Insurance and Public Pensions : (b) Adverse Selection

Adverse selection is said to occur if potential buyers of insurance know their own probabilities of loss better than do insurance companies. So suppose that people are identical in every respect, except that some people are more likely to suffer a loss than other people. If some people have a 10 percent chance of suffering a loss, and others have a 30 percent chance, what are the **actuarially fair** policies? If a policy requires a premium X if there is no loss, and pays out (net of premiums) Y if the policy holder suffers a loss, then the expected profit from the policy is $0.7X - 0.3Y$ if a "high–risk" person buys the policy, and $0.9X - 0.1Y$ if a low–risk person buys. If there were equal numbers of both types of buyer, then the policy would make a profit of $(0.5)(0.7X - 0.3Y) + (0.5)(0.9X - 0.1Y) = (0.8X - 0.2Y$ if all the people bought the policy : this is just the average of the profit on high–risk customers and the profit on low–risk customers.

Certainly a firm will make more money from selling a policy to a low–risk buyer than from selling to a high–risk buyer. In fact, if $0.9X - 0.1Y > 0 > 0.7X - 03.Y$ in the example, a firm would lose money (on average) if a high–risk customer bought the policy, but would make an expected profit if a low–risk customer bought.

Of course, if firms could identify who are the high–risk buyers, and who were the low–risk buyers, then they would choose to sell this policy only to the low–risk buyers. But the problem of adverse selection arises when firms **cannot** identify the different risk classes. If the firms can identify each risk type, then there is no adverse selection problem : firm will sell to different groups at different prices.

The problem with adverse selection, when there is perfect competition in private insurance markets, is that firms will try to design policies to attract low–risk customers. This is possible, even if the firms cannot identify directly which customers are which, because low–risk customers have different preferences over insurance policies than do high–risk customers. Take any insurance policy, which will give rise to an income combination (y_G, y_B). That is, if the person had an initial income of y_0, faced the possibility of loss of L, could buy a policy which reimbursed her Y in the event of a loss, and had to pay X in premia for the policy, then

$$y_G = y_0 - X$$

$$y_B = y_0 - L + Y - X$$

Even if everyone has the same utility–of–income function $u(\cdot)$, the shape of their indifference curves through (y_G, y_B) will not be the same. The low–risk people will have indifference curves with slope

$$-\frac{0.9}{0.1}\frac{u'(y_G)}{u'(y_B)}$$

and the high–risk people will have indifference curves with slope

$$-\frac{0.7}{0.3}\frac{u'(y_G)}{u'(y_B)}$$

in the example. In general, the indifference curves of the low–risk people through any income combination will be steeper than those of the high–risk people (when y_G is graphed along the horizontal, and y_B along the vertical).

Insurance companies, even if they cannot observe directly pople's risk class, can try and do this in an indirect manner. What if every insurance company in the market is offering a policy which requires premia of X, and which pays out Y in case of a loss? If the policy is sold to both high–risk and low–risk people, then what a company would like to do is to introduce a new policy $(\hat{X}, \hat{Y})$ which attracts only the low–risk people, who are much more profitable customers.

The insurance company cannot simply offer the new policy $(\hat{X}, \hat{Y})$ only to low–risk customers. It cannot identify who the low–risk customers are. [1] But what if some company introduced a new policy $(\hat{X}, \hat{Y})$, such that every low–risk person preferred this new policy to the existing policy (X, Y) in the market, and such that every high–risk person would prefer to stick with the existing market policy (X, Y)? Then the insurance company would not have to identify its customers directly. It could offer the new policy $(\hat{X}, \hat{Y})$ to everyone: but only the low–risk customers would prefer to buy the new policy.

What kind of policy would work in this way? In order for the low–risk people to prefer the new policy, it must be the case that

$$(0.9)u(y_0 - \hat{X}) + (0.1)u(y_0 - L + \hat{Y}) > (0.9)u(y_0 - X) + (0.1)u(y_0 - L + Y) \tag{1}$$

If the high–risk people are unwilling to move to the new policy, then it must be the case that

$$(0.7)u(y_0 - \hat{X}) + (0.3)u(y_0 - L + \hat{Y}) < (0.7)u(y_0 - X) + (0.3)u(y_0 - L + Y) \tag{2}$$

Now clearly, either $X > \hat{X}$ and $Y > \hat{Y}$, or $X < \hat{X}$ and $Y < \hat{Y}$. The new policy must offer either lower premia, or a higher payout if it's going to attract any customers. It can't offer both lower premia and a higher payout, because then it would attract all the customers, high–risk as well as low–risk.

