CHAPTER 4

CHAPTER 4

PORTFOLIO THEORY AND RISK MANAGEMENT

The essence of portfolio theory is captured in such well-known sayings as "Don't put all your eggs in one basket" or "You will gain on the swings what you lose on the roundabouts." The intuitive logic of these statements warns us not to bet all our money on one horse, nor to purchase a single stock with all our capital. It tells us that single-engined aircraft are inherently more dangerous than multi-engined aircraft and informs us that we should take a large (random) sample if we wish to measure popular support for competing presidential nominees. Diversification, it seems, helps us to avoid, or at least minimize, the probability of extreme outcomes. The same mechanism helps to explain how insurance functions and can be put to work to identify strategies for corporate risk management. More formally, these mechanisms have been analyzed under the heading of portfolio theory.

The growth and evolution of portfolio theory can be illustrated in three subject areas. In statistics, the properties of samples and their relationship to the population from which they were drawn have been intensively studied. The properties of the sample, which is simply a portfolio of observations drawn from a parent population, enable us to make statements about the population and to evaluate the degree of confidence we can place in such statements. For example, if we choose a suitable large sample of female Americans in a random fashion, not only can we estimate the average height of all female Americans, but we can also place a confidence interval around our estimate. This permits us to make a statement such as "The estimated average height of American females is 5 feet 4 inches plus or minus inches at the 95% confidence level". Subject to proper sampling techniques, the larger the sample, the more we can narrow the range of error. More loosely speaking, the larger the portfolio, the more we cut down the risk of being wrong.

A second area in which portfolio theory has been developed and successfully put to work is security analysis. The rate of return on securities of all kinds is rarely, if ever, free of risk. In the case of common stocks, the variance of the return is usually fairly large. Thus, it is not uncommon for a stock to lose or gain a large proportion of its initial value over a short period. The investment risk can be partially controlled by a number of securities. Over a given period, some security prices may fall while others rise. Even this diversified stock portfolio can suffer a large loss in value if there is common movement in stock prices. Most stocks do tend to track movements in a representative price index (e.g., the Dow Jones Index or the Standard & Poors Index) to a greater or lesser degree, and this element of risk is not amenable to simple diversification. The common movement or correlation between such variables as stock prices is of concern to us and will be considered in some detail presently. Moreover, the analysis of security markets is also important to our treatment of risk management. Investor behavior determines the value of a firm. Since our financial approach to risk management requires that decisions be appraised in terms of their contribution to the value of a firm, we must set our decision criteria with reference to investor behavior.

The third area that has provided a fertile field for the application of portfolio theory is actuarial science. Its roots lie in the study of mortality rates, typically classified by age and sex groupings the population. For large samples of individuals, such as portfolios of insureds underwritten by life insurance companies, the actual mortality rates usually show little deviation from the expected rates calculated from large banks of mortality data. While the insurance company knows little of the time and place of death for any individual, its prediction of average mortality is fairly accurate. Consequently, it can budget for claims liabilities with reasonable confidence. The life insurer has diversified and has accordingly reduced risk. More recently, attention has been paid to the portfolio properties of other forms of insurance, and a sophisticated branch of mathematics, risk theory, has been developed. The initial focus of risk theory was on the probability (or risk) that aggregate claims liability will exceed the insurer's reserves, although risk theory now encompasses a wider study of the risk properties of a diversified insurance portfolio.

The study of diversification within an insurance portfolio provides a useful starting point for our analysis of portfolio theory in risk management. However, the terms of reference of risk management go beyond the study of insurance. The same basic relationships that explain risk spreading in an insurance fund can be used to show how the individual exposure units under a corporate umbrella combine to form the aggregate loss distribution for the firm. Knowledge of the properties of the corporate portfolio of risk exposures is essential to the formation of proper risk management strategy.

DESCRIPTION OF AN INSURANCE PORTFOLIO

In an insurance fund, the basic units that are insured are usually called exposure units. These are not synonymous with the policies in an insurance portfolio. A policy might provide insurance coverage on several automobiles, but for practical purposes, each automobile is a unit of coverage. Each vehicle is insured along with many other vehicles, and each is assumed to introduce its own risk properties to the portfolio. It would be convenient to define exposure units as independently exposed to the prospects of damage or loss by the perils covered. Since each is independent of the other, it may be introduced to the portfolio with predictable and acceptable risk effects. However, the assumption of independence is too restrictive. Instead, we will think of the exposure units as being separately exposed to the possibility of loss or damage. Thus individual houses may be considered exposure units within a homeowners' insurance portfolio. However, if two homes are condominiums separated by a non-fireproof party wall, it may be convenient to consider these as a single risk from a fire insurance viewpoint, since a single fire could easily destroy both homes. A home a block or two away may be considered as a different exposure unit, even though common factors could jointly affect the probability of loss in each (examples are a local arsonist, local weather conditions, construction hazards since both were erected by the same contractor, etc.). The relevant criterion is the degree of independence, although no hard and fast rule can be applied.

Each exposure unit insured by an insurance company represents a liability of the insurer for potential claims that may arise. At the inception of the policy, the value of this liability is not known, but over time, its value will be revealed. For most policies, the eventual liability turns out to be zero, since no loss arises and therefore no claim is filed. However, for some policies, a loss does arise, and the insurer will be faced with a set of claims that will vary in size according to the intensity with which the peril struck, the inherent protection, the value at risk, etc. If we include the “zero” cases, we can represent the final set of outcomes for the insurer having “n” exposure units as:

$$L_1 \; ; L_2 \; ; L_3 \; .... L_n$$

where $L_i \geq 0$ is the revealed liability of the insurer under exposure unit i. The total liability of the insurer is therefore:

$$L = L_1 + L_2 + L_3 + ... + L_n = \sum_{i=1}^{n} L_i$$

For our purposes, it is useful to think of the insurer's liability on a per risk basis. This is determined by dividing L by the number of exposure units n:

$$\frac{L}{n} = \frac{1}{n}\sum_{i=1}^{n} L_i$$

This will be called the mean or average loss, and we will refer to the distribution of average loss. Most policies run for a year, but the value of L may not be known for some considerable time after the expiration of the policies. Liability claims in particular can take an extremely long time to settle, and it may be many years before final values are known.[1] At the time of underwriting the policies, the insurer knows none of the outcomes L_i. Each liability is a random variable that may assume one of a range of values.[2]

The professional interest of the insurer focuses on the properties of the variable L/n. The component values L_i may be of intense interest to the policyholders, but the interest of the insurer arises solely from the insurance contract. The financial performance of the insurer rests on the distribution of average claim payments. However, the distribution of L/n is determined by the distributions of L_i. Schematically, this interdependence may be represented as shown in Figure 1. The exposure units are represented on the top row by a set of probability distributions from which the values L_1 , L_2 , ... L_n will eventually be revealed. The probability distribution for losses on individual exposure units will generally be skewed to the right. However, their shapes may differ considerably according to size, value, degree of peril, and related factors. In combination, these distributions form the average distribution L/n, as shown by the arrows. The process of aggregation accomplishes one very interesting and useful result if the individual exposure units are statistically independent of each other. We will assume for the moment that this condition strictly holds. The very useful property that emerges is that the distribution of the mean loss value will approach a normal distribution. The result is stated by the Central Limit theorem.

---

[1]Recent litigation on asbestos claims suggests that policies written 30 or more years ago may still carry liabilities for cases of asbestos related disease that are currently being diagnosed.

[2]An alternative notation is to use tildes, ˜ , to denote random variables. Since most of the discussion in this book relates to random variables, we can conveniently omit the tilde, taking care to point out particular cases in which we refer to a realized value from a random distribution.

The *Central Limit theorem* states that

> The distribution of the mean value of a set of "n" independent and identically distributed random variables each having mean μ and variance σ^2 approaches a normal distribution with mean μ and variance σ^2/n as n tends toward infinity.

To apply this theorem to insurance, consider an insurer writing "n" automobile policies in some homogeneous rating category, e.g., male urban drivers over 25. These policies can be viewed as a sample of the total population of drivers in this category, the others being insured by other companies or self-insuring. We assume that the risks are sufficiently similar that they can be seen as identically distributed, each having the same expected loss $E(L_i)$ and the same variance $\sigma^2(L_i)$. Our task is to find out what happens to the riskiness of the insurer's portfolio as the insurer increases the number of risks it insures. To undertake this task we consider two different definitions of risk:

1. The first definition consists of the variance or standard deviation of the distribution of $\sum L_i/n$. This is a conventional measure of the spread of the distribution.

2. The second definition consists of the probability that $\sum L_i/n$ exceeds some critical value. The critical value can be set at the level of the insurer's total reserves plus surplus averaged over all policies. If average losses turn out to be below this value, the insurer is able to discharge its claims liabilities. However, should average losses exceed this value, the firm will fail in its obligations and will be insolvent. The probability that such an eventuality will arise is known as the *probability of ruin*, and this is considered to be an important risk measure for the insurer.

