FINANCIAL SIMULATION MODELS IN GENERAL INSURANCE

BY PETER D. ENGLAND

(Presented at the 5[th] Global Conference of Actuaries, New Delhi, India, 19-20 February 2003)

Contact Address

Dr PD England, EMB, Saddlers Court, 64-74 East Street, Epsom, Surrey, KT17 1HB, UK. E-mail: Peter.England@emb.co.uk

Abstract

Increases in computer power and advances in statistical modelling have conspired to change the way financial modelling is performed in general insurance. Not only have the techniques used to tackle certain problems changed, but also the problems themselves have become increasingly complex. The main driver of change has been the use of simulation techniques to incorporate uncertainty into financial models, or to provide solutions to problems that are intractable or impracticable when attempted analytically. This paper provides an introduction to some of the issues, and provides examples through two case studies.

1. Introduction

Actuarial science is concerned essentially with uncertainty of potential future outcomes associated with frequency of insured events, and severity of the events, given they have occurred.

In life assurance and pensions, frequency is concerned with mortality (or its corollary, survival), and severity is concerned with amounts paid out on death or survival. The amounts paid might be fixed, as defined by the contract, or variable, depending on investment returns.

In general insurance, considerations of frequency and severity components will depend on many factors, most notably, type of insurance. For example, for motor insurance, it is common to consider frequency and severity of possible components of loss, including vehicle damage, bodily injury, theft, fire, third party property and so on. Other classes, such as employer's liability, public liability, aviation hull, marine cargo, space, etc, will have their own characteristics. With reinsurance, it is necessary to consider frequency and severity of the underlying business before superimposing the terms of the reinsurance.

Traditional actuarial mathematics has been concerned with deterministic modelling approaches, looking at "best estimates" of potential outcomes, and stress testing key model inputs to investigate a few adverse scenarios. One of the problems with this approach is the difficulty of assigning probabilities to adverse scenarios, or of answering the question "What is the probability of future experience being worse than our best estimate"?

As soon as we start asking questions about probability, or uncertainty, we are considering risk. The field of Risk Theory has developed to try and answer questions concerning uncertainty. Early work in risk theory moved away from deterministic models to incorporate assumptions concerning uncertainty, but did so in an analytic mathematical framework. That is, complex formulae were developed for aggregate claim size distributions, the probability of ruin and so on. If closed form solutions were not available, advanced mathematical techniques were adopted using expansions, transforms or recursive algorithms to simplify the computations.

Due to the complexity of many insurance problems, traditional risk theory has remained a largely academic discipline, concentrating on relatively straightforward insurance problems (even if the mathematics is complex).

Rather than working with pen and paper, and a vast array of sophisticated mathematical techniques, actuaries are increasingly harnessing the power of computers to build computer models that simulate potential future outcomes of insurance contracts and systems. Computers can generate thousands of possible outcomes, both good and bad, of relevant insurance variables under fairly limited assumptions, in seconds or minutes. Complete distributions of outcomes can then be used to make statements about risk, or used in further calculations. As Daykin, Pentikainen and Pesonen (1996) state:

> *"Modern computer simulation techniques open up a wide field of practical applications for risk theory concepts, without the restrictive assumptions, and sophisticated mathematics, of many traditional aspects of risk theory".*

Simulation (or stochastic) modelling is being used increasingly for asset and economic modelling, catastrophe modelling, reinsurance pricing and optimisation, capital adequacy and allocation, reserving, business planning, enterprise wide risk modelling, and satisfying the demands of regulators, and to some extent, ratings agencies.

This paper illustrates different aspects of simulation modelling through 2 case studies. Insurance pricing is considered in Section 2, and a liability model of an insurance enterprise in Section 3.

2. Insurance Pricing

This case study considers pricing of an insurance contract covering Employer's Liability and Public Liability for a City Council (local government). The problem is considered from the perspective of the insurance company providing the insurance cover.

Employer's liability losses arise from employees suing the council for damages as a result of negligence in respect of health, safety and welfare at work. Claims can be made, for example, for compensation due to injury or death while at work if negligence can be proven. The council pays for smaller claims itself, but buys insurance cover for claims in excess of £100k per event up to a limit of £20million.