So which is it, lower premia or higher coverage? If both equations (1), and equation (2) hold, then (subtract one equation from the other)

$$(0.2)[u(y_0 - \hat{X}) - u(y_0 - X)] > (0.2)[u(y_0 - L + Y) - u(y_0 - L + \hat{Y})] \tag{3}$$

Equation (3) says that the new policy $(\hat{X}, \hat{Y})$ will succeed in attracting — only — the low–risk buyers away from the existing policy only if it offers lowers premia, in exchange for less coverage.

[1] Of course, insurance companies would like to be able to identify customers directly. That is why they offer policies which give a discount on life insurance to non–smokers, or a discount on fire insurance to people with smoke detectors. But the whole problem of adverse selection arises (only) when there are some risk factors which cannot be identified directly. The assumption underlying this whole sub–section is that sometimes there is just no way in which different risk classes can be identified directly.

This makes sense : the low–risk people are less likely to suffer a loss. So other things equal, they are the ones who are more willing to give up a little coverage in exchange for a lower premium. Put otherwise, a high level of coverage is the sort of policy most likely to appeal to high–risk buyers, since they are more likely to have to file a claim.

If a policy actually attracts both types of buyer, then it is usually called a **pooling** policy. There is risk–pooling going on. If the policy breaks even on average, and attracts both types of buyer, then the low–risk buyers are subsidizing the high–risk buyers.

When people's risk types cannopt be observed directly, then an insurance contract (X, Y) would be a **pooling equilibrium** contract, if it attracted all buyers, high– and low–risk, and it was actuarially fair on average, and if there were no other contract being chosen. Of course, as described above, each company would like to introduce a new contract $(\hat{X}, \hat{Y})$ to "steal" only the low–risk buyers away from the existing contract (X, Y). Such behaviour by firms is sometimes called "cream skimming" : attempting to steal the good buyers but not the bad.

If such cream skimming succeeds, it would tend to upset any pooling equilibrium in the market. Companies will break even with a pooling contract only if they are selling to both low– and high–risk buyers : the profits from low–risk buyers offset the losses on high–risk buyers.

In the presence of adverse selection, there **cannot** be a "pooling" equilibrium, in which every person (low or high risk) buys the same insurance policy. Figure 3 illustrates why not. If there were a pooling equilibrium, in which the same policy were chosen by every person, then the policy would have to make zero profits in equilibrium, if there are many competitive firms. Why? Because if a firm's policy made positive profits, then some other firm would undercut it slightly in order to steal all the customers. (If the policy is making positive profits, on average, then a competitior would not have to worry about attracting only low–risk buyers. Attracting all the buyers,high– and low–risk, would be profitable as well.)

So if all the people were buying the same insurance policy, and if that policy were the result of competition among firms, then the policy would make zero profits. That mean it would have to lie on the "zero–profit pooling" line : the set of policies which break even when bought by all the people. In the example, in which half the people had a loss probability of 0.3 and half the people had a loss probability of 0.1, that would be a line with slope $-(0.8)/(0.2)$.

In figure 3, the heavy (pink) dot is a policy which breaks even when all people buy the policy. No firm would want to undercut the policy, and attract **all** the buyers, because they would lose money by doing so.

But they would only lose money if the new policy attracted all the buyers. What if some firm cleverly designed a policy which attracted only the low–risk customers, leaving existing insurance firms with only the high–risk customers? How would the clever new firm do this? Low–risk customers have steeper indifference curves. That means they want to buy less insurance than do high–risk customers. So if a firm offers a new policy which low–risk buyers prefer to the old policy, but which high–risk buyers think is worse than the old policy, then the new policy will attract only low–risk buyers. These are the buyers that firms want to attract, since they are more profitable.

In other words, the smart thing to do, if (X, Y) were a pooling equilibrium, is to introduce the sort of policy $(\hat{X}, \hat{Y})$ described above, which is more attractive to low–risk buyers than (X, Y), but less attractive to high–risk buyers.

The policy marked by the (brown) square in figure 3 is such a policy. It is above the (pink) indifference curve of the low–risk people through the old policy, but it's below the less steep (black) indifference curve of the high–risk buyers. Low–risk buyers would switch from the old policy, the one marked by the (pink) circle, to the new one, but high–risk buyers won't. Notice that the new policy yields an income combination below and to the right of the old : it's offering less coverage, but at a lower price per dollar of coverage.