In order to examine how these risk measures are affected by the size of the insurance fund, it is necessary to define the relationships between the individual loss distribution and the average distribution. Consider a portfolio formed just by adding two random variables L_1 and L_2. The aggregate distribution will have the following properties:

$$E(L) = E(L_1) + E(L_2)$$

and

$$\sigma^2(L) = \sigma_1^{\,2} + \sigma_2^{\,2} + 2\,\sigma_{1,2}$$

where $\sigma^2(L)$ = the variance of the portfolio L
σ_i^2 = the variance of L_i
$\sigma_{1,2}$ = the covariance of L_1 and L_2

Although variance is useful, the more useful summary description of risk is standard deviation, which is simply the square root of the variance:

$$\sigma(L) \quad = \quad \sqrt{\sigma^2(L)}$$

Covariance is a measure of the common variations between L_1 and L_2 and is defined by

$$\sum p_i \, (L_{1i} - E(L_1))(L_{2i} - E(L_2))$$

where p_i is the probability of the joint occurrence of L_1 and L_2.

Examination of covariance reveals that it looks very much like the formula for variance, which is

$$\sum p_i \, (L_{1i} - E(L_1))(L_{1i} - E(L_1)) \quad = \quad \sum p_i \, (L_{1i} - E(L_1))^2$$

In effect, variance is simply covariance of one variable with itself.

If there are more than two variables, the formulas for mean and variance are

$$E(L) = E(L_1) + E(L_2) + E(L_3) + \ldots\ E(L_n) = \sum_i E(L_i)$$

and

$$\sigma^2(L) \quad = \quad \sigma_1^2 + \sigma_2^2 + \sigma_3^2 + \ldots\ldots\ \sigma_n^2 + \ldots\ldots 2\sigma_{1,2} + 2\sigma_{1,3} + \ldots\ldots\ldots$$
$$= \quad \sum_{i=1}^{n} \sigma_i^2 \quad + \quad \sum_i \sum_{j \neq i} \sigma_{i,j}$$

Our first task is to calculate the risk in a portfolio of identically distributed independent exposure units. Since the risks are independent, the covariances are equal to zero. The variance of the distribution of average loss is therefore

$$\sigma^2\left(\frac{L}{n}\right) = var\left(\frac{L_1 + L_2 + L_3 + \ldots\ldots + L_n}{n}\right)$$

$$\left(\frac{L_1}{n}\right) + var\left(\frac{L_2}{n}\right) + \ldots\ldots + var\left(\frac{L_n}{n}\right) = \sigma^2\left(\frac{L_1}{n}\right) + \sigma^2\left(\frac{L_2}{n}\right) + \ldots\ldots +$$

since the covariances are zero. Now,

$$\sigma^2\left(\frac{L_i}{n}\right) = \sum p_i \left(\frac{1}{n}L_i - \frac{1}{n}E(L_i)\right)^2$$

$$= \frac{1}{n^2}\sum p_i \left(L_i - E(L_i)\right)^2 = \frac{\sigma_i^2}{n^2}$$

Thus,

$$\sigma^2\left(\frac{L}{n}\right) = \frac{1}{n^2}\left(\sigma^2\ (L_1) + \sigma^2(L_2) + \ldots\ldots + \sigma^2(L_n)\right)$$

Since all risks are identically distributed, they all have the same variance, that is,

$$\sigma_1^2 = \sigma_2^2 = \ldots\ldots\sigma_n^2$$

Therefore:

$$\sigma^2\left(\frac{L}{n}\right) = \frac{1}{n^2} n\sigma_i^2 = \frac{\sigma_i^2}{n}$$

This is the result stated by the Central Limit theorem.

To put some meat on these bones, let our automobile portfolio with male urban drivers over 25 comprise exposure units that each have an expected loss of \$500 and a standard deviation of \$800. The variance of each unit is $\sigma_i^2 = (\$800)^2 =$ \$640,000. The variances and standard deviations of the insurer's liability per exposure unit are given in Table 1.

TABLE 1
RISK REDUCTION IN PORTFOLIOS OF INDEPENDENT RISKS

n	$\sigma^2(L/n)$	$\sigma(L/n)$
1	640,000	800
10	64,000	253
100	6,400	80
1,000	640	25
10,000	64	8
∞	0	0

In general, the effect of portfolio size on portfolio risk as shown in Figure 2. As "n" gets larger and larger, the risk of average loss gets smaller and smaller eventually approaching zero.

Diversification does benefit the insurer because it reduces risk. However, a word of caution is needed. It is the measure of risk per policy that is reduced. The standard deviation of the absolute dollar liability of the insurer increases. As a simple exercise, consider an increase in size of the insurer's portfolio from 100 to 1000 exposure units. The standard deviations of absolute dollar liability before and after the increase in portfolio size are, respectively

$$\sigma_p = \sqrt{100(800)^2} = 8000 \qquad \textit{when} \quad n = 100$$

$$\sigma_p = \sqrt{1000(800)^2} = 25{,}298 \qquad \textit{when} \quad n = 1{,}000$$

However, it will be noticed that the standard deviation has increased at a much smaller proportionate rate (just over 3 times) than has the portfolio size (10 times). This less-than-proportionate increase in aggregate total risk explains why the riskiness per policy falls.

Our result on diversification can be summarized graphically. Figure 3 shows possible distributions of L/n at different portfolio sizes. An "n" increases, not only does the distribution exhibit a smaller standard deviation, but it becomes more symmetrical, approaching a normal distribution. This last feature becomes important for our examination of the probability of ruin. Figure 3 reveals that as n increases, the distribution of L/n huddles closer and closer around the mean value μ. This property is known as the *law of large numbers.*

The second measure of riskiness we wish to examine is the probability that the distribution of the average loss will exceed the financial resources of the insurer. This is the so-called *probability of ruin*. When the insurer receives the premium on a policy, it is allocated to a reserve fund for unearned premiums (after deduction of expenses). Funds in this reserve are held to match liabilities developing on the maturing policies. At the end of the accounting period, funds may be reallocated to other reserve funds, such as outstanding claims and claims incurred but not reported. Procedures also govern the allocation of any residual values to surplus, which can be either distributed to shareholders or retained. Thus, an insurer can meet aggregate claims up to the total value of its reserves plus surplus. If this total value is denoted L*, the insurer will be able to pay a maximum of L*/n per policy. We wish to know what is the probability that the average claims liability will exceed L*/n per policy. This is the probability of ruin. Figure 3 depicts such a critical value. The probability of ruin is depicted by the area of the tail to the right of L*/n. Apparently, this is systematically related to the size of the portfolio, although it behooves us to demonstrate this effect.

As a rule of thumb, when a portfolio is formed of more than 30 independent random variables, the distribution of the mean value is considered to be sufficiently close to normal that statistical tests based on the normal distribution can be used. The normal distribution has the convenient property that the total probability (or area under the curve) can be segmented symmetrically if we know the mean and the standard deviation. We know, for example, that the probability that a randomly selected value will be above $\mu_p+\sigma_p$ (that is, one standard deviation above the mean) is 0.1587, as shown in Figure 4. Similarly, the probability that a randomly selected value will be below $\mu_p - \sigma_p$ (one standard deviation below the mean). is also 0.1587, since the distribution is symmetric. Consequently, there is a 1 - 0.1587 - 0.1587 = 0.6826 probability that a randomly selected value will lie within the range of one standard deviation on either side of the mean. We can work out the probabilities in any other segments, and they are described in terms of multiples of standard deviations from the mean. Thus, the probability that a randomly selected value is greater than (or less than) μ plus some multiple of σ is as follows:

Probability that L > $\mu + (1.5)\sigma_p$ = 0.0668
$\mu + (1.96)\sigma_p$ = 0.025
$\mu + (2.33)\sigma_p$ = 0.01 and so forth

These, and other, values are conveniently tabulated in "z" tables, which are available in many statistics textbooks. A segment of such a table is reproduced at the end of this text. Given this pattern, and the tendency of the portfolio toward a normal distribution, we can easily calculate the probability that the insurer's average claim liability exceeds the ruin value L*/n simply by defining this value as a multiple of "z" standard deviations from the mean. The value L*/n will be defined to be z standard deviations above the mean:

$$\frac{L^*}{n} = \mu_p + z\sigma_P$$

or

$$z = \frac{\frac{L^*}{n} - \mu_p}{\sigma_P}$$

Values can now be calculated using the automobile portfolio example considered earlier. Recall that each exposure unit had an expected value of $500 and a standard deviation of $800. The expected loss per policy is:

$$\iota_p \quad = \quad E\left(\frac{L}{n}\right) \quad = \quad \frac{E(L_1) \ + \ E(L_2) \ + \ .....+ \ EL(_n)}{n} \quad = \quad \frac{n(500)}{n} \quad = \quad 50$$

The portfolio standard deviation has already been calculated:

$$\sigma_p \quad = \quad \sigma\left(\frac{L}{n}\right) \quad = \quad \frac{\sigma_i}{\sqrt{n}} \quad = \quad \frac{800}{\sqrt{n}}$$

Substituting these values into the equation for z gives

$$z \quad = \quad \frac{(L^*/n) \ - \ 500}{800/\sqrt{n}}$$

Now suppose that the insurer's reserves and surplus permit it to pay up to a maximum of $660 per policy. We have

$$z \quad = \quad \frac{660 \ - \ 500}{800/\sqrt{n}} \quad = \quad 0.2\sqrt{n}$$

Thus if n = 9, z = 0.6. Reading from the table of z values, the probability that L/ n > L*/n is 0.2743. This value is recorded in Table 2 alongside values calculated for other n.