In addition, the council can be sued by members of the public for compensation due to injury or pecuniary loss, again where the council can be proven to be negligent in some way. The council pays for smaller claims itself, but buys insurance cover for claims in excess of £100k per event up to a limit of £25million.

The council is aware that it could experience above average levels of loss from a high frequency of smaller losses, not just a single large loss, so it also buys aggregate deductible protection that provides cover for the sum of policy deductibles exceeding £1.25m.

The insurance company has been asked to quote a single premium for the cover, and has been provided information on historic losses for background information.

The company can address the problem using a 'component pricing' approach, that is, consider different types of potential loss and model them separately. Potential future losses can be simulated, from small to large, and the effect on the insurance contract calculated. The risk profile to the insurance contract can then be used to set a premium taking account of risk.

A frequency/severity approach can be used for each type; that is simulate the number of claims, and amounts of claims separately. In this case study, the types of loss considered are:

- Employer's Liability
    - Minor attritional (slips, cuts, etc through to more serious injury) involving only one individual per event
    - Major attritional (major disability or death) involving only one individual per event
    - Minor catastrophe (fire, explosion etc) involving a small number of people
    - Major catastrophe involving a large number of people
- Public Liability
    - Minor attritional (very high frequency, low severity)
    - Major attritional
    - Minor catastrophe
    - Major catastrophe (for example, fire, building collapse at a large public building)

Before simulating, it is necessary to choose statistical distributions for the frequency and severity components of each type of loss, and to choose appropriate values for the parameters of the distributions.

Where historical loss data is available, it can be used to fit a variety of statistical distributions, and the "best" chosen for simulation purposes, taking account of goodness-of-fit, and the underlying characteristics of the distribution being fitted. Parameter estimates can be obtained using a variety of techniques, including method-of-moments and maximum likelihood (see, for example, Klugman *et al*, 1998). Judgement is also needed in finalising parameter selection.

Where data on historical losses is unavailable, for example with the catastrophe loss components, judgement is necessary in choosing suitable frequency and severity distributions and parameters.

In this case study, a Poisson distribution was used for the frequency of each type of loss, where the expected frequency was obtained from observed data. For example, historical experience suggested there were about 30 minor attritional employers liability claims per year, and about 500 minor attritional public liability claims per year. There was no historical information on minor and major catastrophic losses, so their frequency was guessed after discussion with underwriters. For example, a major catastrophic event was assumed to occur with a "return period" of 150 years. From this information, a suitable frequency parameter was obtained. Where a catastrophe event was simulated, the number of people involved was also simulated using a Poisson distribution.

For the severity components, again, (inflated) historical losses were used (where available) to suggest a suitable distribution for simulation purposes. For example, for public liability minor attritional losses, a Gamma distribution was used with a mean of £1,250 and a standard deviation of £1,500, truncated from below such that all simulate losses were above a minimum of £150. For the catastrophe losses, where historic data was unavailable, a simple one parameter Pareto distribution was used, with a lower bound of £50,000 per loss.

After the gross losses have been simulated, it is straightforward to net them down according to the terms of the insurance contract, and calculate the reinsurance recoveries. From the insurer's perspective, the reinsurance recoveries represent losses to the contract, and the distribution can be used to price the contract from the aggregate loss distribution.

Given the aggregate loss distribution, there are several ways of calculating a risk-adjusted price. A basic approach might be to use expected loss plus a multiple of standard deviation: the higher the standard deviation, the higher the price. Certainly, it would be foolhardy to charge less than expected loss. At the other extreme, setting the price at the maximum simulated loss to the contract would result in no risk to the insurer, but would clearly be uneconomical to the insured. We would expect the risk-adjusted price quoted to lie between these extremes, and it is desirable if the pricing methodology adopted can reflect this. One such methodology that has gained popularity was proposed by Wang (1999) and adopts a "proportional hazards"

transform. Essentially, the aggregate loss distribution, expressed as a survival function, is transformed by raising it to a power $1/\rho$. The integral of the transformed survival function over the range gives a risk-adjusted expected value, and in a simulation world, it is straightforward to perform the integration numerically. The level of risk adjustment is controlled through the choice of ρ. When $\rho = 1$, the "adjusted" price is simply the untransformed expected value, and when ρ tends to infinity, the risk-adjusted price tends to the maximum simulated value. Choosing a suitable value of ρ is more art than science, although an analysis of contracts that have been placed previously by the insurer can help. In this case study, a value of $\rho = 1.4$ was used.

