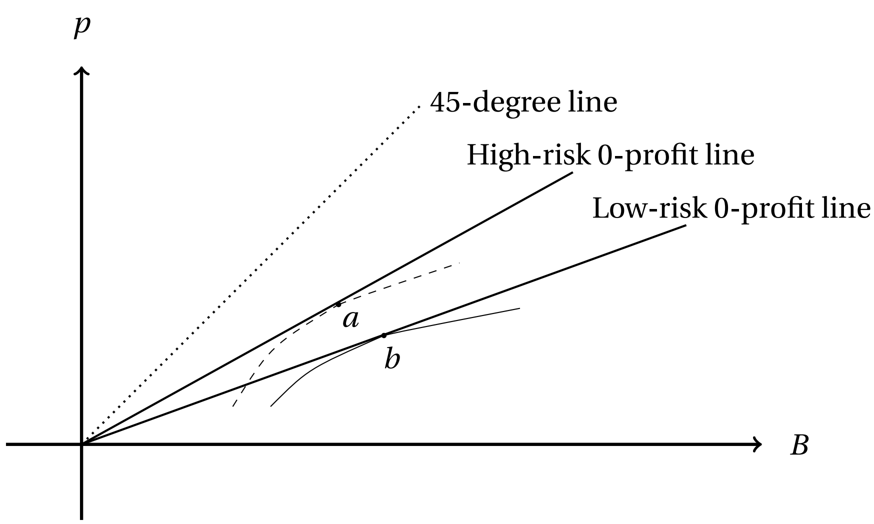

In this equilibrium each consumer insures fully and the insurance company makes zero profits from each. But what is the equilibrium without costless observability? We start by a basic property of any sequential equilibrium in this game.

Lemma 1.6. *Let* $(\psi_l, \psi_h, \sigma(\cdot), \beta(\cdot))$ *be a sequential equilibrium and let* u_l^*, u_h^* *be the equilibrium utility of the low- and high-risk consumer. Let* $\tilde{u}_l \equiv \max_{(B,p)} u_l(B,p)$ *such that* $p = \overline{\pi}B \leq w$ *and* $u_h^c \equiv u_h(a)$, *where* a *is the point at the full-information equilibrium shown above. Then: (1)* $u_l^* \geq \tilde{u}_l$ *and (2)* $u_h^* \geq u_h^c$.

It follows that the high-risk consumer must purchase insurance at equilibrium (since u_h^c is strictly preferred to the origin due to the strict risk aversion assumption).

Definition 1.7. A pure strategy sequential equilibrium $(\psi_l, \psi_h, \sigma(\cdot), \beta(\cdot))$ is **separating** if $\psi_l \neq \psi_h$. Otherwise, it is **pooling.**

The basic idea: each type of consumer must not find it preferable to the truth to pretend it is the other type.

Theorem 1.8 (Separating Equilibrium Characterization)**.** *The policies* $\psi_l = (B_l, p_l)$ *and* $\psi_h = (B_h, p_h)$ *are proposed by the low and high risk types, respectively, and accepted by the insurance company in some separating equilibrium if and only if*

1. $\psi_l \neq \psi_h = (L, \overline{\pi}L)$;

2. $p_l \geq \underline{\pi} B_l$;

3. $u_l(\psi_l) \geq \tilde{u}_l \equiv \max_{(B,p)} u_l(B,p)$ s.t. $p = \overline{\pi}B \geq w$;

4. $u_h^c \equiv u_h(\psi_h) \geq u_h(\psi_l)$.

The proof of this is instructive, but even more instructive is Figure 8.6 in JR. Basic features of any low-risk policy in any separating equilibrium:

- It is above the low-risk zero-profit line, so the firm will accept it.

- It is above the high-risk consumer's indifference curve through the high-risk equilibrium policy so the high-risk consumer does not have an incentive to mimic the low-risk one. (Note: above the indifference curve in this graph means less preferred.)

- It is below the indifference curve that gives utility $\tilde{u}_l$ to the low-risk consumer so that the low-risk consumer has no incentive to deviate and be identified as a high-risk consumer.

In a separating equilibrium, the policy proposal that a consumer makes to the firm is a signal of the consumer's type. So if we expect that a separating equilibrium will prevail, we expect that it will improve the efficiency of the market outcome.