TABLE 2
RUIN PROBABILITIES IN PORTFOLIOS OF INDEPENDENT RISKS

n	z	probability that L/n > L*/n
1[a]	0.2	0.4207
9[a]	0.6	0.2743
50	1.41	0.0793
100	2.0	0.0228
1000	6.3	negligible
10,000	20.0	negligible
∞	∞	0

a: These sample values are too small for reasonable approximation with the normal distribution.

The results in Table 1 and Table 2 confirm the pattern in Figure 3 that diversification reduces risk, both in the sense of reducing variance per policy and by reducing the probability of ruin. We will now try to apply these ideas to more complex portfolios.

RISK REDUCTION IN AN INSURANCE PORTFOLIO OF INDEPENDENT, HETEROGENEOUS RISKS

The second type of portfolio preserves the assumption that risks are independent of each other, but relaxes the restriction that all units are homogeneous in the sense of having identical probability distributions. The heterogeneous portfolios we have in mind here could comprise policies of different types, such as a mixture of fire, automobile, and liability policies, or they could comprise policies of the same type but having different risk levels, such as a fire insurance portfolio drawn from different industries. Or the portfolio could combine a firm's exposure to insurable risk, to financial risk and to market risk. ln order to measure the effects of size on this type of portfolio, it is important to distinguish between two possibilities, one relating solely to the size of the portfolio and the other combining both size and composition effects:

1. Growth of a portfolio might arise without change in the mix of the various types of exposure units. The firm might simply write different types of new policies in the proportions that those policy types represent in the existing portfolio.

2. Conversely, the insurer may grow through acquisition of new policies which are different from existing business. For example, an insurer with a heavy automobile portfolio might acquire another company with a predominantly fire portfolio. Such growth changes the mix of business.

Our main concern is with the effect of size on portfolio risk; thus we will consider size changes that preserve composition. The combined size and mix effects will be examined presently.

Insurers typically subdivide their insurance portfolios into risk classes that are fairly homogeneous internally. Fire, automobile, marine, and liability policies are primary subdivisions. In turn, these lines are sub-classified by insured characteristics relating to hazard level (e.g., by age and sex in life insurance, by industry, protection and location in fire insurance, or by vehicle type, driver age and sex, location, etc. in automobile insurance). The number of homogeneous categories can be very large, but the purpose is to group units that are *a priori* similar in order to aid the rating and underwriting decision process. For our purposes, let us suppose that there are three categories, labeled R, S, and T. The total number of policies is n, where n = r + s + t, and r, s, and t are the numbers of R, S, and T policies, respectively. Within each group, units are identically distributed, having the same mean and variance, and they are subscripted by the group code, for example, $E(L_R)$, $E(L_T)$, σ_S , etc.. Let us combine the policies to examine the risk-return characteristics of the portfolio (L/n).

Recalling the calculations of expected value and variance of the distribution of mean loss, we have

$$\mu_p \quad = \quad E\left(\frac{L}{n}\right) \quad = \quad \frac{1}{n}\sum_{t=1}^{n} \; E(L_i)$$

and, since the covariances are zero,

$$\sigma_p^2 = \sigma^2\left(\frac{L}{n}\right) = \sum_{i=1}^{n} \sigma_i^2 + \sum \sum_{i \neq j} \sigma_{i,j} = \frac{1}{n^2}\left(r\sigma_R^2 + s\sigma_s^2 + t\sigma_t^2\right)$$

To put flesh on this skeleton, consider the following portfolio:

R refers to fire policies each with $E(L_R) = 300$ and $\sigma_R = 400$
S refers to automobile policies each with $E(L_S) = 400$ and $\sigma_S = 350$
T refers to liability policies each with $E(L_T) = 200$ and $\sigma_T = 300$

The portfolio comprises the following numbers of policies in each category:

r = 0.5n ; s = 0.3n ; t = 0.2n

This permits us to examine the effect of growth in "n" without change in the composition of the component units. Now,

$$\mu_p = E\left(\frac{L}{n}\right) = \frac{0.5n(300) + 0.3n(400) + 0.2n(200)}{n} = 310$$

$$\sigma_p^2 = \frac{1}{n^2}\left(0.5n(400)^2 + 0.3n(350)^2 + 0.2n(300)^2\right)$$

$$= \frac{0.5(400)^2 + 0.3(350)^2 + 0.2(300)^2}{n} = \frac{134{,}750}{n}$$

giving

$$\sigma_p = \sqrt{\frac{134{,}750}{n}} + \frac{367}{\sqrt{n}}$$

The forms of the variance and standard deviation clearly reveal that riskiness of the portfolio will fall in the pattern depicted in Figure 2 for the portfolio of homogeneous risks. Furthermore, the standard deviation will tend toward zero as the number of exposure units tends toward infinity. Some values for n and the corresponding variances and standard deviations are shown in Table 3.

TABLE 3
RISK REDUCTION: INDEPENDENT HETEROGENEOUS RISKS

n	σ^2 (L/n)	σ (L/n)
1	134,750	367
10	13,475	116
100	1,348	36.7
1,000	135	11.6
10,000	14	3.7
∞	0	0

The reader might experiment with other portfolios that show different initial compositions. However, caution must be exercised to ensure that the proportionate mixture of units does not change. This will reveal that the law of large numbers continues to apply for all portfolios of independent risk units, regardless of proportionate mix, a long as diversification does not affect the relative composition.

The second measure of risk is the probability of ruin. We can apply this in an identical manner to that shown for homogeneous portfolios. The reason is that the normal approximation continues in view of the independence of the risk units. To illustrate the process, let us assume that the insurer's reserves and surplus permit it to meet claim liabilities up to a limit of $400 per policy. We now have sufficient information to calculate ruin probabilities for different portfolio sizes. Using the formula for z, and the previously calculated values for μ_p and σ_p , we have

$$z = \frac{L^*/n - \mu_p}{\sigma_p} = \frac{400 - 310}{\frac{367}{\sqrt{n}}} = 0.245\sqrt{n}$$

Values for ruin probabilities at different levels of n are given in Table 4.

TABLE 4
RUIN PROBABILITIES: INDEPENDENT HETEROGENEOUS RISKS

n	**z**	**Probability that L/n > L*/n**
1	0.25	0.4013
9	0.74	0.2296
10	0.77	0.2206
100	2.45	0.0071
1,000	7.75	negligible
10,000	24.5	negligible
∞	∞	zero

The ruin probability behaves in a fashion similar to that for the portfolio of independent and homogeneous risks; it diminishes as "n" increases, approaching zero when n approaches infinity. Combining this with the result for the behavior of standard deviation, diversification achieved by increasing the size of the insurance portfolio reduces portfolio risk. This result is valid if exposure units are independent, regardless of whether the units are identical or not, as long as the portfolio composition remains unchanged. To see how important this last assumption is, we will consider an example in which the portfolio size grows but the composition also changes. Independence is retained in this example.

EXAMPLE: CHANGING PORTFOLIO COMPOSITION

Consider an initial portfolio of, say, 10,000 automobile policies, each with a standard deviation of \$300 and an expected loss of \$300. To calculate the ruin probability, assume the insurer can pay claims up to \$400 per policy. The mean, standard deviation, and ruin probability of the insurer's portfolio are

and

$$z = \frac{L^*/n - \mu_p}{\sigma_p} = \frac{400 - 300}{3} = 33.33$$

Reading from Table 3, the probability of ruin is negligible (less than 0.001).

