

3. Insurance Supply

Till Stowasser

University of Stirling

Insurance Economics (LMU, 2024)

Simple approach

Most of the economics literature on insurance markets takes a very simple approach:

- It is assumed that the market is “competitive”.
- The “production costs” of insurance are zero.
- As a result there is a perfectly elastic supply of insurance cover at a fair premium.

The technology of insurance has two aspects:

- Conceptually, insurance can come into existence via the **pooling** and **spreading** of risk.
- Physically, rather than conceptually, activities involved in “producing” insurance matter: drawing up and selling new insurance contracts, administering the stock of existing contracts, processing claims, estimating loss probabilities, calculating premiums, and administering the overall business. We will call these costs **insurance costs**.

- **Risk pooling**

When an insurer enters into insurance contracts with a number of distinct individuals, the probability distribution of the aggregate losses they suffer differs from the loss distribution facing any one individual. We are interested in the distribution of claims on the insurer to which it gives rise.

- **Risk spreading**

The insurer is typically not a single individual, but rather a group of individuals. Each member of this insurance **syndicate** or insurance **company** or **firm** takes on an equal share of the total loss. How does this affect the premium that would be set, given that the individuals may be risk averse.

Risk pooling

- insurer enters into insurance contracts with n individuals
- distribution of claims costs \tilde{C}_i under each contract is identical, and independent across contracts with mean μ and variance σ^2 both finite and with zero covariance between any pair of values $C_i, C_j, i, j = 1, \dots, n, i \neq j$
- assumption of **identically and independently distributed** (*i.i.d.*) risks is not essential but is very helpful in simplifying the technicalities.

- Standard property of the sum of *i.i.d.* random variables: $\tilde{C}_n = \sum_{i=1}^n \tilde{C}_i$ is also a random variable with **mean** $n\mu$ and **variance** $E[(\sum_{i=1}^n \tilde{C}_i - n\mu)^2] = E[\{\sum_{i=1}^n (\tilde{C}_i - \mu)\}^2]$
- since the covariances between the \tilde{C}_i are all zero

$$E[\{\sum_{i=1}^n (\tilde{C}_i - \mu)\}^2] = \sum_{i=1}^n E[(\tilde{C}_i - \mu)^2] = n\sigma^2.$$

- **Note:** The variance of the total claims cost increases linearly with n while its standard deviation, $\sqrt{n}\sigma$, is strictly concave in n .
- **Implication:** If the insurer sets the premium on each contract equal to the expected value of cover or claims cost μ (and insurance costs are zero) it will just break even **in expected value**, since total premium revenue $n\mu$ will equal the expected value of claims costs. **Thus we call μ the fair premium.**

- **Note:** Any one realization of \tilde{C}_n , i.e. actual aggregate claims costs in any one period, may be larger or smaller than $n\mu$ since the variance $n\sigma^2$ is always positive and increases with n .
- To avoid **insolvency**, i.e. the situation in which claims costs exceed the funds available to meet them, the insurer has to carry technical or insurance reserves.
- Assume that each contract has a maximum cover C_{\max}
 \Rightarrow maximum aggregate claims cost nC_{\max} . The insurer can assure a zero probability of insolvency by setting a premium amount P per contract and carrying reserves $R_{\max} = n(C_{\max} - P)$
- **However:** Typically it will be extremely unlikely that claims costs will be in the region of nC_{\max} , while, for a large insurer, attempting to raise a capital of R_{\max} could be extremely costly.

- \Rightarrow insurers choose a so-called **ruin probability** ρ by choosing a level of reserves $R(\rho) = C_\rho - nP$, where C_ρ satisfies

$$\Pr[\tilde{C}_n > C_\rho] = \rho$$

- I.e., reserves are set such that the probability that actual claims costs will exceed premium revenue plus reserves and the insurer will be insolvent is ρ .
- The aggregate loss claims distribution is bounded below by zero and above by nC_{\max} , and C_ρ is the value of aggregate claims such that with probability ρ the insurer will be insolvent.
- For a given value of ρ , the value C_ρ will increase with the number of contracts n . Since μ is independent of n , this means that the value of the required reserves $R(\rho)$ must also increase with n .