A diagram of the modelling process is shown in Figure 1. The basic components of Employer's Liability and Public Liability, Attritional and Catastrophe appear as boxes on the left hand side. At this stage, gross losses are simulated, and are then fed through to the "Dependencies" box that imposes any dependencies between the different loss components. One form of dependency might be to consider that the loss components are correlated in some way. In this case study, it has been assumed that the loss components are independent (a reasonable assumption). It should be noted that in this kind of simulation modelling, assumptions regarding dependence are always made, either explicitly or implicitly. For example if dependence is ignored, it is highly likely that simulated variables are, in fact, considered to be independent, which is a strong assumption. After considering dependence between gross losses, the terms of the insurance contract are imposed, which provides net losses, and losses to the contract. The profile of losses to the contract is then used to calculate a risk-adjusted price.
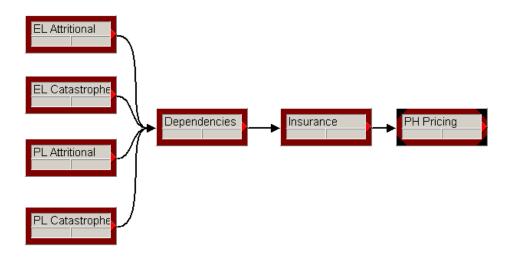

Figure 1: Project diagram for City Council Insurance Case Study

Figure 2 shows the cumulative distribution function of gross losses (red line on right), net losses excluding aggregate deductible protection (green line in the middle), and net losses including the aggregate deductible protection (blue line on left). Essentially

the insurer covers the difference between the blue and red lines, and it can be seen that there are no losses to the insurance contract in about 50% of simulations.
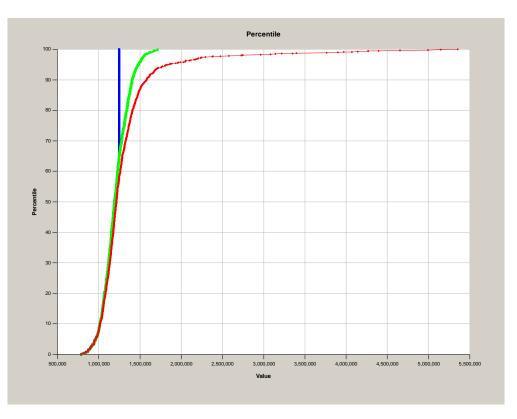

Figure 2. Gross and net loss profile

Table 1 shows a summary of results. The expected loss to the contract is £148,059, with the expected loss above £100k being £105,045. Note this is the expected loss across all simulations: for many simulations, there will be no losses above £100k, and for some simulations, the losses above £100k could be substantial. With a proportional hazards parameter of 1.4, the total risk-adjusted price is £316,043, which is more than twice the expected loss. The difference between the risk-adjusted price and the expected loss represents the cost to the insured of the risk transfer. Using this approach, the risk-adjusted price can also be allocated back to the loss components, such that the sum is the total risk-adjusted price: this is also shown in Table 1.

Table 1. Output Summary

Average loss above £100k	105,045
Average loss to Aggregate Deductible Protection	43,014
Average Total Loss	148,059
Proportional Hazards parameter	1.40
Risk Adjusted Price for losses above £100k	254,394
Risk Adjusted Price for Aggregate Deductible Protection	61,649
Total Risk Adjusted Price	316,043

3. Modelling a Lloyd's Syndicate

This case study considers a liability model of a Lloyd's syndicate. The syndicate was interested in the:

- overall risk profile for the next underwriting year given its business plan, natural catastrophe exposures and reinsurance programme;
- overall capital required to support underwriting in the next year;
- allocation of capital by business division;
- efficacy of the reinsurance programme.

