The Staple Inn Actuarial Society

Modelling Copulas: An Overview

Martyn Dorey
Phil Joubert

1 Introduction

Identifying and modelling dependencies is a key skill of the risk practitioner. An accurate model of dependency enables us, amongst other things, to model assets and price financial instruments more fairly. Historically measuring and modelling dependencies has centred on correlation. However, it is rare for distributions to follow the strict spherical assumptions with a constant dependence across the distribution implied by correlation.

The good news is that practitioners can now draw on copula techniques to overcome the limitations of correlation. Copula techniques will be familiar to general insurance actuaries, however investment and pensions actuaries are now getting to grips with the technique and applying it to their field.

In this paper we will look at some of the issues in modelling dependencies. We will describe some of the measures used to estimate the strength of relationships and some of the approaches for modelling. Finally, as an example, we discuss the application of copulas to price Collateralised Debt Obligations (CDO's).

A brief history of dependence

In 1888, Francis Galton developed the concept of correlation while working on heredity. However there were gaps in Galton's work, and shortly afterwards a group came forward to fill up the gaps in Galton's work and to extend it in various directions. The most prominent member of this group was Karl Pearson who developed the familiar form of the correlation coefficient in 1896 [1].

$$r(X,Y) = \frac{Cov(X,Y)}{s(X)s(Y)} \qquad (1)$$

An early example of correlation involved an anthropologist investigating whether different bones belonged to the same skeleton by calculating the correlation between the lengths of various bones for each skeleton divided by the length of the skeleton [2]. We will call this type of correlation event correlation.

Since then correlation analysis has been used intensively in the social sciences to ascertain the relationship between occurrences of economic or social events. In addition, correlation's role within regression analysis has led to it's familiarity across a wide range of statistical fields.

The qualitative notion of dependence has a much longer history. In 1686 data maps were being used to infer a physical cause of monsoon rains [3]. Further use of spatial models came later in [4] to model the distribution of particles through a liquid. Although spatial dependence was observed in agricultural fields, most efforts were aimed at removing dependencies prior to any analysis [5]. However, spatial models have gained popularity in the last decade and have been used in areas such as ecology, geology, climatology and environmental science. In forestry, for example, spatial models are used to model patterns of tree growth.

Temporal correlation, dependence of measurements taken at different points from the same process in time, has grown in use since the 1950's. Although many observations have a time dimension, often temporal correlation is ignored, and a cross-section of data is analysed. In the past few years, spatio-temporal models have been used to describe dynamic systems such as ecological and climatic phenomena.

Nowadays, the notion of correlation is central to financial theory. The Capital Asset Pricing Model (CAPM) and the Arbitrage Pricing Theory (APT) use correlation as a measure of dependence between different financial instruments [6]. Furthermore, the importance of correlation has often been emphasised in the context of the pricing of derivatives instruments whose pay-offs depend on the joint realisations of several prices or rates. Examples of such derivatives products are basket of options, swaptions and spread options [7]. The increasing complexity of insurance and reinsurance products has led to interest in the modelling of dependent risks, especially with the emergence of intricate multi-line products [6]. With the ICA requirements of the FSA, robust and defensible approaches to modelling dependencies are required. In Enterprise Risk Management (ERM) the modelling of dependencies between lines of business is critical.

The insurance premium cycle can result in dependent loss ratios from different classes of business. Concentration of risks in a given sector, for example the energy sector, can result in increased claims such as directors and officers (D&O), errors and omissions (E&O), surety and others. Extreme events, such as hurricanes, can also result in dependencies between classes of business which are unrelated in normal conditions. Dependencies can be very tricky to model and sometimes counter intuitive.

In ecology complex dependence structures, built up from several factors, are common. In the financial world, the dependence structures vary with the volatility of the market. The estimation of dependence in non-volatile conditions can be very tricky depending on the amount of data available, the quality of the data available and complexity of the dependence structure.

2 Spurious correlations

Pragmatists may impose a form of dependence rather than trying to empirically determine and validate a structure. Whether a dependency structure is determined empirically or pragmatically the modeller should be wary of spurious structures that may emerge.

In particular, it is possible to obtain a significant value for a coefficient when in reality the two functions are absolutely uncorrelated. An important example is the calculation of correlation coefficients of two price indices. Raw prices may show a strong correlation, but stripping inflation out may result in two uncorrelated indices.