Now suppose that this insurer also writes a set of fire insurance policies in order to diversify its line base. The firm writes 1000 fire policies each covering fairly large industrial risks. Each risk has an expected loss of \$10,000 and a standard deviation of \$15,000. Now the firm estimates that its reserves and surplus permit it to meet claims up to \$1300 per policy. Now the portfolio mean, standard deviation, and ruin probability are

$$\mu_p = \frac{10{,}000(300) + 1{,}000(10{,}000)}{11{,}000} = \$1{,}181.8$$

$$\sigma_p = \left(\frac{1}{11{,}000^2}\left[10{,}000(300)^2 + 1{,}000(15{,}000)^2\right]\right)^{1/2} = \$43.21$$

and

$$z = \frac{1300 - 1181.8}{43.21} = 2.74$$

Hence, the probability of ruin is 0.0031.

In this case, diversification has increased the portfolio risk both in terms of the standard deviation averaged on a per policy basis and in terms of the ruin probability. A little reflection will reveal what has gone wrong. Originally, the insurer had written a book of automobile risks that each had a small expected value and a small level of risk. In diversifying, the insurer changed the composition of the portfolio to include a significant number of high-value, high-risk policies. The impact is to "swamp out" the already diversified risks of the automobile portfolio. The mean and variance of the sub-portfolio of fire policies alone is

$$\mu_p\,(fire) = \$10{,}000 \quad \text{and} \quad \sigma_p\,(fire) = \$474.3$$

Thus, diversification means that an automobile portfolio having a standard deviation of \$3 per policy is combined with one having a standard deviation of \$474.3 per policy. Not surprisingly, the fire policy risks dominate the portfolio, increasing overall portfolio risk.

MEASURING CORRELATION BETWEEN RISK UNITS

With independence, the covariance terms between risk units are defined to be zero. Now we consider portfolios in which there is correlation between exposure units. Consider some examples:

1. Close proximity of buildings implies that a fire in one structure immediately increases the probability of a fire in each of the remaining buildings, since the natural progress of the fire or the blowing of burning embers may cause the fire to spread. The probabilities of loss in each are interdependent.

2. Bad weather conditions increase the probability that any one vehicle may crash, but the same conditions increase the accident probability for every other vehicle in the same meteorological region. Thus the probabilities are interdependent.

3. Many types of losses, e.g., fires, industrial accidents, collisions, burglaries, etc., may bear some relation to the economic cycle. For example, arson increases during recession; people drive more miles during peak periods in the cycle, increasing exposure to loss; etc. These sorts of effects may lead to relationships between the loss probabilities in different types of portfolios. For example, fire losses might be correlated with automobile losses.

4. A non insurance firm has the risk of liability of directors and officers from

shareholders lawsuits. Such lawsuits are more likely when the firm's stock price is depressed. And the stock price is related in part to financial and business conditions. Thus, the firm considers its directors and officers liability risk to be correlated with its business and financial risk.

Whether such relationships do exist or not is an empirical matter.

To illustrate the operation of covariance, consider the loss experience on two fire policies for industrial risks. The historical record for a 16-year period is given in Table 5. The data are expressed in real terms to remove the effect of inflation.

TABLE 5 SAMPLE FIRE LOSS DATA

Total losses for exposure unit, $

Year	1	2
1985	5,000	500
1986	300	600
1987	6,200	3,200
1988	7,000	2,900
1989	-	1,500
1990	500	-
1991	5,300	
1992	25,000	5,600
1993	5,000	2,000
1994	7,000	1,800
1995	200	-
1996	-	-
1997	1,000	500
1998	8,000	-
1999	32,000	3,000
2000	2,300	3,500
μ	6,550	1,569
σ	8,831	1,625

The formula for covariance was given earlier, but the following form will make the computation somewhat simpler:

$$\sigma_{1,2} = \sum p_i (L_{1i} - E(L_1))(L_{2i} - E(L_2)) = E(L_1 L_2) - E(L_1)E(L_2)$$

$$= 19{,}373{,}125 - 6550(1569) = 9{,}096{,}175$$

The relationship is positive. Large losses on exposure unit 1 tend to arise at the same time as large losses on exposure unit 2; small losses similarly coincide. Intuitively, we might expect that if all units in a portfolio exhibited such positive relationships, the portfolio risk would tend to be high. Bad weather, or depressed economic conditions, would bring a flood of claims. "It never rains but it pours." However, good weather or a strong economy might reverse the flow of claims. This correlation could cause large fluctuations in the insurer's total claims experience, even if the insurer holds a large number of policies. The impact of covariance on portfolio risk depends on the relative strength of the covariance. Unfortunately, the measure $\sigma_{i,j}$ does not directly reveal the strength of the relationship; it is an absolute number that is not placed on any scale. Certainly, the number 9,096,175 seems a large number, but a casual inspection of Table 5 shows the relationship to be less than perfect. If covariance were very strong, we always find that above-average losses on unit 1 are associated with above-average losses on unit 2. This is clearly not so; witness 1987, 1993, 1998, and 2000. The good news is that a convenient measure of the strength of the relationship is available. This is the *correlation coefficient*, which is defined as

$$r_{i,j} = \frac{\sigma_{i,j}}{\sigma_i \, \sigma_j}$$

Thus, for the information in Table 5, the correlation coefficient is

$$r_{i,j} = \frac{9{,}096{,}175}{8831(1625)} = 0.634$$

The correlation coefficient is defined on a scale of minus unity to plus unity. Negative values show a negative relationship, with - 1 suggesting a "perfect negative correlation". A perfect negative relationship of this form means that there is an inversely proportional relationship between the variables; if one variable falls to half its previous value, the other variable doubles in value. A correlation coefficient of zero indicates that there is no association between the variables. Positive values denote a positive relationship, with + 1 indicating a "perfect positive correlation." A perfect positive relationship is a directly proportional relationship between the two series. The example in Table 5 reveals a reasonably strong positive relationship. Insurers would worry about this strong relationship because it implies that there is a fire conflagration hazard. In a diversified insurance portfolio, it would be expected that correlations would average much lower values.

Using the correlation coefficient instead of covariance, the variance of a portfolio is

$$\sigma^2(L) = \sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i \neq j} \sigma_{i,j} = \sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i \neq j} r_{i,j}\, \sigma_i\, \sigma_j$$

Thus, the standard deviation of the mean loss is

$$\sigma^2\left(\frac{L}{n}\right) = \left[\frac{1}{n^2}\left(\sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i \neq j} r_{i,j}\, \sigma_i\, \sigma_j\right)\right]^{1/2}$$

FIGURE 5

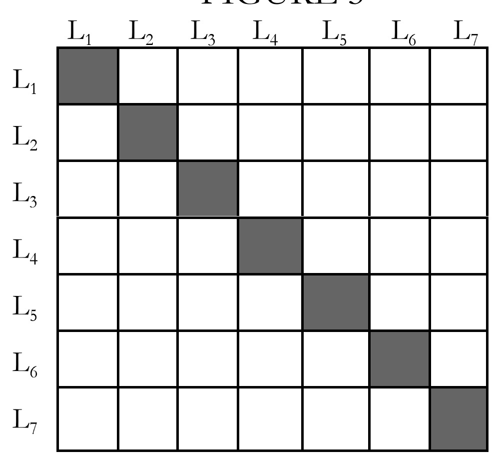

This is the expression we must now use to calculate the variance of the insurance portfolio. This task will be more complex in view of all the correlation terms that had previously dropped out. To perform such calculations, it is helpful to know how many terms there will be in this expression and, in particular, how many co-variances or correlations.

The variances and covariances in a portfolio can be represented schematically, as shown in Figure 5. All the shaded terms in the leading diagonal are the variances, for example, the pair L_4, L_4. The remaining expressions represent all the possible covariances, for example, L_3 and L_7. In total, the matrix must have n times n elements, since it is square (in this picture n is 7). Of these elements, there are n variances, leaving (n^2 -n) or n(n - 1) covariances. The double summation sign in the previous variance formula contains n(n - 1) terms. This will be very convenient to remember. The behavior of a portfolio having correlated exposure units can now be examined.

RISK REDUCTION IN AN INSURANCE PORTFOLIO OF INTERDEPENDENT RISKS

For a portfolio of "n" identical exposure units in which the correlation coefficient between any pair i and j is $r_{i,j}$, the mean and standard deviation are

$$\mu_p = E\left(\frac{L}{n}\right) = \sum_{i=1}^{n} \frac{E(L_i)}{n} = E(L_i)$$

$$\sigma_p = \sigma^2\left(\frac{L}{n}\right) = \left[\frac{1}{n^2}\left(\sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i\neq j} r_{i,j}\,\sigma_i\,\sigma_j\right)\right]^{1/2} = \left[\frac{1}{n^2}\left(n\sigma_i^2 + n(n-1)r_{i,j}\sigma_i\sigma_j\right)\right]^{1/2}$$

$$= \left[\frac{\sigma_i^2}{n^2} + \left(\frac{n-1}{n}\right) r_{i,j}\sigma_i\sigma_j\right]^{1/2}$$

If $r_{i,j} > 0$, then portfolio risk will not converge on zero when n approaches infinity because the term (n-1)/n approaches unity as n approaches infinity. Thus, the second term in the bracket will approach the positive constant $r_{i,j}\,\sigma_i\,\sigma_j$ as n approaches infinity. This can be seen in the following example.