The aims of the analysis required a complex model that could be used to investigate many aspects of the syndicate's structure and operation. For simplicity, this case study considers the liabilities of only one underwriting year on an ultimate basis. A more sophisticated model would also consider the cash flows to ultimate, reserving risk on prior business, asset and liquidity risk, reinsurance bad debt and other credit risks, and so on. All of these additional components can be considered in a simulation environment, but adds to the model complexity.

For business planning and management information purposes, the syndicate was organised into four main divisions, each writing a number of distinct classes of business. The liability model was built to mirror the syndicate's internal structure, since data were available at that level of detail, and results at that level would be meaningful and could easily be compared to traditional performance and planning measures.

To help build and parameterise the model, historical information was obtained regarding aggregate losses by year of account for each class of business within each division. In addition, individual loss information was obtained for large losses for each class, to enable large losses to be modelled individually. Each division operated fairly autonomously, and purchased its own reinsurance cover. The reinsurance arrangements for each division were also obtained to enable accurate modelling of gross and net losses. The reinsurance arrangements were complex in some cases and included features such as:

- Excess-of-Loss
- Multiple Layers
- Different numbers of reinstatements per layer
- Different cost of reinstatements
- Co-insurance (that is, not 100% placed)
- Maximum aggregates
- Backup policies
- Reinstatement protection policies

In addition to the division level reinsurance programmes, there was an Umbrella programme for losses breaching the reinsurance arrangements at division level.

A high level view of the model structure is shown in Figure 3, which shows the division structure, umbrella reinsurance, and capital modelling components. There

are 4 business divisions labelled, UK, Commercial, Marine and Aviation. The UK division writes predominantly motor business, with some small commercial liability business. The Commercial division writes predominantly large commercial property risks on a worldwide basis, and also acts as a reinsurer on property risks. The Commercial division is exposed to natural catastrophes. The Marine and Aviation divisions are self explanatory, and also have significant large loss exposure.
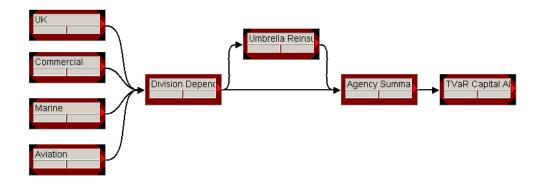

Figure 3: Lloyd's Syndicate High Level Diagram

For more detail, it is necessary to 'drill down' into the boxes shown at the top level, as shown in Figure 4 for the Aviation division. The boxes on the left hand side represent gross results at a class of business level. Dependencies between gross losses at a class of business level are considered before passing the gross losses through the reinsurance programme. For this division, there is a whole account excess-of-loss reinsurance programme that provides cover for aviation event losses aggregated across all classes (that is, a single aviation event can impact multiple classes). Reinsurance recoveries are allocated back to class level to enable loss profiles at that level to be estimated.
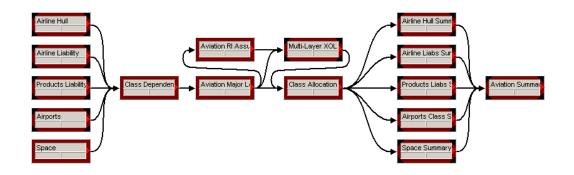

Figure 4: Lloyd's Syndicate Aviation Division

The model was built by first simulating gross losses at a class of business level within each division. The high frequency/low severity "attritional" losses were modelled in aggregate by simulating a gross attritional loss ratio and multiplying by forecast gross premium income within each class. The large losses were simulated individually using a frequency/severity approach, where the frequency of loss was linked to an exposure measure.

Dependencies between aggregate gross losses at a class level within each division were considered before feeding the gross losses through the division level reinsurance programme. Once netted down, the gross and net loss profiles for each class within each division can be investigated, together with the gross and net profiles at a division level.