Fallacies can also be caused by mixing different records. Suppose that a drug is effective only on women and the population tested is predominantly men. In this case a spuriously high correlation is obtained only because some women are present in the sample [8, 9, 10]. Correlating time series can also produce spurious correlation especially due to noise or finiteness of the time series [11].

3 Measures of dependency

Suppose we have a pair of datasets (X,Y) and we wish to empirically determine the dependency between them. There are a number of methods for estimating the correlation coefficient and we will look at some of the most common ones. The most common approach is the Pearson's moment approach. This assumes a bivariate normal distribution and a linear relationship such as the one illustrated below.

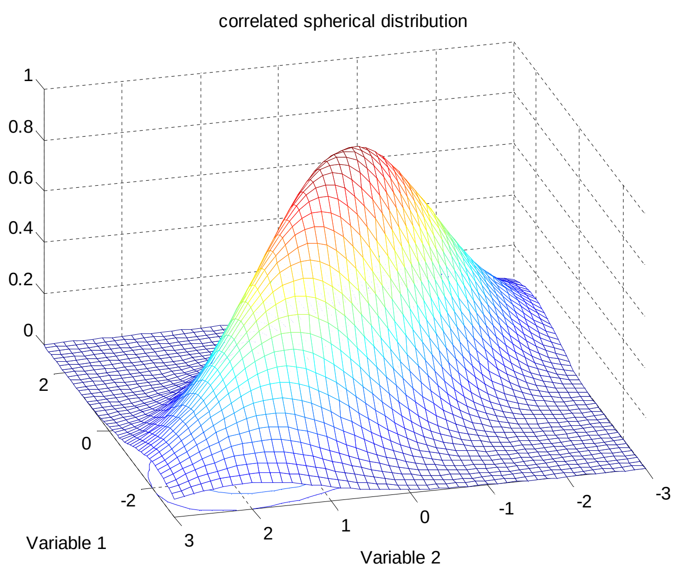

However, what if our data are drawn from a joint distribution which has a non-linear relationship? The familiar correlation coefficient might be zero, despite the existence of a strong link between the data sets. Consider as an example X ~ uniformly on (0,1). The (Pearson) correlation coefficient between the sets X and X^3 is around 90% - despite the fact that our data are <u>deterministically</u> related!

Let us briefly think about the desirable characteristics of a measure of dependence:

- it should be normalised to take values between -1 and 1. While the exact normalisation is arbitrary, this range would agree with our intuition gained through familiarity with the traditional correlation coefficient
- it should be scale invariant
- it should take the value 0 for data sets which are independent, and one of the two extreme values for deterministically related data, whether monotonic or not.

If you are thinking that even meeting this small list is a fairly tall order – you would be correct! Several alternative "correlation coefficients" have been proposed. Two of the most popular are Kendall's tau and Spearman's rho.

Kendall's tau is defined as

$$t = 1 - \frac{4Q}{n(n^2 - 1)}$$

where Q is the number of inversions between the rankings of x and y. An inversion is any pair of objects (i,j),(i',j') such that rank_i-rank_j and rank'_i-rank'_j have opposite signs. (Rank is the "order" of the observation in the set, the rank of the largest observation is 1, the second largest has rank 2, etc)

For example, since $f(x) = x^3$, $x \varepsilon (0,1)$ is monotone, the difference between all pairs of ranks in any random draw from this data set will be zero. Thus no 2 pairs will have a different sign, i.e. Q = 0. Hence $\tau = 1$, as desired.

The Spearman's rho is defined as

$$u = 1 - \frac{6D^2}{n(n^2 - 1)}$$

where

$$D^2 = \sum_1^n \left(r_{y_i} - r_{x_i}\right)^2$$

where r_z denotes the rank of observation z.

Again, (X, X^3) has $\rho = 1$. The point about the monotonicity of $f(x) = x^3$ is critical here – both Spearman's rho and Kendall's tau are measures of monotone dependence, in the same way as Pearson's r is a measure of linear dependence. Data which is deterministically related by a non-monotonic function will still fool these measures. (X, sin(X)) where X ~ Uniformly on $(0,2\pi)$ will return $\tau = 0$.

