Seminar for Insurance Economics
Wintersemester 2001/02

Insurance Markets: Lecture 5

The Supply of Insurance

1. In the analysis of the supply side of a market, we usually proceed by presenting a model of the firm. We specify its objective function, technology and market constraints, and analyse its supply decision as the outcome of a constrained optimisation problem. We interest ourselves in the details of the firm's technology only to the extent that these are relevant to determining the essential structure of its optimisation problem. For example, in the theory of the profit maximising firm, a lot of attention is paid to the way in which the underlying properties of the firm's production function, which embodies everything we need or want to know about its technology, determine the nature of the firm's cost and profit functions.

2. In much of the economics literature on insurance markets, a much simpler approach is taken. It is just assumed that the market is "competitive", the "production costs" of insurance are zero, and that as a result there is a perfectly elastic supply of insurance at a fair premium. This approach is perhaps justified when the purpose of the model is to analyse issues which would only be unnecessarily complicated by a more complete specification of the supply side of the market, for example the implications of information asymmetries for the existence and optimality of insurance market equilibrium. It clearly will not suffice however when we are concerned with the general analysis of insurance markets as such. In fact, we will try to show that an analysis of the supply of insurance is an interesting matter in its own right, as well as being necessary for discussion of a number of important issues, such as insurance market regulation.

3. There is one type of supply situation in which we do not need a model of the insurance firm. This is when two or more individuals, endowed with *ex ante* risky incomes, simply exchange state contingent income claims with each other. In effect they sell each other insurance. We call this the *pure exchange* model of the insurance market. This is an abstract and, empirically speaking, not particularly common type of market, but it is worth studying. We begin with this. We then go on to develop models of the insurance firm and its supply decision.

The Pure Exchange Model

4. We start with the simplest case. There are two individuals, indexed $i = 1, 2$, with initial incomes y_i. They face identical, independent risks of losses L with probability π There are therefore four possible states of the world with total incomes and probabilities as set out in the following table.

$x_1 = y_1 + y_2$	$x_2 = y_1 + y_2 - L$	$x_3 = y_1 + y_2 - L$	$x_4 = y_1 + y_2 - 2L$
$\theta_1 = (1-\pi)^2$	$\theta_2 = \pi(1-\pi)$	$\theta_3 = (1-\pi)\pi$	$\theta_4 = \pi^2$

Note that aggregate income x varies across states of the world. We say therefore that there is *social risk*. The individuals have utility functions $u_i(y_{is})$, $i = 1, 2, \quad s = 1, .., 4$, which are state independent. Let

$$\begin{aligned} U_i &= (1-\pi)u_i(y_i) + \pi u_i(y_i - L) && (1) \\ &= [(1-\pi)^2 + \pi(1-\pi)]u_i(y_i) + [\pi^2 + \pi(1-\pi)]u_i(y_i - L) && (2) \end{aligned}$$

denote $i's$ expected utility in the absence of trade, and y_{is} his income in state $s = 1, .., 4$ after trade. Note that if we regard these state contingent incomes as goods, we have a perfectly standard Edgeworth general equilibrium model. If the two individuals are rational and exhaust all the possibilities of mutually beneficial exchange, the equilibrium allocation can be found as a solution to the problem:

$$\max_{y_{is}} \sum_{s=1}^{4} \theta_s u_1(y_{1s}) \ \ s.t. \ \sum_{i=1}^{2} y_{is} \leqslant x_s \quad (3)$$

$$\sum_{s=1}^{4} \theta_s u_2(y_{2s}) \geqslant U_2^0 \geqslant U_2 \quad (4)$$

$$\sum_{s=1}^{4} \theta_s u_1(y_{1s}) \geqslant U_1 \quad (5)$$

Assuming an interior solution, the first order conditions include

$$\begin{aligned} \theta_s u_1^{'}(y_{1s}^*) &= \lambda_s \ \ s = 1, .., 4 && (6) \\ \alpha\theta_s u_2^{'}(y_{2s}^*) &= \lambda_s \ \ s = 1, .., 4 && (7) \end{aligned}$$

where λ_s and α are Lagrange multipliers. These yield the standard Pareto efficiency conditions

$$\frac{\theta_s u_1^{'}(y_{1s}^*)}{\theta_t u_1^{'}(y_{1t}^*)} = \frac{\theta_s u_2^{'}(y_{2s}^*)}{\theta_t u_2^{'}(y_{2t}^*)} \quad (8)$$

i.e. equality of marginal rates of substitution between incomes contingent on states s and t. Figure 1 illustrates.