The syndicate's own data were used to obtain the parameters of the model, supplemented by expert judgement. For example, information was available on approximately 100 major aviation event losses (where a major event was defined as an aggregate loss above £250,000). A Generalised Pareto distribution was fitted to the historical data and used for simulation purposes. A graph of the observed and simulated cumulative distribution functions is shown in Figure 5. It can be seen from Figure 5 that the fit is reasonable. It can also be seen that by using an analytic distribution, it is possible to simulate losses that are more extreme than any observed in the past. This is useful for ensuring that the syndicate has sufficient financial strength to withstand severe adverse deviations from its business plan.
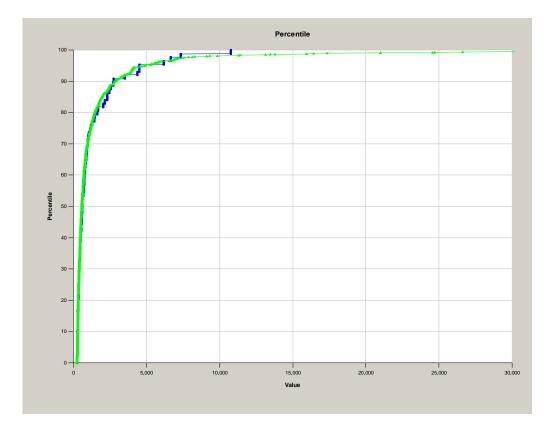

Figure 5. Aviation Large losses – Observed and Fitted Cumulative Distributions

By simulating losses and passing them through the various reinsurance programmes, the efficiency of the programmes can be assessed. For example, a summary of the Aviation reinsurance programme is shown in Table 2. The programme has six layers, with the base premium shown in the first column. After simulating losses to the programme, the average cost of the programme allowing for reinstatements (Column 2) can be compared with the average recovery (Column 3). Allowing for the fact that the reinsurer would be expected to charge more than expected cost for the risk transfer, the cost of the cover seems reasonable at the lower layers. The higher layers seem to cost proportionately more, reflecting the increased uncertainty associated with the higher layers. It could be argued that the top layer is not really needed, but provides 'sleep-easy' cover since the cost is not large relative to the total reinsurance spend.

Table 2: Aviation Reinsurance Summary

(£000)	Initial RI Premium excluding reinstatements	Ave RI Premium including reinstatements	Ave RI Recoveries
Layer 1	2,616	7,009	5,548
Layer 2	5,303	7,714	6,726
Layer 3	1,697	1,907	1,312
Layer 4	2,015	2,067	733
Layer 5	919	923	156
Layer 6	207	207	0

Finally, results from the business divisions were combined to consider an overall profit/loss profile, taking account of dependencies between aggregate gross losses at a division level. Once the overall profit/loss profile is obtained, it is possible to estimate overall capital required, allocate capital to each division (and class of business, if required), and measure the trade-off between reinsurance and capital.

There is a number of ways of assessing capital requirements and adequacy. In this case study, capital was set at a global level using a 'Tail Value-at-Risk' approach such that the expected capital shortfall is zero for a given risk tolerance, in this case a 5% level. If *Free Capital* is defined as:

Free Capital = Starting Capital + Net Premium – Net Claims – Expenses

(ignoring investment income) then Tail Value-at-Risk at the 5% level is the expected value of Free Capital, given Free Capital is below the 5^{th} percentile. This is easily calculated in a simulation environment by calculating the average of the lowest 5% of simulated values of Free Capital. Tail Value-at-Risk is an example of a "coherent" risk measure that has a number of desirable properties making it particularly suitable for capital calculations (see Artzner *et al*, 1999).

Having investigated capital at a global level, it is straightforward to allocate capital to different business divisions, if required. However, overall capital requirements and allocation are very sensitive to assumptions regarding dependencies between classes

of business and between divisions, which highlights the importance of considering dependencies in financial modelling.