Both of these measures are statistics, and are subject to sampling error (the equalities in the last example will be true as the sample size tends to infinity). In applications to experimental data the standard error of these coefficients can be estimated using bootstrapping, see [12].

Note that while the former can give values which are very different from the Pearson r, the latter can be numerically identical to it, especially if applied blindly

There are other alternatives for estimating the correlation when non-linearity is suspected. In [13] the authors note that for a continuous function, f(x) = y, "close" values of x give rise to "close" values of y, and suggest the statistic given by

$$K = \sum_{i,j} I\left[\left|X_i - X_j\right| < d\right] I\left[\left|Y_i - Y_j\right| < e\right]$$

Where I[.] is the indicator function of event . . A large value of K, which can be compared to the value calculated for a reference distribution, will indicate strong relationship. The Moran's coefficient [14] replaces the Xs and Ys by their ranks and then calculates a moment correlation coefficient using the ranks. The coefficient is then compared to a reference distribution. Another rank correlation that is robust to outliers is introduced in [15] and in [16] another robust correlation coefficient is proposed and is based on the median.

4 Copulas

One method of modelling dependencies which has become very popular recently is the copula. The word copula is a Latin noun which means ‘a link, tie or bond’, and was first employed in a mathematical or statistical sense by Abe Sklar [19].

Mathematically, a copula is a function which allows us to combine univariate distributions to obtain a joint distribution with a particular dependence structure.

To illustrate how copulas work, we first recall how one uses the CDF of a distribution to generate a random sample: to simulate a univariate distribution you would start by sampling from a uniform U(0,1) distribution. Then by treating this as an observation of your variables CDF you can obtain a sample from a PDF. See below:

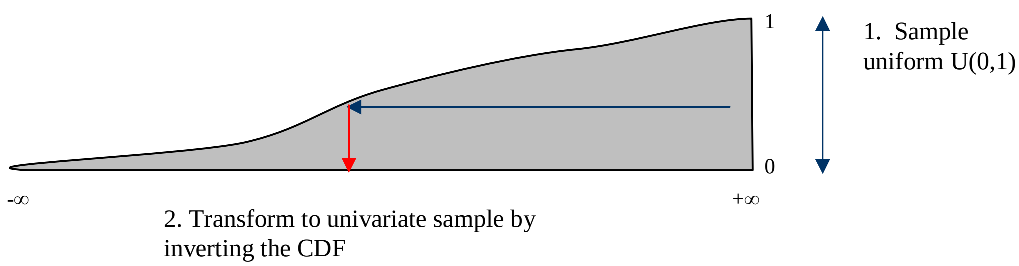

Copulas extend this method to two distributions:

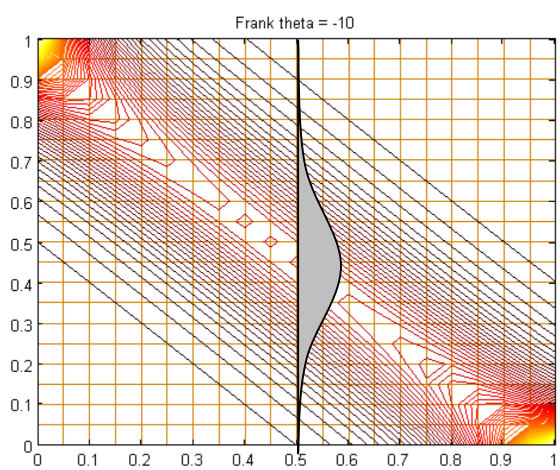

The 50th percentile on the x axis is probabilistically linked to a CDF sample on the y axis. The copula dictates how we go from one distribution to another. This may not necessarily be in a straight-line manner, as is the case in a multivariate normal distribution.

Sklar's theorem, which is the foundation theorem for copulas, states that for a given joint multivariate distribution function and the relevant marginal distributions, there exists a copula function that relates them. In a bi-variate setting:

Sklar's theorem

Let F_{xy} be a joint distribution with margins F_x and F_y. Then there exists a function $C:[0,1]^2 \text{->} [0,1]$ such that

$$F_{xy}(x,y)= C(F_x(x), F_y(y)) \qquad (4.1.1)$$

If X and Y are continuous, then C is unique; otherwise, C is uniquely determined on the (range of X) x (range of Y).

Conversely if C is a copula and F_x and F_y are distribution functions, then the function F_{xy} defined by 4.1.1 is a joint distribution with margins F_x and F_y.

