Sri Lankan Journal of Applied Statistics, Vol.10, 2009, p.1-7

A Very Applied Approach to Risk

J. HENSTRIDGE

**Data Analysis Australia, Adj. Professor of Statistics,
University of Western Australia**

ABSTRACT

When faced with a statistical problem, compromises must often be made between theoretical ideals and practical realities. When most statistical practice was restricted to research contexts, typically academia and institutions, the balance was heavily in favour of the theoretical ideas. Today many applied statisticians work in the commercial world where the balance has shifted in the other direction. Recent world events suggest that the modelling of risk is one area where there has not been a proper balance of theory and practice.

INTRODUCTION

Many modern societies have become very risk adverse, with public statements that they want to achieve zero deaths, injuries or failures in many areas ranging from finance to road safety. To a statistician familiar with handling distributions, the idea of achieving zero risk may seem unrealistic, with a more practical aim being to reduce risk to acceptable levels.

At the same time, reducing risk almost always has a cost and almost always there is more than one way to achieve such reductions. Choices need to be made to decide what level of risk reduction is reasonable and how it may be best achieved. Decisions need to be made, and these decisions will only be rational if it is possible to quantify risk.

Risks are almost always associated with the extremes or tails of distributions. As such, they present a statistical challenge since, by definition, extremes are infrequent and little direct data is available. The usual approach is to adopt a model for the system so that more data or a greater proportion of the existing data can be used estimating the risk parameters of interest. These situations create a second type of risk, that the model used is not sufficiently accurate for the purpose. This is particularly the case where the models are complex and there is the concern that the model choice may have too much influence in determining the outcomes.

George Box's famous quote "Essentially, all models are wrong, but some are useful" is relevant here. The challenge is to find models that are not too wrong, but are sufficiently powerful to make use of actual data to quantify risk.

Before moving to discuss these issues further, it should be noted that these problems are not unique to risk quantification. A classic example in the use of models is the use of microdiamonds in understanding potential diamond deposits where data on the numbers of very small diamonds is used to estimate the likely numbers of large diamonds, the only ones of major commercial interest. The connection between the two sizes of diamonds takes the form of an assumed size distribution, commonly lognormal. If the assumption is correct the result is an enormous reduction in the cost of data collection. If it is wrong, the estimates might be severely biased.

THE ROLE OF SIMULATION

The feature of risk studies is the need to understand the tails of the distributions. Assuming parametric models for the distributions of key variables creates the possibility that any evaluation of risk is heavily influenced by the choice of parametric form. Simulation is seen as a non-parametric method of avoiding such problems. However simulation means that the role of statistical models is somewhat different from when they are used in estimation or hypothesis testing.

Most simulation models can be considered within the structure illustrated below:

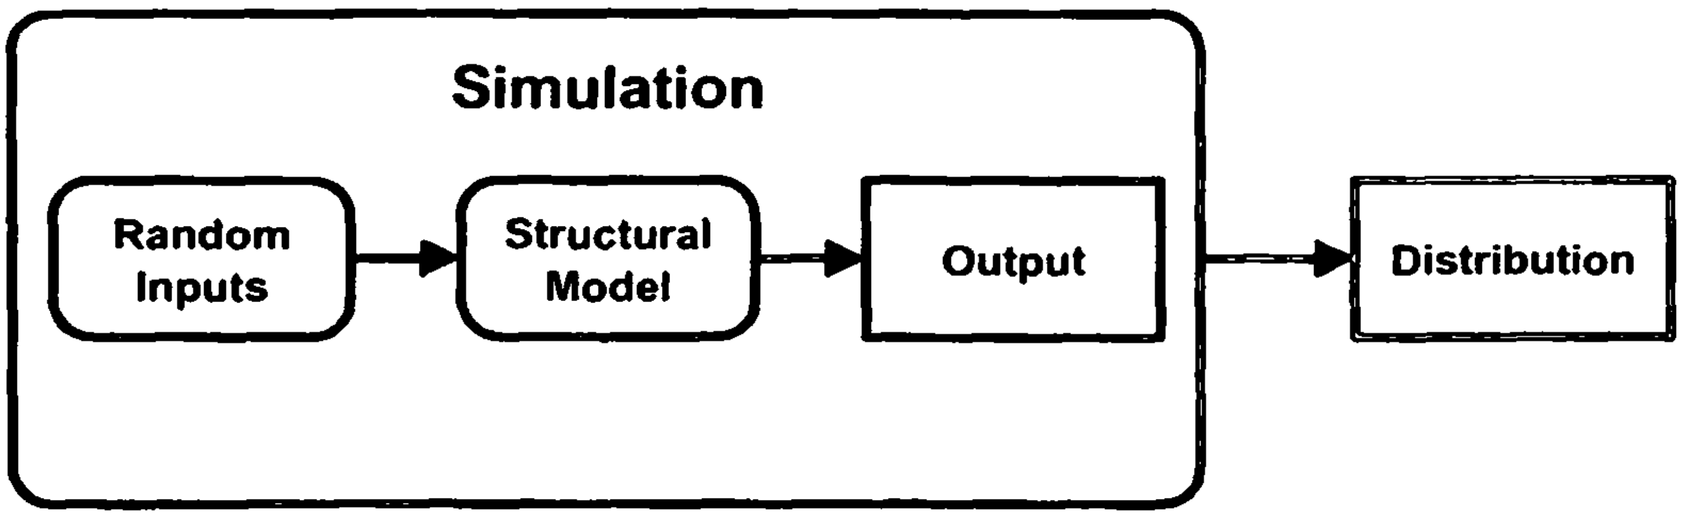

The model itself typically consists of a set of random inputs that drive a structural model that transforms the inputs to the outputs of interest. The whole is imbedded in a simulation environment where the multiple random inputs are applied, each giving the corresponding outputs. The collection of outputs then provides an empirical estimate of the distribution of interest.

Actual simulation studies vary in both their approaches to these components and the level of detail they apply. The two case studies below differ in their solutions to these components.

Case Study 1 – Modelling Investments

Australia has a system by which every employer must contribute to a superannuation fund for each employee, an amount of at least 9% of salary. While much of this is managed through large superannuation investment companies, often industry based, a significant number of employees have "self-managed" funds. In practice these funds usually invest in standard securities – typically company shares, property trusts and cash deposits. Paradoxically, many such employees use a management company to manage their self-managed funds, usually with the aim of tailoring the investments to their specific career and retirement plans.

Data Analysis Australia was approached to build a better model to predict how a particular set of investments might perform for an employee. The aim was to replace a simple "average only" model. What made this complex was that the model had to be able to account for the impact of Australia's complex taxation laws and the need to handle a whole family or partnership. This complexity virtually dictated a simulation approach, with the aim that if the stochastic aspects could be reasonably modelled then they could feed into an arbitrarily detailed financial model to measure the impacts. In principle, the financial model is simply determined by the rules of taxation and accounting.

It was evident at the beginning that any stochastic model would need to work with very limited data. While investments could be reduced to 13 separate asset classes, these were only consistently defined over less than twenty years. Any calibration of a model for likely returns had to be sufficiently simple to be calibrated from this dataset. At the same time, a realistic model needed to account for dependence between asset classes, autocorrelation over time and long tailed distributions. Exploratory data analysis suggested that:

- On a log scale, returns were not readily distinguished from normal distributions, with no evidence of long tails;

- For most asset classes, annual returns showed minimal autocorrelation; and
- Correlations between asset classes were significant.

The first of these findings was actually most problematic, since the power to detect such long tails was meaningless when only such a small amount of data was involved. This contrasts with the more theoretical modelling of long tails in financial transaction problems with, for example, Levy processes.[1] However the difficulty in quantifying the longer tails did not absolve us from including them in our risk calculations – it was not reasonable to have normality or log normality as a null hypothesis to be used unless the evidence to the contrary was strong.

The final choice was to use approximate hyperbolic distributions[2], obtained by considering a variable of the form $z+\alpha z^3$ where z has a normal distribution. By appropriate choice of the coefficient α, it is possible to match the fourth moment of a hyperbolic distribution. (By taking an appropriate higher order odd polynomial it is clearly possible to approximate the distribution to arbitrary accuracy. A theoretical framework for this is found in Hermite polynomials.)

Similarly, the pair wise correlations between asset classes were individually barely significant but collectively they were significant. This led to a model of correlations that had substantial uncertainty in the detail but which captured most of the larger scale variability. Irrespective of the uncertainty, it was clearly a reasonable model for simulation.

Autocorrelation was the one area where a null model was used, since economic theory implies that market driven returns should be unpredictable. The exceptions were cash since these tend to be driven by central bank policy rather than the markets. These did display autocorrelation, albeit small.

The final simulation started with normal varieties, cross correlations created by Choleski factorisations of the required correlation matrix, autocorrelations by simple autoregressions and the long tails by the odd polynomial transformation above.

[1] See for example Schoutens, Wim (2003) *Lévy Processes in Finance: Pricing Financial Derivatives*, Wiley.

[2] Barndorff-Nielsen, Ole (1977). Exponentially decreasing distributions for the logarithm of particle size. *Proc. Royal Society of London. Series A, Mathematical and Physical Sciences* 353 (1674), p 401-409

The initial project plan was to use an off the shelf spreadsheet programme True Blue[3] (then called DSS) to develop a prototype, with a subsequent transition to a purpose writing programme (probably in C) that could carry out the simulations efficiently. The development of this prototype highlighted that this structural side needed to be substantially more complex than originally thought (a typical spreadsheet had up to 2000 rows) and there was the fear that translating the model to a lower level language would reduce the ability to reliably update it. At the same time, the simulation of the returns was far more efficient in True Blue than expected.