That means that the firm which introduced the new policy will make a profit. The new policy's income combination is below the zero–profit line for the low–risk types, so it's not actuarially fair, and makes a profit off them. The new policy is above the zero–profit pooling line, so it makes a profit **only** because it does not attract the high–risk customers.

Also, the introduction of the new policy $(\hat{X}, \hat{Y})$ now causes the old policy (X, Y) to go from a break–even proposition to a money loser. The old policy made zero profit if it was chosen by **all** the buyers, low and high risk. If the low–risk buyers abandon the old policy, firms offering it will lose money on average with the remaining high–risk customers.

What figure 3, and the above argument shows, is that there can be **no** market equilibrium in which everyone buys the same policy — at least not when there is adverse selection and when private insurance firms behave competitively. To repeat why : any "pooling equilibrium" policy would have to break even, and if a policy breaks even when everyone buys it, then some firm can steal away the more profitable customers with some other policy.

So what does happen in competitive equilibrium? One possible equilibrium is illustrated in figure 4. In that equilibrium,**two** different policies are offered. One policies, marked by the circle, is chosen by the high–risk people. This policy offers **full insurance**. It offers $Y = L$, so that the resulting income profile (y_G, y_B) for buyers is on the 45–degree line in the diagram. Competition among firms has also driven profits to zero, so this policy is actuarially fair for high–risk buyers. The other policy is the one chosen by low–risk buyers. It offers incomplete insurance ; it's the square in the diagram. In the diagram, the indifference curve for the high–risk types suggests that they are indifferent between the two policies. But they should prefer — ever so slightly — the circle. The policy chosen by the low risk buyers also breaks even : it's on the zero–profit line for the low–risk people.

If there is any market equilibrium under adverse selection and competition, it's got to be the one in figure 4. But it some cases, if there are not very many high–risk people, there may actually be no equilibrium at all. Some clever firm could introduce a new policy which would steal **all** the customers away from the circle and the square, and still make a profit. This new policy would be a pooling policy. So in turn, it would be vulnerable to cream–skimming, and so on.

That is, if profit maximizing firms behave competitively, the outcome depends on the number of people in each risk class. If the number of high–risk people is relatively high, then there will be a

market equilibrium, in which two different policies are sold in equilibrium. One offers full coverage at a high price, and attracts only high–risk buyers. The other offers incomplete coverage, but at a low price. This second policy attracts only low–risk buyers. Given the type of people who choose to buy the policies, each breaks even on average.

If the number of high–risk buyers is relatively low, however, then there will be no market equilibrium at all. That is, there is no set of contracts which firms can offer, for which, both : (i) all contracts offered break even, on average ; (ii) no firm can make a profit by introducing a new, competing contract.

What are the consequences of the adverse selection? If firms knew people's risk probabilities, then there would be a competitive equilibrium in which each person got full insurance at actuarially fair odds. So — if there is a market equilibrium — the asymmetric information hurts the low–risk people, since they now have to settle for incomplete insurance (at actuarially fair rates).

Could the government improve the situation? Of course, it could change things if it had better information than firms, but that does not seem plausible.

However, in this case, it is competitive behaviour by firms which lead to pooling contracts not being sustainable. If the insurance is publicly provided, the government has effectively banned competitive behaviour. The government could offer a contract which offered full insurance, and which broke even over the whole population. That is, it could offer — all — buyers a contract which gave them an income profile at the intersection of the pooling line with the 45–degree line, in figure 3. Low–risk buyers still subsidize high–risk buyers here.

If there are many high–risk buyers, then the outcome under public compulsory provision of this nature will be better for high–risk people, and worse for low–risk people, than the competitive outcome. In other words, if the government wants to redistribute from low–risk buyers to high–risk buyers, it must limit the contracts which can be offered. Any cross–subsidization (low–risk buyers subsidize high–risk buyers) will collapse if firms are allowed to introduce new contracts on their own — even if they make the new contracts available to all buyers. So, if such redistribution is a government goal, then it must intervene in the insurance market : either by restricting the contracts which can be offered by private firms, or by simply taking over the insurance industry, and offering a single contract to all.

We do observe considerable cross–subsidization in publicly provided insurance. Employees in industries with low layoff rates are subsidizing workers in industries with high layout rates, through the public employment insurance system. Healthy people with small families are subsidizing unhealthy people, and people with large families, through the medical insurance system. To the extent that some of the characteristics determining people's risk class for these types of insurance cannot be verified easily by insurers, this sort of redistribution would not be possible under competitive private provision.

Of course, if there are few high–risk people, then there may be no stable equilibrium outcome at all under private competitive provision, providing an even stronger justification for public provision.