Consider a portfolio with n policies that each have

$\mu_i = 500$; $\sigma_i = 650$; $r_{i,j} = 0.1$

$$E\left(\frac{L}{n}\right) = 500$$

$$\sigma_p = \left[\frac{650^2}{n} + \frac{n-1}{n}\,0.1(650)(650)\right]^{1/2} = \left[\frac{380{,}250}{n} + 42{,}250\right]^{1/2}$$

Values of for the mean, variance and standard deviation are shown in Table 6 for different portfolio sizes.

TABLE 6
RISK REDUCTION: INTERDEPENDENT RISKS

n	$\sigma_2(L/n)$	$\sigma(L/n)$
1	422,500	650
10	80,275	283
100	46,052	214.6
1,000	42,630	206.5
10,000	42,288	205.6
∞	42,250	205.5

Consequently, the portfolio risk is seen to converge on some positive value, that is, 205.5, as n approaches infinity. Diversification brings some risk reduction, but there is a limit to this process. This limit is determined by the size of the correlation coefficient. The higher the correlation coefficient, the higher the lower limit on portfolio risk. In Figure 6 we show how the correlation coefficient affects diversification. We assume the same portfolio of units that each have $\sigma_i = 650$, but we allow the correlation coefficient to assume the values 0.1, 0.2, 0.5, and 1.0. The calculation follows the previous case. For each correlation coefficient, the portfolio standard deviation asymptotically approaches the value shown. There is a clear scope for risk reduction as long as units are not perfectly correlated. However, the risk-reducing benefits become more pronounced as the correlations get lower.

The second aspect of risk that has been explored is the ruin probability. When risks are independent, the combined distribution approaches normal as the number of risks approach infinity. The distribution cannot be assumed to be exactly normal, but the approximation is sufficiently close for reasonable estimation when n exceeds, say, 30. *The assumption that the portfolio distribution approaches normal does not hold when the exposure units are correlated*. This means that use of the normal distribution to estimate ruin probability is not appropriate for correlated exposures even when n is very large. The degree of error in using the normal approximation for such cases depends on the degree of correlation. If the correlation coefficient is close to zero, the normal approximation may be reasonable if n is fairly large. However, when the correlation coefficient diverges widely from zero, the error may be very substantial. Our purpose here is not to give accurate estimates of ruin probabilities; rather, it is to show how diversification reduces the risk of ruin. This will be illustrated on a portfolio with fairly low correlations, and we will use the normal distribution as a crude approximation. However, we recognize that the results may contain a significant error factor, and we will therefore supplement with alternative calculations based on Chebyshev's inequality. The Chebyshev method enables us to establish an upper bound for the ruin probability without actually providing a direct calculation.

The normal approximation is calculated as before, recognizing the additional error:

$$z = \frac{L^*/n - \mu_p}{\sigma_p} \quad : \quad \textit{probability that} \ \frac{L}{n} > \frac{L^*}{n}$$

Chebyshev's inequality states that

> *For any random variable x, the probability of realizing an outcome above k standard deviations above the mean is at most $1/k^2$.*

This implies that an upper limit can be set on the ruin probability for any critical value L*/n. The actual ruin probability will be lower than this value. However, unlike the calculation from the normal approximation, the upper-limit calculation is valid for any distribution. Given this generality, the upper limit calculated by the Chebyshev method tends to be conservatively high. Nevertheless, we will observe how the estimated ruin probabilities behave as the number of items in the portfolio increase. To simplify the calculation, we will specify the Chebyshev inequality in the following terms:

$$\mathit{Probability}\left(\left|\frac{L^*}{n}\right| > \mu_p + k\sigma_p\right) \leq \frac{1}{k^2}$$

Verbally, the probability that the absolute value of L*/n is greater than $\mu_p + k\sigma_p$ is at most $1/k^2$. Therefore, an upper limit on the ruin probability, given some maximum payout L*/n, is given by α in the following expression:

$$\frac{L^*}{n} = \mu_p + \sigma_p\sqrt{\frac{1}{\alpha}}$$

Rearranging gives

$$\alpha = \frac{1}{\left(\frac{L^*/n - \mu_p}{\sigma_p}\right)^2} = \frac{1}{z^2}$$

where z is defined as earlier for the normal distribution. The Chebyshev method will only provide meaningful answers (probabilities below unity) if k>1.

Now consider the portfolio discussed earlier with $\mu_i = 500$, $\sigma_i = 650$ and $r_{i,j} = 0.1$, to which we will add the information that the insurer can pay losses up to maximum of $800 per policy. Calculating rough ruin probabilities with the normal approximation gives:

$$z = \frac{800 - 500}{\sqrt{380{,}250/n + 42{,}250}}$$

Table 7 shows sample calculations using this expression and the z tables.

TABLE 7
RUIN PROBABILITIES: INTERDEPENDENT RISKS

n	z	Normal approximation	Chebyshev upper limit ($1/z^2$)
1	0.462	0.3228	a
10	1.059	0.1446	0.892
100	1.398	0.0808	0.512
1,000	1.453	0.0735	0.474
10,000	1.458	0.0721	0.470
∞	1.460	0.0721	0.469

"a" value exceeds unity.

Both the normal-approximation values and the Chebyshev upper limits reveal that ruin probabilities decline as the number of exposure units increase. Furthermore, both the normal approximation values and the Chebyshev upper limit converge on positive values, revealing no tendency for the ruin probability to approach zero as n approaches infinity. This may be contrasted with the portfolio of independent units that does converge to zero risk.[3] In view of the latent error in the normal approximation and the severe conservatism of the Chebyshev method, these results are not as strong as we would like.[4] However, they may be interpreted as providing some support for the claim that diversification will still provide limited benefits in the form of risk reduction even when exposure units are positively correlated. The reduction in risk arises in all cases in which the correlation coefficient is less than

---

[3]For independent exposures, it is easily verified that the Chebyshev upper limit on ruin probability also approaches zero. Tables 2 and 4 showed that, for such portfolios, $z \to \infty$ as $n \to \infty$. Since the Chebyshev value α equals $1/z^2$, then clearly $\alpha \to 0$ as $n \to \infty$.

[4]There are many other methods of estimating the tails of distributions. One of the newer methods is extreme value theory (see Embrechts).

unity.

RISK REDUCTION WITH CORRELATED RISKS

For completeness, we should show that limited risk reduction can be achieved in heterogeneous portfolios of positively correlated risk units as long as the correlation coefficients are less than unity. The method of tackling this should be well established by now, and we will content ourselves with a quick example showing the calculation of portfolio standard deviation.

Let us consider a mixed portfolio of fire and automobile policies in which there are "s" fire policies and "t" automobile policies. To preserve the portfolio mix, we assume that s and t bear a constant ratio to each other. The total number of policies is s + t = n. In order to calculate the portfolio standard deviation, we need to know how many terms will be present in the formula. Clearly, there are s + t variance terms for individual policies. In addition, there will be:

$s(s - 1)$ correlations between different fire risks
$t(t - 1)$ correlations between different automobile risks
$2st$ correlations between fire and automobile risks

The respective correlation coefficients for these groups are

$r_{F, i, j} = 0.1$ for fire policies
$r_{A, i, j} = 0.1$ for auto policies
$r_{A, F} = 0.05$ for fire and auto policies

Further information for the portfolio is

$s = 20{,}000$	$\mu_F = 500$	$\sigma_F = 700$	for all fire risks
$t = 10{,}000$	$\mu_A = 300$	$\sigma_A = 400$	for all automobile risks

The standard deviation of the distribution of mean loss can now be calculated in the usual manner:

$$\sigma_p = \sigma^2\left(\frac{L}{n}\right) = \left[\frac{1}{n^2}\left(\sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i \neq j} r_{i,j}\, \sigma_i\, \sigma_j\right)\right]^{1/2}$$

= { $1/30{,}000^2$ [$20{,}000(700)^2 + 10{,}000(400)^2$ + (20,000)(19,999)(0. 1)(700)(700)
+ (10,000)(9,999)(0. 1)(400)(400) + 2(20,000)(10,000)(0.05)(700)(400)]}$^{1/2}$

= $29{,}789^{1/2}$

= 173

The interested reader might experiment with different values of s and t (but retaining the ratio of s=2t) to confirm that risk does indeed fall as n increases and that it converges on a positive value as n approaches infinity.