For a proof we direct the reader to [22], the standard introductory text on the subject.

So far we have been rather vague about the function C. Clearly C must be a rather particular type of function.

Definition

C is a Copula if $C:[0,1]^2 \text{->} [0,1]$ and

(i) $C(0,u) = C(v,0) = 0$

(ii) $C(1,u) = C(u,1) = u$

(iii) $C(u_2,v_2) - C(u_1,v_2) - C(u_2,v_2) + C(u_1,v_1) \geq 0$ for all $v_1 \leq v_2$, $u_1 \leq u_2$.

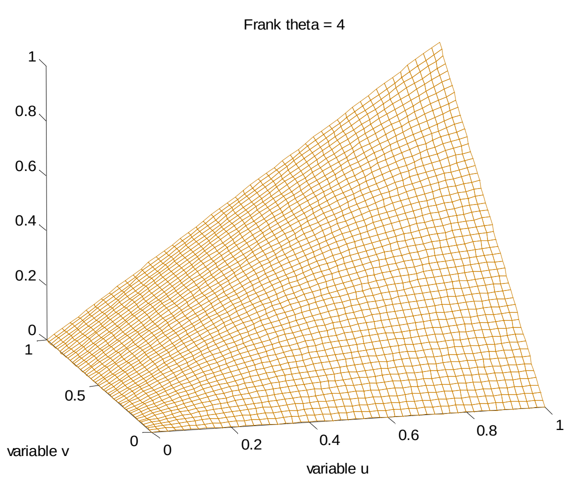

The shape of the Frank copula of the left illustrates that these definitions hold, essentially that it is increasing monotonically on u and v.

If C is differentiable once in its first argument and once in its second, (iii) is equivalent to $\int_{v1}^{v2}\int_{u1}^{u2}\frac{\partial^2 C}{\partial u \partial v}dudv \geq 0$ for all $v_1 \leq v_2$, $u_1 \leq u_2$ in the range. From this observation we see that the definition simply states that a copula is itself a distribution function, defined on $[0,1]^2$, with uniform marginals. Each of the marginal distributions produces a probability of the one dimensional events. The copula function then takes these probabilities and maps them to a joint probability, enforcing a certain relationship on the probabilities.

Using a copula to build multivariate distributions is a flexible and powerful technique, because it separates choice of dependence from choice of marginals, on which no restrictions are placed [20].

Sklar's theorem provides one of the primary ways of constructing copulae. In this case, if X and Y are the marginals, then $C(u,v) = F_{xy}\left(F_X^{-1}(u), F_Y^{-1}(v)\right)$. Other methods use geometric arguments and the definition above to find appropriate functions.

Choosing a copula

A wide, one might say bewildering range of copulas exists. This prompts the question: how does a practitioner choose which one to use? Frequently the choice is based on the usual criteria of familiarity, ease of use and analytical tractability.

The most commonly used copulae are the Gumbel copula for extreme distributions, the Gaussian copula for linear correlation, and the Archimedean copula and the t-copula for dependence in the tail [21, 22]

The Gaussian copula is derived from the multivariate Gaussian or Normal distribution. Other methods of construction may use geometry and the definition above to construct copula functions, such as the Frank copula.

For example, if we have two marginal distributions - one with a beta distribution with parameters α and β, and the other with a lognormal distribution with parameters μ and σ, then we can use a copula, say a member of the Frank's family given by

$$\mathrm{C(u,v)} = \tfrac{-1}{d}\ln\left(1+\frac{(\mathrm{e}^{-d\mathrm{u}}+1)(e^{-dv}-1)}{(e^{-d}-1)}\right)$$

and by substituting the relevant distribution functions, (ie applying 4.1.1) we generate a new joint distribution.

The parameter δ determines the level of dependence of between the marginals.