While it was not considered in the initial design, it was found that the nature of the random inputs – annual returns over forty years for each of thirteen asset classes – had sufficient opportunity for the long tail effect to show through in the simulations. Hence it was decided to leave the production version in True Blue, particularly once it was found that as few as 200 to 500 simulations gave the required risk envelopes.

The resulting simulation has performed surprisingly well. It cannot claim to have predicted the Global Financial Crisis, but it did give a fair indication that such an event could occur and hence the possibility should be included in financial plans.

Case Study 2 – Electricity Demand

A second example concerns the understanding of electricity consumption in Perth, Western Australia. Electricity demand is, in the Australian context, heavily dependent on the weather, particularly with air conditioning loads. Furthermore, in planning and operating an electricity supply system, it is necessary to allow for a range of demands – the peaks demands must be managed as well as the average since the system has absolutely no energy storage. In this context, we have developed a number of models of varying complexity to answer different needs, so what I discuss here is a hybrid of several of these, one that considers daily demands and uses daily maximum and minimum temperatures.

Here the structural models are very "black box", totally empirical since while it is possible to consider what may affect demand for electricity, the actual processes are beyond a mechanistic model and can only be quantified by reference to the data itself. Fortunately there is ample data, decades of daily demands. The models are therefore fitted using standard statistical techniques of regression and typically have around 200 parameters since they must incorporate very strong annual and weekly

[3] Information is available at www.trueblue.com.au.

cycles, weather terms and the interactions of the three. Fitting such models requires the usual care in selecting variables and deciding on the most appropriate level of complexity. The final models used a seasonally stationary[4] structure for the error term to account for the strong day of the week effect on serial dependence.

The random inputs model was much more problematic, with an initial question being on how to simulate the weather inputs. While in recent decades there has been much work on modelling weather patterns, this has been at a level that gives large area averages rather than local variation. Furthermore, creating a simulation model for weather would have been more difficult than all the other parts of the project combined, and even then it would be seen as the least credible part of the model. Subsequently it was decided to use actual historical weather records, which fortunately were available for over sixty years with reasonable consistency.

Sixty years of weather data would give just sixty simulations, not sufficient by itself to give good tail probabilities. However two steps could extend the value of this beyond what might at first seem a severe limitation. First, the weather is only one part of the random input; the other is the seasonally stationary error component. The effect of having both is to convolute the distributions and smooth out the irregularity of the empirical simulation distribution due to weather alone. Second, a feature of this problem is the strong interaction between the weather and the day of the week, largely due to the effect of commercial air conditioning. Shifting the weather series used in the simulations ranging from lags of three days through to leads of three days does not give seven times as much data but does remove one chance aspect of the simulations, the interaction between extremes of weather and weekdays when the effect is the greatest.

These two steps have maximised the information that can be extracted through simulation from sixty years of data but that is still barely enough. Fitting simple models to the upper tail of the distributions was still necessary for some purposes. But this was enough to help plan maintenance activities while ensuring that there would be a sufficiently low risk of not meeting demand.

One interesting side effect of this study was the discovery of apparent non-stationarity in the historical weather record. This displayed through apparent biases in the simulation results compared with the period over which the model was fitted.

[4] Parzen, Emanuel and Pagano, Marcello, (1979) An approach to modelling seasonally stationary time series, *J. Econometrics*, B, p 137-153

Further investigation linked this non-stationarity to changes in the site where Perth weather was collected – movements of only one hundred metres were sufficient to make a difference. For the simulations to work properly the historical weather record needed to be adjusted for these changes. Subsequent work has used multiple weather stations, both to reduce the dependence on unrecorded changes and to improve the model sensitivities.

DISCUSSION

These two case studies illustrate a non-traditional approach is often required when assessing risk. For example:

- The structural component of the model can often be quite complex, and there is frequently no strong reason to keep it simple, even when empirically estimated. What is more important is that the structural model fully reflects the sensitivity of the outputs to the random inputs.
- The full understanding of the random inputs requires great care. Whereas in many statistical problems the random component is the "noise" that interferes with the understanding of the "signal", in many risk analysis problems it is the random inputs that create the problem itself. Modelling these to a level that allows the simulation of the whole distribution is critical.

The problems discussed here have interest beyond risk. With all statistical modelling, the assessment on whether a model is accurate enough to be useful is often made difficult by the statistician's problem of knowing what is useful in an application area. A model for risk can however be assessed in terms of accuracy of probabilities, an area that is well within the statistician's domain.