- The optimal ruin probability ρ solves the problem of the optimal trade-off between the costs of insolvency and the cost of holding reserves.
- Consider the implications of the *Law of Large Numbers* for the value of the loss and insurance reserves *per contract*.
- Consider a particular realization C_1, C_2, \dots, C_n of the claims under the n individual contracts. We can regard this as a random sample from a distribution with mean μ and variance σ^2 , both finite.

- Let \bar{C}_n denote the sample mean, or average loss per contract, i.e.
$$\bar{C}_n = \frac{1}{n} \sum_{i=1}^n C_i.$$
- Then the version of the Law of Large Numbers (there is more than one) relevant for present purposes says that for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \Pr[|\bar{C}_n - \mu| < \varepsilon] = 1 \quad (1)$$

- As n becomes increasingly large the average loss claim per contract, will be arbitrarily close to the value μ with probability approaching 1.
→ For a sufficiently large number of insurance contracts, it is virtually certain that the loss per contract is 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 arbitrarily small interval around μ goes to zero.

It is also useful to look at the variance of \bar{C}_n given by

$$\begin{aligned} E\left[\left(\frac{1}{n} \sum_{i=1}^n C_i - \mu\right)^2\right] &= E\left[\frac{1}{n^2} \left(\sum_{i=1}^n C_i - n\mu\right)^2\right] \\ &= \frac{1}{n^2} E\left[\left(\sum_{i=1}^n C_i - n\mu\right)^2\right] = \frac{n\sigma^2}{n^2} = \frac{\sigma^2}{n} \end{aligned}$$

Thus the variance of the realized loss per contract about the mean of the individual loss distribution goes to zero as n goes to infinity.

- The Law of Large Numbers statement in (1) becomes, by setting $\varepsilon = \sigma^2/n$
 $\lim_{n \rightarrow \infty} \Pr[|\bar{C}_n - \mu| < \frac{\sigma^2}{n}] = 1.$
→ As the number of individual insurance contracts sold by an insurer becomes very large, the risk that the claims cost per contract will exceed the fair premium becomes vanishingly small.
- We can interpret this as a type of **economy of scope/scale**: although the variance of aggregate claims increases with n , so the insurance reserves will have to increase in *absolute* amount, the required reserve *per contract* converges to zero.

Risk spreading

- Suppose the insurer is either a syndicate with N members or a company with N shareholders. These individuals are all identical and share the net income from the insurance business equally, so that each receives a share $s = 1/N$. We also assume that the individuals are risk averse.
- **Question of interest:** What are the implications of increasing the number of individuals N in the syndicate or company, i.e. of spreading the risky income over a larger number of individuals?
- **Intuition:** This reduces the riskiness of the individual incomes and therefore reduce the risk premium that they demand for taking a share in the insurance, thus reducing the insurance premium. Again this is a type of **economy of scale**.

- Remember the Arrow-Pratt measure of absolute risk aversion.

$$r \approx -\frac{1}{2} \frac{u''(y_0)}{u'(y_0)} \sigma_z^2 \approx \frac{1}{2} A(y_0) \sigma_z^2$$

- r is the individual's risk premium for accepting a small risky income z with zero mean and variance σ_z^2
- y_0 is his income in the absence of this risk

Now suppose this risky income z consists of the share $s = 1/N$ in a given aggregate risky income \tilde{Z} , which has mean zero and variance σ_Z^2 . Then we have

$$\sigma_z^2 = E\left[\left(\frac{\tilde{Z}}{N}\right)^2\right] = \frac{\sigma_Z^2}{N^2}$$

Clearly $\lim_{N \rightarrow \infty} r = 0 \rightarrow$ For a sufficiently large number of individuals, each becomes essentially risk neutral.

- **Moreover:** If we consider Nr , the *sum* of the risk premia

$$Nr \approx \frac{N}{2} A(y_0) \sigma_z^2 = \frac{1}{2} A(y_0) \frac{\sigma_z^2}{N}$$

the total cost of risk bearing to the syndicate as a whole, Nr , goes to zero as the syndicate size grows.