Contour maps of a Frank Copula function

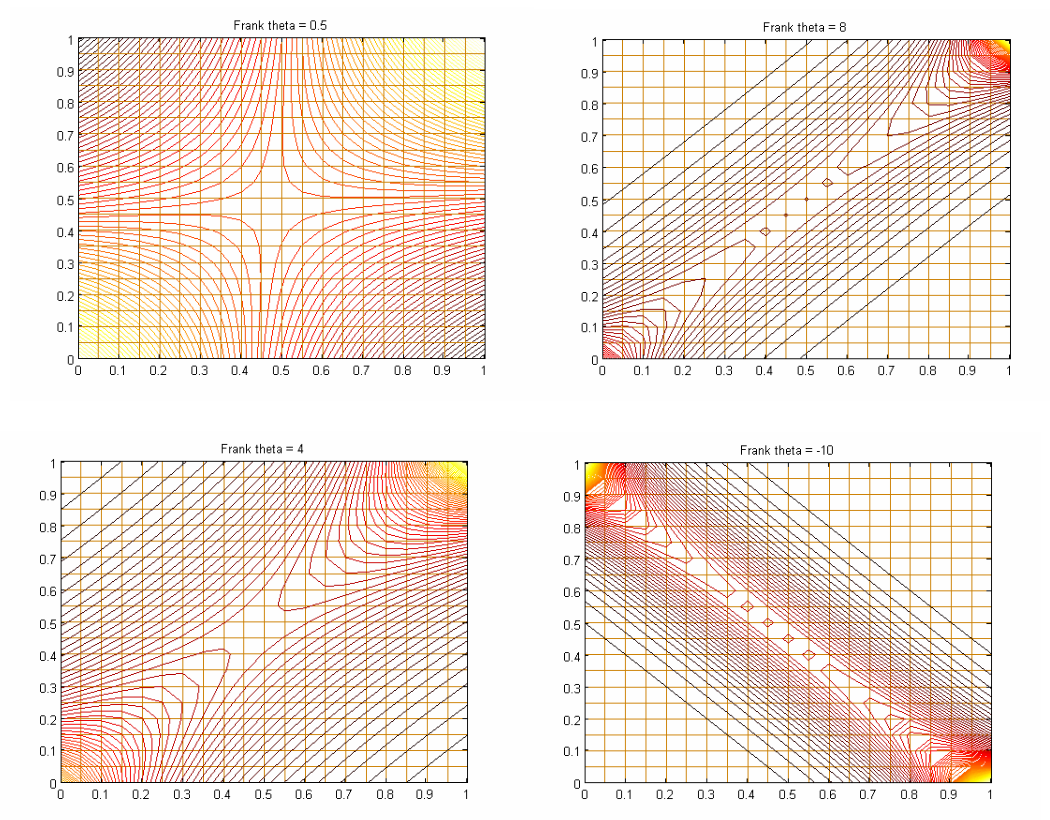

The t copula

The t copula is based on the multivariate t distribution, in the same way as the Gaussian copula is derived from the multivariate normal. This is a straightforward copula to simulate from: first generate samples from the multivariate t distribution (well known algorithms exist to do this), and then invert to cumulative distribution samples.

The t copula is fast growing in usage because the degree of tail dependency can be set by changing the degrees of freedom parameter 'v' [24,25]. Large values for v, say v = 100, approximates a Gaussian distribution. Conversely small values for v, say v==3, increases the tail dependency, until for v = 1 you simulate a Cauchy distribution .

Sampled contour maps of a T copula function

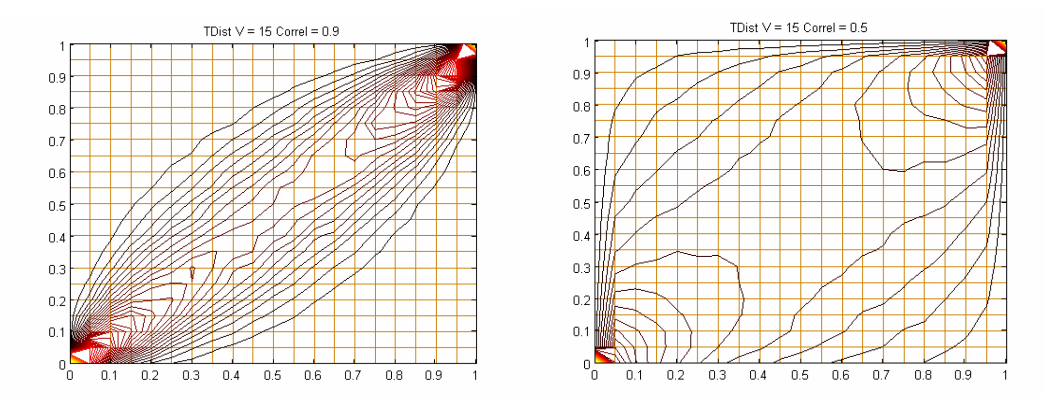

To illustrate the application of the t copula consider the monthly movements on two equity indices, FTSE All World and FTSE All World Specialist Finance as shown below. In this example the historic downward movements in the indices are more dependent than the upward movements, which fan out.