The efficiency improvement may be unimpressive. If the indifference curve of the low risk consumer that passes through the null policy (the policy $(0,0)$) is not steeper than the high-risk break-even line, then there is a separating equilibrium in which the low-risk consumer gets the null policy, which means that the good customers get no insurance! Even worse, this is true (when the slopes are as stated in this paragraph) even if there are one billion low-risk consumers and only one high-risk consumer.

But there are also separating equilibria in which the low-risk consumer gets some insurance. The best one of these for the low-risk consumer is shown in Figure 8.7 in JR.

We see the problem of multiple equilibria here in a particularly strong form. But, even worse, we are not done seeing how many equilibria we have in the insurance game. We turn now to pooling equilibria.

When both types of consumer propose the same policy in equilibrium, the insurance company has no signal to use to tell them apart. In such a situation, the high-risk consumer type is masquerading as a low-risk one.

Remember that $\pi = \alpha\underline{\pi} + (1-\alpha)\overline{\pi}$ is the slope of the break-even line when the firm cannot tell the consumer types apart.

Theorem 1.9 (Pooling Equilibrium Characterization)**.** *The policy $\psi' = (B', p')$ is the outcome of a pooling equilibrium if and only if it satisfies the inequalities*

$$\begin{aligned} u_l(B', p') &\geq \tilde{u}_l, \\ u_h(B', p') &\geq u_h^c, \\ p' &\geq \pi B'. \end{aligned}$$

The beliefs underlying the separating and the pooling equilibrium are extreme and worth discussing.

In each equilibrium, if the firm received a policy proposal that was not part of the equilibrium, it would believe that this policy proposal came from a high-risk consumer. This is a valid belief system according to the definition of sequential equilibrium, but is it intuitively reasonable?

This question morphs into the question: “what additional restrictions might be reasonable to impose on a sequential equilibrium to achieve a good refinement of sequential equilibrium?”

Consider the following figure.

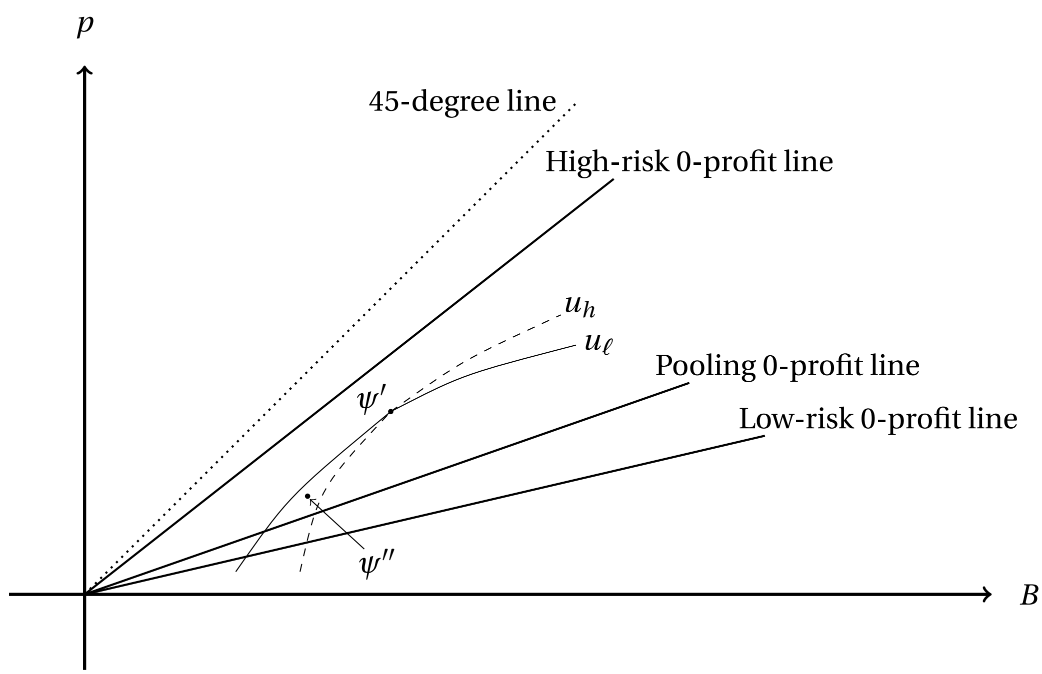

Here, ψ' denotes a typical pooling equilibrium policy. Suppose that all players (including the firm) think that this is the equilibrium being played, and yet suddenly the firm receives the proposal ψ''. According to the belief system we discussed, the firm should now believe that it was the high-risk consumer who made this proposal. But a glance at the indifference curve configuration shows that this is not reasonable, as the high-risk consumer would become worse off with this proposal relative to the equilibrium proposal ψ'. This leads to the following refinement of sequential equilibrium, due to Cho and Kreps, strangely named “the intuitive criterion”.