- → for sufficiently large N the risk aversion of the individuals can be ignored, and the insurer can be treated as risk neutral.
- Intuitively, the individual risk premia fall at a rate determined by N^2 , while the sum of risk premia grows at a rate determined by N , and so overall this sum goes to zero as N grows.

The Arrow-Lind Theorem

- This theorem has many applications over and beyond insurance markets, but is also of central importance here.
- It confirms the intuitive idea that the larger the number of syndicate members who share in a given distribution of income from a risky insurance business, the smaller the cost of the risk associated with that business, even though the individual syndicate members are risk averse.
- More importantly, it makes clear a necessary condition for this result:
The covariance between the member's income from the insurance business, and his marginal utility of income if he does not share in this business, must be zero.
- **Note:** If this covariance were positive, implying, since $u'' < 0$, a negative covariance between y and \tilde{Z} , the aggregate value of the insurance business to its shareholders would exceed its expected value, and conversely if the covariance were negative.
- In the former case, the insurance business offers the shareholders a way of diversifying their asset portfolio.

The Raviv Model

We now analyze the implications of introducing **insurance costs** for an insurer supplying insurance in a perfectly competitive market. To focus the analysis on the effect of insurance costs, we assume

- insurance buyers are identical and their relevant characteristics - utility function, income, loss distribution - are fully known to the insurer
- the insurer can be treated as risk neutral
- we can ignore the cost of reserves in pricing individual insurance contracts.
- The loss distribution is $\{0, L_1, L_2, \dots, L_S\}$, with corresponding probabilities $\{\pi_0, \pi_1, \pi_2, \dots, \pi_S\}$, all positive, and with $0 < L_1 < L_2 < \dots < L_S$.
- The premium (amount) is P .

Insurance costs and the Raviv Model

- The insurance buyer's incomes are

$$y_0 = y - P \quad \text{and} \quad y_s = y - P - L_s + C_s \quad \text{for } s = 1, \dots, S$$

where $C_s \geq 0$ is cover in state s .

- The insurer's profits on any one contract in the respective states are

$$x_0 = P - F \quad \text{and} \quad x_s = P - C_s - K(C_s) \quad \text{for } s = 1, \dots, S$$

- The **insurance cost function** $K(C_s)$ has $K'(\cdot) \geq 0$, $K''(\cdot) \geq 0$, and $K(0) = F \geq 0$, a fixed cost.

Note: This cost function relates to one individual insurance contract. The model is assuming that the cost of having n contracts is just n times the cost of one contract.

That the insurance market is perfectly competitive implies two things:

- the insurer will make zero profits in expected value, $\sum_{s=0}^S \pi_s x_s = 0$, since otherwise entry or exit of risk neutral insurers will take place. In other words, the contract we derive is a long run equilibrium contract
- this equilibrium contract must also maximise the expected utility $EU = \sum_{s=0}^S \pi_s u(y_s)$ of the insurance buyer since, if not, a competitor could offer a superior contract and compete away the business

Summary of results

- A buyer will choose full cover if offered a fair premium, while otherwise he prefers a contract with a deductible over all other types of contract with the same expected cost to the insurer.
⇒ application of SOSD
- The existence of insurance costs suggests that a fair premium will not be feasible, and thus, given the competitive market assumption, contracts with a deductible are likely to make an appearance.

Summary of results

- If $K'(\cdot) = 0$, so that the insurance costs take the form of a fixed cost per contract, $K(\cdot) = F > 0$, there is full cover, implying that the insurer offers a premium $P = F + \sum_{s=1}^S \pi_s C_s$.
Thus *at the margin* the premium is fair, inducing the buyer to choose full cover, and the insurer covers its costs by making a lump sum charge in addition (this is known as a two-part tariff).
- With $K'(\cdot) > 0$ we find that the optimal contract will give partial cover, and is likely to involve a deductible.
- If $K''(\cdot) = 0$ there is simply a deductible.
- With $K''(\cdot) > 0$ the optimal contract involves *coinsurance above a deductible*, or an increasing gap between loss and cover