A final word in this section concerns negative correlations. Within a real insurance portfolio there may be a large degree of heterogeneity between risks and a wide diversity in the correlations between individual pairs of risks. Some of these correlations might be negative. The incremental effect of including a policy that is negatively correlated with other policies already in the portfolio is to reduce risk substantially. Therefore, negative correlations should be valued by the insurer.

THREE APPLICATIONS

To highlight the process of diversification, three examples are now given below in which common risk management types of activities take advantage of diversification.

EXAMPLE 1: MERGERS AND ACQUISITIONS. *An insurer with some 80,000 industrial fire policies acquires a second firm writing 20,000 automobile policies. The respective policies will be subscripted F and A, and details of individual policies are given as follows:*

Fire policies: *Each policy has mean* $\mu_F = 1000$ *and standard deviation* $\sigma_F = 1500$. *The correlation between any pair is* $r_{F,i,j} = 0.1$
Automobile policies: *Each policy has mean* $\mu_A = 500$ *and standard deviation* $\sigma_A = 700$. *The*

correlation between any pair is $r_{A,\,i,\,j} = 0.1$
Fire and auto policies are uncorrelated.

The standard deviation of the distribution of mean loss of each insurer before acquisition is

$$\sigma_{p,F} = \left[\frac{1}{n^2}\left(\sum_{i=1}^{n}\sigma_i^2 + \sum\sum_{i\neq j} r_{i,j}\,\sigma_i\,\sigma_j\right)\right]^{1/2}$$

$$= \{1/80{,}000^2\,[\,80{,}000(1500)^2 + (80{,}000)(79{,}999)(0.1)(1500)(1500)\,]\,\}^{1/2}$$

$$= 474$$

and

$$\sigma_{p,\,A} = \{1/20{,}000^2\,[\,20{,}000(700)^2 + (20{,}000)(19{,}999)(0.1)(700)(700)\,]\,\}^{1/2}$$

$$= 221$$

And, after the merger, the acquiring insurer will have a portfolio with standard deviation

$$\sigma_{p(F+A)} = \{\ 1/100{,}000^2\,[\,80{,}000(1500)^2 + 20{,}000(700)^2 + (80{,}000)(79{,}999)(0.1)(1500)(1500) + (20{,}000)(19{,}999)(0.1)(700)(700)\,]\}^{1/2}$$

$$= 382$$

This reduction in risk has been accomplished by two effects. First, the acquiring company has obtained a group of policies with significantly lower risk (standard deviation) per policy. Second, the initial fire portfolio and the acquired automobile portfolio are uncorrelated; thus diversification is achieved simply by combining two uncorrelated portfolios.

It might be objected that the reduction in risk calculated on a per policy base is misleading. The original portfolio consisted only of fire policies, but the combined portfolio averages over the "chalk and cheese" of fire and automobile policies. To avoid this problem, consider what happens to total risk (that is, not averaged over policies). This is calculated in a similar fashion to σ_p *, except that we do not divide by* $1/n^2$*. The before and after values are*

$$\sigma_p \quad (total) = 37{,}949{,}466$$

$\sigma_{p,A}$ (total) = 4,428,184
$\sigma_{p,(F+A)}$ (total) = 38,206,947

Not surprisingly, total risk increases. However, the combined value is significantly less than the sum of the parts. In fact, total risk hardly increases at all, despite the fact that insurer has increased its portfolio by some 20,000 policies.

EXAMPLE 2: REINSURANCE. *Two insurers have identical portfolios that each comprise 10,000 identical policies with* $\sigma = 400$ *and* $r_{i,j} = 0.1$*, for any i and j. The correlation between any policy of insurer A and any policy of insurer B is also 0.1. The insurers establish a reciprocal reinsurance treaty in which A receives half the premium on each policy written by B and vice versa. In return, A pays B half the cost of each and every claim suffered by B, and vice versa. That is, each pays half its premiums to the other in return for reimbursement on half of each claim. Before the treaty, each insurer had a portfolio risk of*

$$\sigma_p = \left[\frac{1}{n^2} \left(\sum_{i=1}^{n} \sigma_i^2 + \sum \sum_{i \neq j} r_{i,j} \, \sigma_i \, \sigma_j \right) \right]^{1/2}$$

$$= \{1/10{,}000^2 \, [\, 10{,}000(400)^2 + (10{,}000)(9{,}999)(0.1)(400)(400) \,] \}^{1/2}$$

$$= 126.5$$

To calculate standard deviation after the treaty, consider that the agreement has produced two relevant effects. Whereas each insurer previously had an interest in only 10,000 risk units, it now has an interest in 20,000. Second, each insurer will pay only half the loss on each policy; thus the standard deviation per risk unit will be smaller. Both effects should reduce the portfolio standard deviation. The standard deviation per exposure unit can be calculated as follows:

Before reinsurance:

$$\sigma_i = \left[\sum_i p_i \left(L_i - E(L_i) \right)^2 \right]^{1/2} = 400$$

After reinsurance:

$$\sigma_i = \left[\sum_i p_i \left(\frac{L_i}{2} - \frac{E(L_i)}{2}\right)^2\right]^{1/2} = \tfrac{1}{2}\left[\sum_i p_i \left(L_i - E(L_i)\right)^2\right]^{1/2} = 200$$

Then the portfolio standard deviation for each insurer after the treaty is:

$$= \{1/20{,}000^2\ [\,20{,}000(200)^2 + (20{,}000)(19{,}999)(0.1)(200)(200)\,]\ \}^{1/2}$$

$$= 63.26$$

The reduction in risk is dramatic. Generally, the smaller the correlation coefficients between the portfolios of the participating insurers, the greater the risk reduction achieved by reciprocal reinsurance arrangements of this nature.

***EXAMPLE 3: RUIN IN A SELF-INSURANCE FUND**. A grocery retailer has 100 stores that are located in several states. In view of the geographic separation and similar size and risk of the stores, the firm decides to self-insure. The administration of this program involves the establishment of a fund into which "premiums" are paid and from which losses are financed. Each store has an expected loss of $1000 and a standard deviation of $1500. The risks are assumed to be independent. Ignoring any administrative expenses, how much should the firm contribute to the fund for each store in order to restrict the probability to 10% that the fund's resources will be inadequate to meet losses.*

Since there are 100 units and they are independent, we can assume that the normal approximation applies. Recall that z tables record the probability that any revealed value from a normal distribution will be more than z standard deviations from the mean. The z value corresponding to a 10% probability is 1.28. Thus,

$$z_{10} = \frac{F - \mu_p}{\sigma_p} = 1.28$$

where F is the cutoff contribution that leaves a 10% chance of ruin. Therefore,

$$1.28 = \frac{F - 1000}{1500/\sqrt{100}}$$

giving $F = 1192$.

Since the assumptions for the use of the normal approximation are met in this particular case, the result is fairly accurate (not strictly accurate, since the portfolio distribution only approaches normal as n approaches infinity). Had there been significant correlations, this method would not have produced reliable answers. However, we would generally expect that the required contribution to maintain any given ruin probability would be higher with positive correlation because the portfolio distribution would exhibit higher risk.

DIVERSIFICATION AND THE DISTRIBUTION OF AGGREGATE LOSSES FOR THE (NON-INSURANCE) FIRM

Example 3 provides a bridge between the use of portfolio theory to explain risk reduction in an insurance fund and its application to risk management in other types of corporate organizations. All firms face a variety of risks, and each type of risk can affect the firm's financial performance and indeed its valuation. For some purposes, it may be useful to identify and treat these risk units individually. A narrow focus on individual units provides a starting point in estimating an aggregate loss distribution, or at least summary measures. Furthermore, attention to individual units may reveal possibilities for reducing loss costs by safety and preventive activities. However, in determining the financial performance of a firm, it is the aggregate impact of these risk units on the earnings and value of the firm that is of importance. Since the properties of a portfolio are somewhat different from the sum of its parts, we must derive risk management strategies with reference to these aggregate effects. In the meantime, we consider the properties of the portfolio of exposure units that is owned or controlled by a firm.

All firms are exposed to multiple risks. In addition to the risks inherent in business activity (such as market risk, financial risk and regulatory risk) there are multiple pure risks, (such as fires, liability suits, weather-related perils). Even a very small single-plant firm can consider itself to have a portfolio of risk exposures, and each exposure pertains to a different type of risky event. Just how diversified such a portfolio will be depends on how the risks units are defined and the correlations that exist between them. We have already argued that the identification of separate exposure units is a matter of judgment relating to the degree of correlation between risks. However, it should be clear from the preceding discussion that we cannot usefully define exposure units as statistically independent. Thus, for a one-plant firm, the risk of fire to the plant may exhibit a "small" correlation with the probability that the firm's truck will collide, leading to the common-sense judgment that they are separate exposure units.