Definition 1.10. A sequential equilibrium $(\psi_l, \psi_h, \sigma(\cdot), \beta(\cdot))$ that yields u_l^* and u_h^* to the low- and high-risk consumer, respectively, satisfies the **intuitive criterion** if the following condition is satisfied for each policy ψ such that $\psi \neq \psi_l$ and $\psi \neq \psi_h$ and for $i = l, h,\ j \neq i$:

If $u_i(\psi) \geq u_i^*$ and $u_j(\psi) < u_j^*$, then the belief system β places probability one on risk type i.

Theorem 1.11. *There is a unique policy pair* (ψ_l, ψ_h) *that can be supported by a sequential equilibrium that satisfies the intuitive criterion. This equilibrium is the best separating equilibrium for the low-risk consumer.*

1.4 Screening

What if the insurance company plays first, and proposes different contracts to attract different consumer types? This sequence of events certainly seems to fit reality. The idea is to screen consumers, which means that the high-risk types buy one kind of policy, while the low risk types buy another.

[MATERIAL OMITTED]

2 Moral Hazard and Principal-Agent Problems

[MATERIAL OMITTED]

In the adverse selection situation, one side of an economic relationship faces the problem of attracting undesirable partners (customers, in the insurance example we used). Asymmetric information ensures that the company, or agent, who faces the adverse selection problem cannot easily avoid it by simply refusing to deal with the undesirable customers, since it does not know who they are. We now consider another serious problem that comes from asymmetric information: not knowing what your partner does. This is called **moral hazard** because of the terminology in the insurance field. Think of the old story, to be found in movies even, of someone burning an insured property surreptitiously in order to pocket the insurance.

Another example is in the Mel Brooks movie *The Producers,* in which a Broadway producer and an accountant raise money for a theatrical production which they want to see fail, so they can grab the investors' money and leave for a warm place far, far away. They figure that if the production is a total flop, nobody will bother to audit the books and so they will get away with their plan. In order to ensure the failure of the production, they get a Hitler admirer to write a play about Hitler and they turn it into a hideous musical. The point, regarding moral hazard, is that the *principals,* in this case the investors in the production, cannot see that the actions of the *agents,* the producer and his accomplice, are against the principals' interests. In the end, the audience interprets the musical as a funny satire of Hitler, the play is a great success, and the plot is foiled.

Other examples abound. MWG discuss the owner of a business as the principal, and the manager as the agent. This includes the Mel Brooks example but also many, many other such relationships between principals and agents. JR choose to continue looking into the insurance field and place the principal-agent analysis in the framework that has the insurance company as the principal and the customer as the agent.

Consider a model with one insurance company and one consumer. The consumer might have an accident that results in one of a variety of loss amounts, $l \in \{0, 1, 2, \ldots, L\}$. We include 0 in the possible loss amounts to cover the case of no accident. Let e denote the consumer's level of effort to avoid the accident. For each level of loss l and effort level e, $\pi_l(e)$ stands for the probability of this loss occurring. We require that every $\pi_l(e)$ is non-negative and that for each effort level e, $\sum_{l=0}^{L} \pi_l(e) = 1$.

Assume that $e \in \{0, 1\}$; there are only two levels of effort. This is not needed for the analysis, which reaches similar results when we allow more effort levels, but it makes our life a bit easier. We make an important assumption:

Assumption 2.1 (Monotone Likelihood Ratio)**.** *For every $l, l' \in \{0, 1, \ldots, L\}$, if $l > l'$, then*

$$\frac{\pi_l(0)}{\pi_l(1)} > \frac{\pi_{l'}(0)}{\pi_{l'}(1)}.$$

This says that the probability of a worse accident gets relatively higher when less effort is exerted to avoid accidents. So when we see that a bad accident has occurred, we think that the chance that effort was low is pretty high; but we cannot be *sure* that effort was low.

[MATERIAL OMITTED]

2.1 Symmetric Information

We can only understand the impact of asymmetric information by using the unrealistic case of symmetric information as a benchmark.