Historic Joint distribution between AW Tech AW Specialist

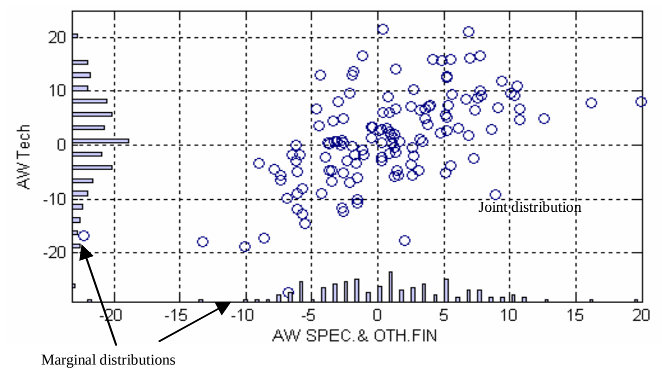

The historic monthly returns are represented as circles on the plot, and the univariate distributions are shown on the x and y axis. From this chart you can see how the term ‘marginal distribution’ came about, the univariate distributions lie on the ‘margin’.

We can model the joint distribution with a t copula and subjectively asses a degree of freedom that is representative of the tail dependency. The charts below show the t distribution with different values for v and 500 simulations.

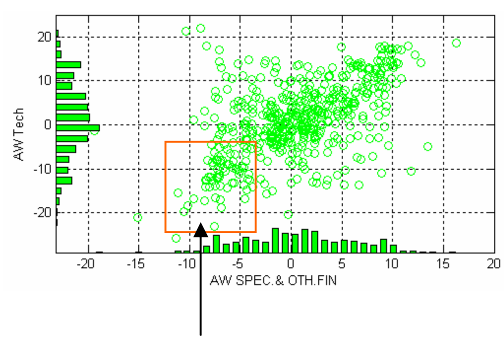

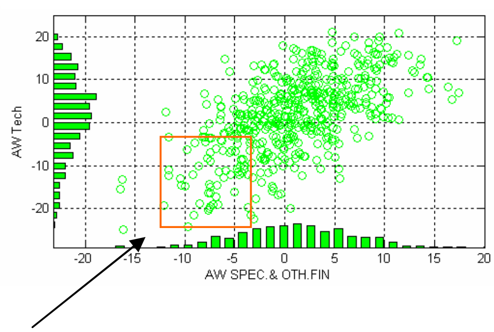

High tail dependency V=3 Lower tail dependency v =500

On these charts we have added some red windows to highlight the downside tails. Clearly the t distribution on the left captures the falls better, within the red box on the left there is higher correlation than the one on the right.

Care needs to be taken when applying the t distribution, as too low a value for the degree-of-freedom parameter v can generate 'wings'. In our example this can be clearly seen when v=1. This may not be a desirable feature.

T Copula, v=1.

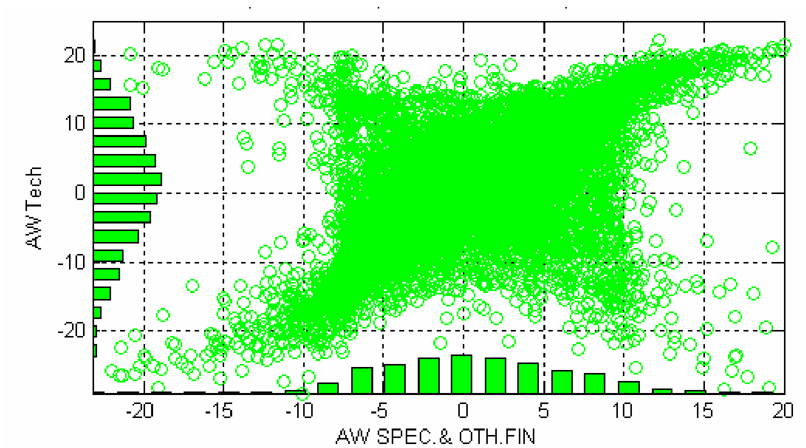

A difficulty with the use of multivariate copulas such as the t copula is that there is only one parameter to control tail association and different pairs of variates might have different tail association [26]. However conditioning offers a route to join different copulas designs together.