Table 3 shows capital requirements and allocation by division for three different scenarios regarding dependence between divisions, keeping dependencies between classes of business within each division fixed. In the first scenario, the divisions are assumed to be independent. In the second scenario a standard correlation approach has been used with the rank correlation between divisions set at 30%. Under the third scenario, the overall rank correlation is approximately 30% between divisions, but extreme "tail" losses tend to occur simultaneously. A discussion of dependence concepts appears in Appendix A.

Table 3: Capital Setting and Allocation – TVaR Risk Measure at 5%

(£000)	Independent	Correlated	Tail Dependent
Rank Correlation	0	0.3	~0.3
Total Capital Required	168,806	228,795	265,445
UK Division	4,092	23,426	30,866
Commercial Division	90,452	85,231	88,521
Marine Division	45,555	60,626	76,075
Aviation Division	33,971	64,459	76,894
Umbrella Reinsurance	-5,264	-4,947	-6,911

£168m of capital required if the business divisions are assumed to be independent. An additional £60m of capital is required if the business divisions are assumed to have pair-wise rank correlations of 30% using a standard correlation approach (tail independence), and a further £37m if the tails are considered dependent, but the rank correlations are maintained at around 30%. It can be seen that the dependence assumptions have a significant effect on estimated total capital requirements

In Table 3, the total capital has also been allocated to business division, in such a way that the divisions contributing the most risk to the global risk profile are penalised the most. The "capital" allocated to the umbrella reinsurance is negative reflecting the trade-off between capital and reinsurance: it is often said that reinsurance is a substitute for capital, and in this case study, that trade-off has been quantified.

If capital is set in this way, it does not mean that there will never be a capital shortfall (that is, capital being insufficient), but that the probability of capital being insufficient is small (and less than 5%). A graph of the cumulative distribution of Free Capital is shown in Figure 6, which shows that the probability of Free Capital being less than zero (giving a capital shortfall) is about 2% in this example. Figure 6 also shows that the potential downside is still large, and adjustments to the reinsurance programmes could be investigated to try and reduce the downside.
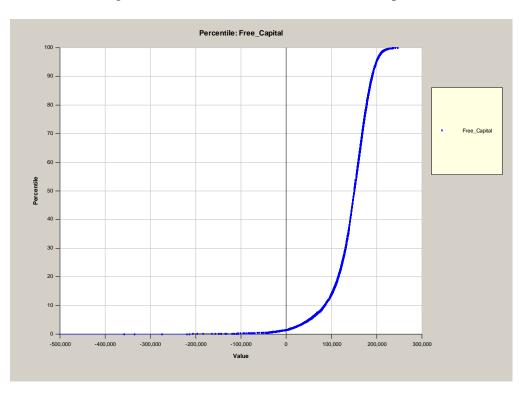
Figure 6: Cumulative Distribution of Free Capital

4. Conclusions

The two case studies described in this paper highlight the potential uses of financial simulation models, and some of their benefits. Although the case studies are relatively simple, they can be extended considerably to include other areas of risk and uncertainty. For example, the liability model described can be extended to include the asset side of the balance sheet, and linked to an economic scenario generator that simulates, for example, inflation rates, short term interest rates and yield curves.

Simulation modelling provides many challenges in model design, assumptions and parameterisation, but also offers significant business benefits. The discipline of looking at all the risks of an insurance operation, and obtaining the relevant data for modelling, provides a thorough understanding of the business. A sound financial model will have many uses, from business planning and quantifying management decisions to demonstrating financial strength to ratings agencies and regulators.

Regulators have seen the potential benefits of "stochastic" modelling and are beginning to encourage, or even insist that insurance companies build their own internal financial models to ensure that companies understand the risks they take, and have sufficient capital to withstand adverse deviations in those risks such that policyholders are protected.

As computer power increases and simulation methods become more widely understood, it is clear that the techniques will be used increasingly in the management

of insurance companies, since simulation modelling removes many of the difficulties previously associated with traditional risk theory.

In fact, many now see stochastic modelling as fundamental to actuarial work. As Chris Daykin, the UK Government Actuary states (see discussion of Wilkie, 1995):

> *"I believe that stochastic modelling is fundamental to our profession. How else can we seriously advise our clients and our wider public on the consequences of managing uncertainty in the different areas in which we work?"*

References

Artzner, P., Delbaen, F., Eber, J.M. and Heath, D. (1999). *Coherent measures of risk*. Mathematical Finance, 9, 3, 203-228.