5. If an individual in this model is fully insured, that means that y_{is}^* is the same in each state. Because of the presence of social risk, it is clear that, in an optimal solution, *both* individuals cannot in general be fully insured, because then $y_1^* + y_2^*$ would have to be constant across states. This is only feasible if $y_1^* + y_2^* = y_1 + y_2 - 2L$, but this would then involve throwing away income in states $1, 2$, and 3, which cannot be optimal when $u_i^{'}(y_{is}) > 0$, which we assume. Thus in general the optimum will involve uncertain incomes. However, there is a special case in which one of the individuals will be fully insured, and that is where the other is risk neutral. Thus suppose $u_2^{''}(y_{2s}) = 0$. Then we can write 2's marginal rate of substitution as θ_s/θ_t, and inserting this into the optimality condition gives

$$\frac{u_1^{'}(y_{1s}^*)}{u_1^{'}(y_{1t}^*)} = 1 \tag{9}$$

implying equality of 1's income across states. The risk neutral individual in the optimal solution fully insures the risk averse individual.

6. It was convenient to analyse the pure exchange case as a standard model because then we could apply the standard results. However it is easy to give the equilibrium outcome an insurance interpretation. Take the case where 1's income is constant across states, at y_1^*. This is equivalent to an insurance contract in which 1 pays 2 for sure a premium of $y_1 - y_1^*$, and receives L if and only if the loss occurs (confirm by showing that 1's income is then y_1^* for sure). Note that this does not imply that 1 receives fair insurance. The premium $y_1 - y_1^* \geqslant \pi L$. We know it cannot be lower, because then 2 would not be prepared to trade, but it could be higher, depending on the outcome of the trading process. For example, if 2 is smart enough to get all the gains from trade, we will have $u_1(y_1^*) = U_1$, and 1 receives only his reservation utility. (see figure). In the case where both are risk averse, so that no-one ends up as being fully insured, we could design contracts that realise the optimal incomes by having 1 selling partial cover to 2 against the event of 2's loss, and 2 selling partial cover to 1 against the event of 1's loss. Alternatively, they might simply agree to pool their incomes and make the required payments in each of the contingencies.

7. If the risks faced by the two individuals were perfectly negatively correlated, then total income is constant at $y_1 + y_2 - L$, and so there is no social risk. In that case it is straightforward to show that the trading

equilibrium implies full insurance for both individuals, whether or not both are risk averse. This is left as an exercise.

8. It is interesting to see what determines the equilibrium incomes across states, but this is not too easy in the discrete case. Thus we now generalise the model somewhat. Let $s \in [a, b]$ now be a continuous state of the world variable, and $y_i(s)$ distributions of incomes which are to be chosen given the aggregate endowed income distribution $x(s)$. The probabilty density function on states is $\pi(s)$. The equilibrium of the state contingent income exchange process will now be the solution to

$$\max_{y_i(s)} \int_a^b u_1(y_1(s))\pi(s)ds \;\; s.t \;\; \int_a^b u_2(y_2(s))\pi(s)ds \;\; \geqslant \;\; U_2^0 \tag{10}$$

$$y_1(s) + y_2(s) \;\; = \;\; x(s) \tag{11}$$

For simplicity we omit the other constraints. The first order conditions now imply, much as before

$$u_1^{'}(y_1^*(s)) - \alpha u_2^{'}(y_2^*(s)) = 0 \tag{12}$$

and so we have

$$\alpha = \frac{u_1^{'}(y_1^*(s))}{u_2^{'}(y_2^*(s))} \tag{13}$$

Substituting from the constraint for $y_2^*(s)$ and differentiating through the above first order condition with respect to s gives

$$u_1^{''}\frac{dy_1^*}{ds} - \alpha u_2^{''}(\frac{dx}{ds} - \frac{dy_1^*}{ds}) = 0 \tag{14}$$

where $u_i^{''} \equiv u_i^{''}(y_i^*(s))$ $i = 1, 2$. Then substituting for α and rearranging gives

$$\frac{dy_1^*}{ds} = (\frac{A_2(y_2^*)}{A_2(y_2^*) + A_1(y_1^*)})\frac{dx}{ds} \tag{15}$$