For larger, more diversified firms, the portfolio of risk units is more obvious. Firms such as Safeway, Ford, IBM, etc. own many plants, stores, warehouses, offices, etc., and each can be considered to be a separate exposure unit, even when contemplating a single peril such as fire. Again, we are not asserting independence. It may be that, due to the quality of management, corporate protection policies, the effect of economic cycles on corporate activity, etc., there is interdependence between the loss distributions for separate exposure units. Nevertheless, these firms do achieve some degree of diversification, since the correlations between units will fall far short of unity. The same general principles of diversification apply to the risk management issues facing a firm and to the formation of an insurance portfolio. Therefore, it is important to examine the properties of a firm's portfolio of exposure units in order to properly measure risk management costs and derive appropriate risk management strategies.

Although the general principles of portfolio theory can be transferred to risk management problems for industrial, commercial, and service firms, the resulting portfolio may not display such convenient properties as those for the insurance portfolio. The main differences that do arise are

1. Many firms comprise only a small number of risk units. Consequently, portfolio risk will not substantially disappear even if risk units are not correlated.

2. However, risky events may be correlated. Risk units within a firm may be exposed to common influences that imply statistical dependence. For example, financial and insurable risks may be commonly related to the level of business activity for the firm. Common management strategies, labor relations, and safety programs may give rise to interdependence. Physical proximity of some units may give rise to conflagration hazard (although at some stage conflagration hazard may be so severe that risks may be combined into a single exposure unit).

3. A third difference is that risk exposures may differ vastly in terms of value and expected loss. Insurance portfolios are usually far from homogeneous, although insurers do try to group exposures into fairly homogeneous classes. The device of "averaging" risk over exposure units was useful in illustrating the nature of risk spreading. For risk management purposes, the main focus of attention is on the distribution of aggregate losses. It is this distribution that determines the risk costs to the firm, and it is this distribution that becomes a primary focus in formulating risk management strategy. We will now show how the distribution can be affected by diversification within the firm's portfolio of exposure units.

These conditions (if present) imply that the degree to which risk is reduced by diversification is limited. Nonetheless, risk can always be reduced to some extent if exposure units are not perfectly correlated. To illustrate the range of possibilities, consider three firms that differ with respect to the degree of diversification. Details of assets held by the firms and summary measures of the fire loss distributions are as follows:

Firm A: Operates 30 fast-food restaurants. The restaurants are similar style and have equivalent value and are exposed to identical risk from fire:

Value of each restaurant	$ 300,000
Total value	$9,000,000
Expected loss, each restaurant	$ 1,500
Total expected loss	$ 45,000
Standard deviation, each restaurant	$ 2,500
Exposure units are independent	

Firm B: Manufactures plastic goods. Its business premises are constructed on a fairly large site with sufficient physical separation of buildings to limit fire conflagration hazard. The building values and summary loss measures are:

	Value	Expected loss	Standard deviation
Factory	$5 million	$30,000	$50,000
Warehouse	$2 million	$10,000	$16,667
Distribution	$1 million	$ 3,000	$ 5,000
Office	$1 million	$ 2,000	$ 3,333
Total	$9 million	$45,000	

The correlation coefficient between exposure units is $r_{i,j} = 0.1$.

Firm C: Also manufactures plastic goods. Manufacture, storage, and related office works are all housed in a single building, which is considered a single exposure unit from a fire viewpoint.

Value of premises	\$9,000,000
Expected loss	\$ 45,000
Standard deviation	\$ 75,000

The three firms described are comparable in terms of the value at risk and the risk characteristics. Each firm has assets at risk valued at \$9 million. The total expected loss in each case is \$45,000, indicating comparable degrees of hazard. Furthermore, the variability of risk for each exposure unit, measured as the ratio of standard deviation to expected loss, assumes the same value (≈1.67). The three firms differ only with respect to the degree of diversification on the riskiness of the firm's *aggregate* loss distribution. Using the standard deviation formula,

$$\sigma(L) = \left[\sum_{i=1}^{n} \sigma_i^2 + \sum\sum_{i \neq j} r_{i,j}\, \sigma_i\, \sigma_j\right]^{1/2}$$

For firm A:

$$\sigma_A(L) = [30(2500)^2]^{1/2}$$

$$= \mathbf{\$13{,}693}$$

For firm B:

$$\begin{aligned}\sigma_B(L) = \; & [(50{,}000)^2 + (16{,}667)^2 + (5{,}000)^2 + (3{,}333)^2 \\ & + 2(0.1)(50{,}000)(16{,}667) + 2(0.1)(50{,}000)(5{,}000) \\ & + 2(0.1)(50{,}000)(3{,}333) + 2(0.1)(16{,}667)(5{,}000) \\ & + 2(0.1)(16{,}667)(3{,}333) + 2(0.1)(5{,}000)(3{,}333)]^{1/2}\end{aligned}$$

$$= \mathbf{\$55{,}633}$$

For firm C, the standard deviation of aggregate losses is already given, since there is only one exposure unit:

$$\sigma_C(L) = \mathbf{\$75,000}$$

Although the expected loss is the same, the variability in losses differs considerably. This difference may substantially affect the type of risk management strategy adopted by each firm. Consider, for example, whether the firms should purchase insurance. The need for insurance is normally considered to arise from the variability of the loss distribution. Insurance certainly does not reduce expected loss cost (in view of transaction costs and insurer's markups, the expected cost of an insurance policy usually exceeds the expected value of the loss). However, insurance does reduce, or ideally eliminate, variability. One might therefore suppose that firm A has the least to gain from insurance and firm C has the most to gain. We should not be surprised to find firms with a wide spread of exposure units choosing not to purchase insurance and firms with a high concentration of value in a small number of exposure units choosing extensive insurance protection.

CONCLUSION

This chapter examines the pooling process that permits insurance companies to diversify much of the risk in their portfolios. Current interest in the insurance process arises, first, because insurance is an important risk management device and, second, because a similar pooling process arises within the portfolio of risk units held by the non-insurance firm. The interaction between these risk units is important in determining the total level of risk facing a firm. An insurance portfolio is a collection of insurance policies, each of which represents a separate contract with an external party. Whether claims arise under each policy and the size of such claims are random processes. The insurer's interest lies in the aggregate claims to be paid under the portfolio, for it is upon this aggregate that the insurer's financial performance depends. Aggregate claims also follow a random process that is determined by the risk and return characteristics of the individual policies. However, as we see in this chapter, the insurer's risk is not equal to the sum of the individual policy risks.

When insurance contracts in a portfolio are independently and identically distributed, the riskiness of the portfolio, as measured by the standard deviation of the average loss, tends toward zero as the number of policies in the portfolio tends toward infinity. Furthermore, by the Central Limit Theorem, the distribution of average loss tends toward a normal distribution as the number policies increases. This last tendency makes possible the use of the normal approximation method to estimate the probability of ruin. The probability that the average loss for the portfolio will exceed some critical value (determined by the insurer's reserves) in excess of the mean falls as the number of policies increases. As n approaches infinity, the ruin probability approaches zero. Thus both risk measures, the standard deviation and the ruin probability, reveal that portfolio risk is related inversely to size in a portfolio of independent and identical policies.

Broadly similar results prevail when the insurer's portfolio comprises heterogeneous but independent policies. Risk tends to disappear as the number of policies becomes very large. However, this conclusion is qualified, since it assumes that the relative composition of the portfolio is independent of its size.

When policies are not independent, the conclusions on portfolio diversification require substantial modification. First, with correlated policies, the distribution of average loss does not tend toward normal as the number of policies increases. This implies that use of the normal approximation to estimate ruin probability is not strictly appropriate. Even more significant, the riskiness of the distribution of average loss does not tend toward zero as the number of policies approaches infinity. Such portfolios exhibit an irreducible minimum level of risk that cannot be diversified away. The size of the remaining risk depends positively on the correlations between the policies. An example of this problem is earthquake insurance. It is not so much the high expected value of loss that makes insurers nervous about insuring this risk, but rather the high degree of geographic interdependence between individual policies.

The same process of diversification arises with the portfolio of exposure units possessed by a non-insurance firm. The riskiness of the firm's distribution of aggregate value is derived from the characteristics of individual cash flows. But risk does not add up in a simple way. The aggregate risk is less than the sum of its parts. Thus, as part of the firm's risk management strategy, it might well consider the natural risk pooling that arises within the firm. A good example of such pooling is an oil company that owns several hundred service stations. The aggregate loss from, say, fire might be fairly predictable, even though the loss to any individual station is highly uncertain. Another example is a firm that does business in many countries but accounts in dollars. It is subject to foreign exchange risk an many transactions. Suppose each transaction were hedged. This would certainly remove risk, but the transaction costs of all these hedges would be enormous. However, a closer look reveals that many of the exposures might have low or even negative correlations. For er example, one branch might be long in sterling and short in dollars whereas another division is short in sterling and long in dollars. These two positions offset each other (i.e. have a negative correlation). Thus it might be that the overall, portfolio, risk to the firm from foreign currency fluctuations is low. Further opportunities for internal pooling arise when it is considered that the business risk of a firm often exhibits low correlation with risk management risk.