Daykin, C.D., Pentikainen, T. and Pesonen, M. (1996). *Practical Risk Theory for Actuaries.* Chapman and Hall.

Frees, E.W. and Valdez, E.A. (1998). *Understanding relationships using copulas*. North American Actuarial Journal, 2, 1, 1-25.

Klugman, S.A., Panjer, H.H. and Wilmot, G.E. (1998). Loss Models: from data to decisions. Wiley.

Wang, S. (1995). *Insurance pricing and increased limits ratemaking by proportional hazards transforms*. Insurance: Mathematics and Economics, 17, 43-54.

Wilkie, A.D. (1995). *More on a stochastic asset model for actuarial use*. British Actuarial Journal 1, 777-964.

Appendix A

Dependence Concepts

When considering dependency, most analysts in the financial community think of correlation (usually Pearson correlation). However, there is more to dependency than correlation. Consider Figure A1, which shows four scatter plots of *y* against *x*. The correlation between *x* and *y* in the plot on the top left is zero, although *x* and *y* show strong dependence. This highlights a fact that is well known to statisticians: if *x* and *y* are independent, their correlation will be zero, but the corollary is not true, zero correlation does not imply independence. In the other three plots, the correlation between *x* and *y* is 80%, but the pattern of dependence is very different. The plot in the top right corner exhibits tail independence (that is, although the overall correlation is 80%, the correlation falls towards the tails, and falls to zero in the limit). The plot in the bottom left corner exhibits upper and lower tail dependence (that is, low values of *x* and *y* tend to coincide, and high values of *x* and *y* tend to coincide), and the plot in the bottom right corner exhibits upper tail dependence only.
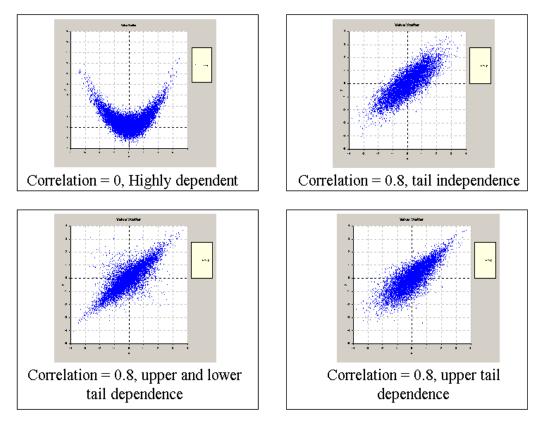

Figure A1. Patterns of dependence

In fact, all three plots with a correlation of 80% share an additional feature: the set of simulated values of *x* is identical in each plot, and the set of simulated values of *y* is identical in each plot. In statistical terminology, the *marginal* distributions of *x* are identical and the *marginal* distributions of *y* are identical, both being Normal with mean = 0 and standard deviation = 1. The only difference between the plots is the pairings of the simulated values of *x* and *y,* which specifies the dependency. The plot

in the top right corner is a standard bivariate Normal distribution, which exhibits tail independence as a theoretical characteristic.

The joint distributions of x and y where the correlation is 80% were generated using *copulas*. A copula is a representation of a joint distribution that separates the marginal distributions from the dependency structure. If $F(x)$ is the cumulative distribution function of X, then

$$F(x_1, x_2, \ldots, x_p) = C\big(F(x_1), F(x_2), \ldots, F(x_p)\big)$$

where $F(x_1, x_2, \ldots, x_p)$ is the corresponding joint distribution and the function $C()$ is a copula function. Any multivariate distribution function can be written in a copula representation (although the form of the copula function is not always obvious). An advantage of using copulas is that each marginal distribution $F(x)$ can be a different distribution.

Although copulas appeared in the statistical literature at least forty years ago, they have only recently been noticed by the financial community. For a good introduction to the use of copulas in actuarial science, see Frees and Valdez (1998).