which says that the way in which income is divided across the distribution of states depends on the absolute risk aversion functions of the two individuals. Thus if 2 is risk neutral we have that $A_2(y_2^*) = 0$, and so 1's income is again constant across states, while if $A_1(y_1^*) = 0$, 1 carries all the risks. If both have constant risk aversion then

$$\frac{A_2(y_2^*)}{A_2(y_2^*) + A_1(y_1^*)} = \beta \tag{16}$$

and integration gives

$$y_1^*(s) = \gamma + \beta x(s) \tag{17}$$

There is a linear "sharing rule", with the slope coefficient β determined by the relative values of risk aversion, and the constant of integration, γ, chosen so as to satisy the constraint on the level of 2's utility. It is then easy to insert specific utility functions into this expression and derive their implications for the optimal sharing of risky incomes.

Risk Pooling and Insurance Syndicates

9. The pure exchange model leads to the idea of insurance as a "mutual pooling" activity. We could however think of the formation of a group of individuals, a syndicate, which would sell insurance to non-members and then divide the risks and profits among themselves. The key aspects of the "technology" here are first the **pooling** of risks, in the sense of the aggregation of a number of separate individual risks, and then the **spreading** of the aggregate risk among the members of the syndicate.

10. **Risk Pooling.** Suppose there are n individuals with identically and independently distributed (*i.i.d.*) risks. That is, each faces the same distribution of losses L_i, where these random variables each have finite mean μ and finite variance σ^2, and zero covariance. It follows from the standard properties of such random variables that their sum, $L_n = \sum_{i=1}^{n} L_i$ has mean $n\mu$ and variance $n\sigma^2$. If the insurance syndicate therefore offers full insurance to these n individuals, these are the mean and variance of its distribution of total claims costs. Note that the variance of the total claims cost, which we can think of loosely as an indicator of the overall riskiness of the business, increases proportionately with n.

11. Consider now however the claims cost per contract, i.e. the average loss $\bar{L}_n = \frac{1}{n}\sum_{i=1}^{n} L_i$. This is of course also a random variable. There will be n specific realisations of the random variable L_i, which can be thought of as a random sample of size n drawn from the n individual loss distributions, and $\bar{L}_n$ is just the mean of this random sample. It follows that any one realisation of $\bar{L}_n$ may not equal μ, the true mean of the loss distribution. However, the mean of the distribution of $\bar{L}_n$ is μ, while the variance of $\bar{L}_n$ about this mean is σ^2/n,with the standard deviation therefore $\sigma/\sqrt{n}$. Note that these go to zero as the sample size, or in this case the number of insurance contracts, goes to infinity. More precisely, the *Law of Large Numbers* implies that for

all $k > 0$,

$$\lim_{n\to\infty} \Pr[\left|\bar{L}_n - \mu\right| < k\frac{\sigma}{\sqrt{n}}] = 1 \tag{18}$$

Put loosely, this says that for a sufficiently large number of insurance contracts, it is virtually certain that the loss per contract is just about equal to μ, the mean of the individual loss distribution. As the number of contracts increases, so the probability that the loss per contract lies outside an interval around μ, defined in terms of so-and-so many standard deviations, goes to zero. It follows that, for a sufficiently large number of insurance contracts, the insurance syndicate can base its premium calculations on the assumption that it is virtually certain that the loss per contract is the mean of the loss distribution μ. Thus a premium equal to μ covers, *in expected value*, the loss per contract. In that case premium revenue would be $n\mu$. Of course, if there are costs of running the insurance business, often just called transactions costs, then a loading will have to be added to μ, and, depending on the competitiveness of the market, a profit mark-up may also be added.

12. However, it must not be forgotten that this applies to the loss per contract. It is still the case that the variance of the total claims cost increases with n. Thus, abstracting from transactions costs, though a premium revenue of $n\mu$ covers the expected value of claims costs, in any particular realisation of total claims, the latter may exceed premium revenue, and so the insurance syndicate would be **insolvent** if it did not also have so-called "technical reserves" to cover this eventuality.

13. We return to this subject later, when we look at the issues surrounding the regulation of insurance markets. For the moment, we just pursue the implications of the fact that an insurance syndicate that charges a premium per contract of μ will, ignoring transactions costs, break even in expected value, but will have an uncertain net income

$$Y_n = n\mu - L_n \tag{19}$$

If the syndicate members are risk neutral, this makes no difference, but if they are risk averse, then we may feel that it ought to. At this point the question of the risk-spreading aspect of an insurance syndicate becomes relevant.

14. **The Arrow-Lind Theorem.**