In this case, the policy can depend on the level of effort. So it can be such that it only pays benefits when a certain level of effort was exerted. The insurance company wants to solve the following maximization problem.

$$\max_{e,p,B_0,B_1,\ldots,B_L} \quad p - \sum_{l=0}^{L} \pi_l(e) B_l, \quad \text{subject to} \tag{2}$$
$$\sum_{l=0}^{L} \pi_l(e) u(w - p - l + B_l) - d(e) \geq \overline{u}.$$

The constraint is called the *participation constraint.*

Note that the choice variable e is discrete, so we cannot take a derivative with respect to it. The solution is to solve the problem twice, once taking $e = 0$ and once taking $e = 1$ and then compare the maximized profits to see the real maximum. For now, then, let us take e as fixed and we will return to consider what happens when it takes its two possible values.

The Lagrangian function for this problem with e fixed is

$$\mathscr{L}(p, B_0, B_1, \ldots, B_L; \lambda) = p - \sum_{l=0}^{L} \pi_l(e) B_l + \lambda \left[\sum_{l=0}^{L} \pi_l(e) u(w - p - l + B_l) - d(e) - \overline{u} \right].$$

The first order conditions are

$$D_p \mathscr{L}(p, B_0, B_1, \ldots, B_L; \lambda) = \quad 1 - \lambda \left[\sum_{l=0}^{L} \pi_l(e) Du(w - p - l + B_l) \right] = 0, \tag{3}$$

$$\forall l,\ D_{B_l} \mathscr{L}(p, B_0, B_1, \ldots, B_L; \lambda) = \quad -\pi_l(e) + \lambda \pi_l(e) Du(w - p - l + B_l) = 0, \tag{4}$$

$$D_\lambda \mathscr{L}(p, B_0, B_1, \ldots, B_L; \lambda) = \quad \sum_{l=0}^{L} \pi_l(e) u(w - p - l + B_l) - d(e) - \overline{u} \geq 0, \tag{5}$$

$$\lambda \geq 0, \tag{6}$$

$$\lambda D_\lambda \mathscr{L}(p, B_0, B_1, \ldots, B_L; \lambda) = \quad \lambda \left[\sum_{l=0}^{L} \pi_l(e) u(w - p - l + B_l) - d(e) - \overline{u} \right] = 0. \tag{7}$$

The first condition follows from the second condition, summed up for all l. Therefore we have at most $(L+2)$ equations in $(L+3)$ unknowns.

From the second condition, we get that $\lambda > 0$ and also that, for all l,

$$Du(w - p - l + B_l) = \frac{1}{\lambda}.$$

This means that the argument on the left hand side is constant as l varies, so for every l we have that $B_l - l$ is constant.

From $\lambda > 0$ we also get that the constraint binds, that is, $\sum_{l=0}^{L} \pi_l(e) u(w - p - l + B_l) - d(e) - \overline{u} = 0$. But in the sum of this equation the term $u(w - p - l + B_l)$ is constant and can be factored out. When we do this and use the fact that $\sum_{l=0}^{L} \pi_l(e) = 1$, we get

$$u(w - p - l + B_l) = d(e) + \overline{u}. \tag{8}$$

Since we have too many unknowns by at least one, we have to choose one arbitrarily. Fortunately, this is easy, as B_0 can be set equal to zero harmlessly. Any B_0 different from zero can be washed out by an appropriate change of p, anyway. But setting $B_0 = 0$ also means that $B_l = l$ for each other value of l, as all $B_l - l$ values have to be equal to each other. So we conclude that it is optimal for the insurance firm to offer full insurance, which makes intuitive sense as the consumer is strictly risk averse and the insurance company is risk neutral.

Using (8), have

$$u(w - p(e)) = d(e) + \overline{u}. \tag{9}$$

(We write $p(e)$ to emphasize that p is determined implicitly by this equation.)

But now it is easy to finish the profit maximization exercise. Simply find $p(1)$ and $p(0)$ from (9) and then choose the larger one among the amounts $p(0) - \sum_{l=0}^{L} \pi_l(0) l$ and $p(1) - \sum_{l=0}^{L} \pi_l(1) l$.

Note that because of the full insurance involved in this equilibrium, the resulting contract leads to a Pareto efficient outcome.

[MATERIAL OMITTED]