The process of diversification is also important in understanding how investors behave and in determining the value of a firm. Similar concepts will now be put to use to help derive sensible financial criteria for evaluating risk management strategies.

REFERENCES

Beard, R. E., *Risk Theory: The Stochastic Basis of Insurance*, Halstead Press, New York, 1977.

Beard, R. E, T. Pentikainen, and E. Pesonen; *Risk Theory*, Methuen, London, 1969.

Breiman, L.; *Probability*, Addison-Wesley, Reading, Mass., 1968.

Cramer, H.; *Mathematical Methods in Statistics*, Princeton University Press, Princeton, N.J., 1946.

Embrechts, Paul; Kluppelberg, Claudia; Mikosch, Thomas. *Modelling Extremal EventsFor Insurance and Finance.*. *Applications of Mathematics: Stochastic Modelling and Applied Probability*, vol. 33. Heidelberg and New York: Springer. p xv, 645. 1997.

Feller, W.; *An Introduction to Probability and Its Applications*, vol. 2, Wiley, New York, 1966.

Friefelder, L. R.; *A Decision Theoretic Approach to Insurance Ratemaking*, Huebner Foundation Monograph, no. 4, Irwin, Homewood, III., 1976.

Gerber, H.; *An Introduction to Mathematical Risk Theory*, Huebner Foundation Monograph no. 8, Irwin, Homewood, III., 1979.

Houston, D. B.; "Risk, Insurance and Sampling," *Journal of Risk and Insurance*, vol.31, 1964, pp.511-538.

Markowitz, H.: *Porfotlio Selection: Efficient Diversification of Investments*, Wiley, New York, 1959.

Spurr, W. A., and C. P. Bonini: *Statistical Analysis for Business Decisions*, Irwin, Homewood, 111., 1973.

Figure 1: Individual and Aggregate Distributions

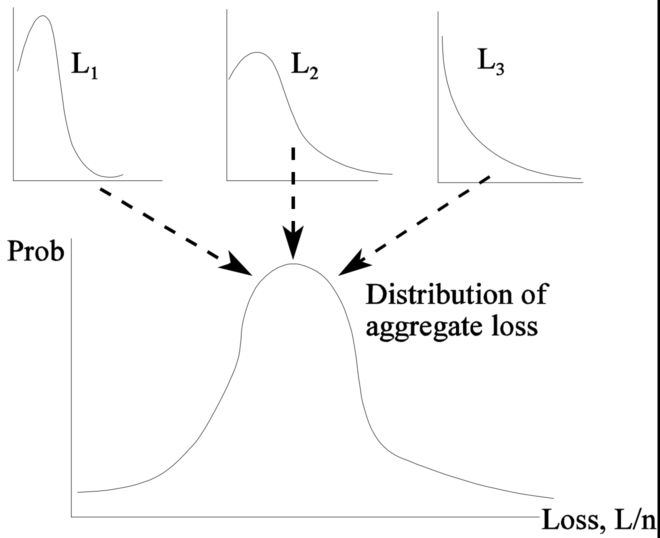

Figure 2: Portfolio Risk Reduction

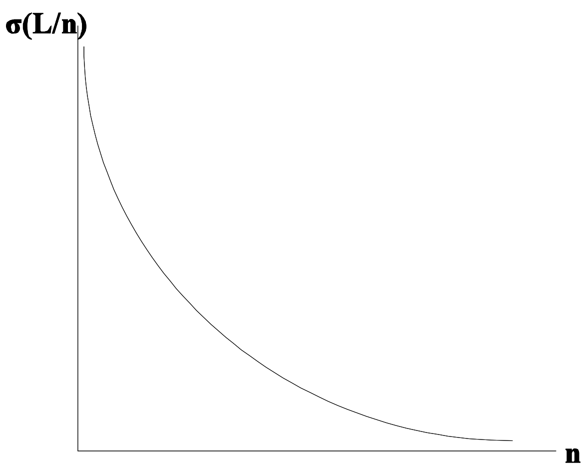

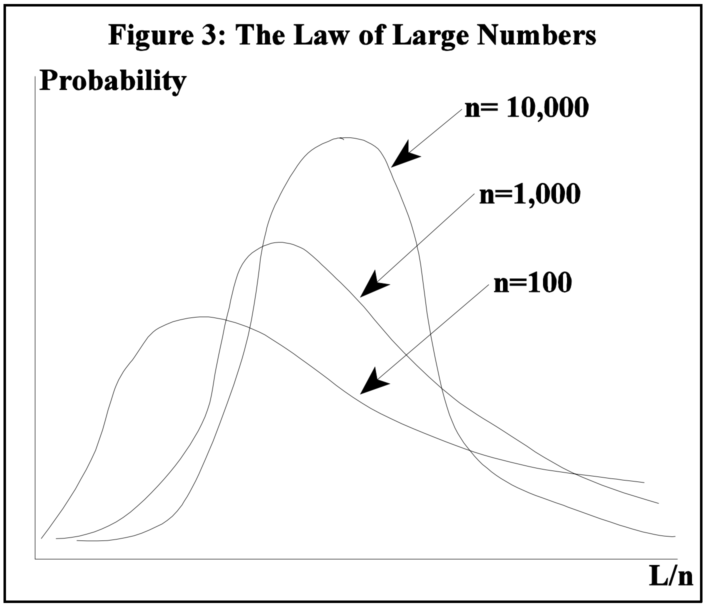
Figure 3: The Law of Large Numbers
Probability
n= 10,000
n=1,000
n=100
L/n

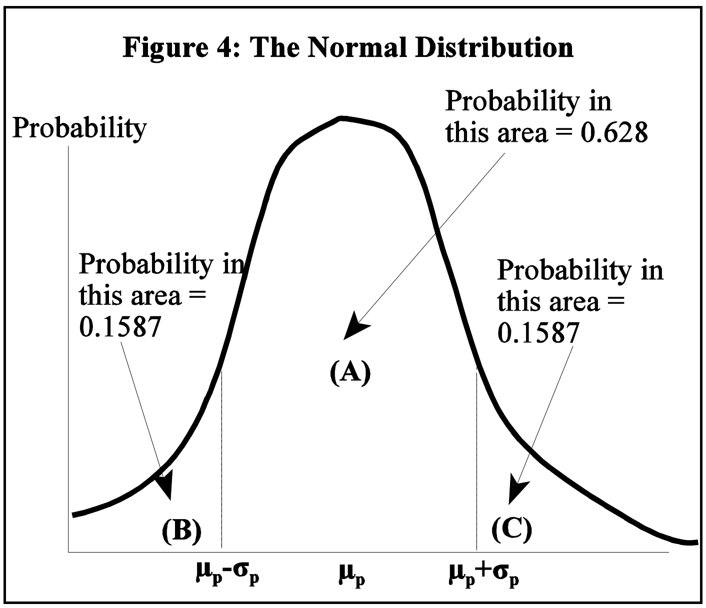
Figure 4: The Normal Distribution
Probability
Probability in this area = 0.628
Probability in this area = 0.1587
Probability in this area = 0.1587
(A)
(B)
(C)
μp-σp
μp
μp+σp

FIGURE 5

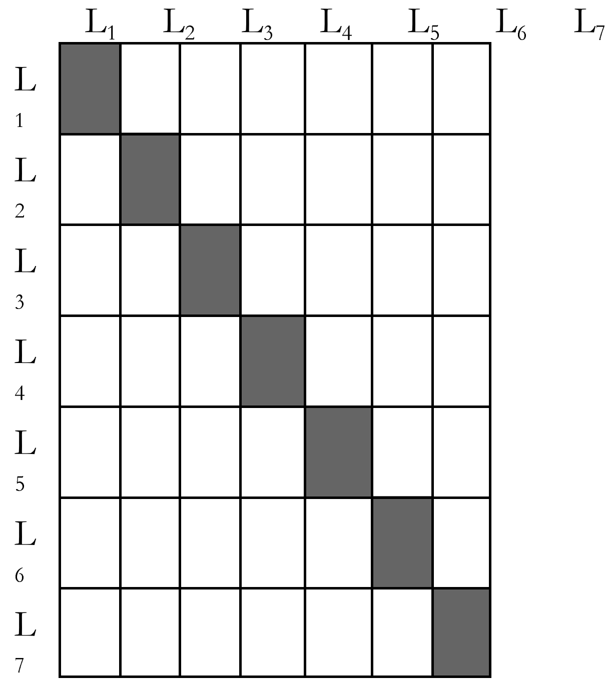

Figure 6. Risk Reduction with Correlated Risk