A possible drawback of the t-copula is that it has symmetric tail dependency. However more recent work shows that it can be generalised to give asymmetric dependence [23].

Marginal Distributions

Although a great deal of the literature considers the dependency structure between variables, the practitioner will still have to build the marginal distributions. Different approaches can be taken, such as empirical distributions or parametric best fit.

Using empirical distributions results in a cumulative histogram of steps, The discrete nature of the steps is often not desirable. As a result many practitioners start with an empirical distribution and apply a cubic spline or kernel smoothing technique to interpolate between the steps, as can be seen below.

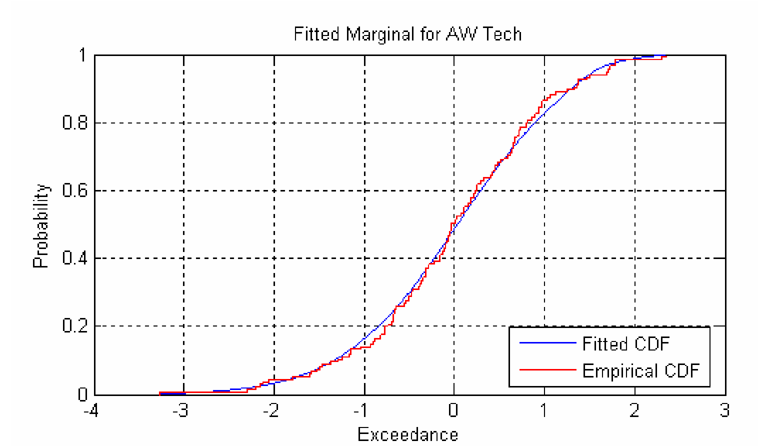

Consideration also needs to be given to the tails. The tails could be an abrupt minimum or maximum, or they can be fitted using Extreme Value Theory (EVT) related techniques; such as a Gumbel distribution [27]. Fitting an EVT tail would be appropriate for equity returns, but would not be appropriate for unemployment, which has a minimum value of 0% and a maximum of 100%.

In addition the modeller can improve the flexibility of a copula through the smart use of pre-processing, for example [24,25] suggest applying a GARCH filter to give i.i.d. observations. The authors found that a simple normalisation of the form (x-mean)/stdev offers an improvement in the dependency estimation and modelling of tails. Indeed pre-processing would remove the need for complicated copulas designs such as the EV t copula described in [23].

.

Conditioning copulas

From the perspective of the practitioner the ability to condition copulas seems to be a very powerful tool. Fitting the copula to all the data is equivalent to fitting a non-linear regression. Forcing certain variables to take a particular value allows the modeller to generate an expected distribution given a variable value Y. This offers benefits over traditional Markov chains.

The authors have found that conditioning allows different copula forms to be bolted together in one model, i.e. different tail dependencies. As the scope of this subject is beyond this paper we have decided to address this subject in more detail separately.

6 Case Study: Pricing CDO’s

The market for credit derivatives and structured products based on those instruments has been growing in recent years. In this section we discuss how copula based models are used to price Collateralised Debt Obligations, or CDO’s. We believe this will be of interest to many actuaries given that many life and pension funds are investing in these products.

In this case study we are describing what we believe to be an emerging industry standard. However, the methods described below have many theoretical shortcomings, and the standard will no doubt change in this fast developing field [26].

What is a CDO?

CDO’s are a wrapper around a basket of corporate debt instruments, for example high-yield bonds. The CDO is divided into “tranches”, or levels of seniority, each with a promised coupon. Payments from the underlying instruments are passed on, via a Special Purpose Vehicle (SPV), to the purchasers of the tranches, starting with the most senior level, and proceeding down. If defaults occur in the underlying assets, the senior levels are (at first) unaffected, with the lower levels losing some or all of their investment. More defaults will mean more layers are “burnt through”, but the senior level’s principle and interest payments will continue unless almost all the underlying instruments have defaulted – generally considered to be an unlikely scenario.

This is illustrated in figure 6.1.

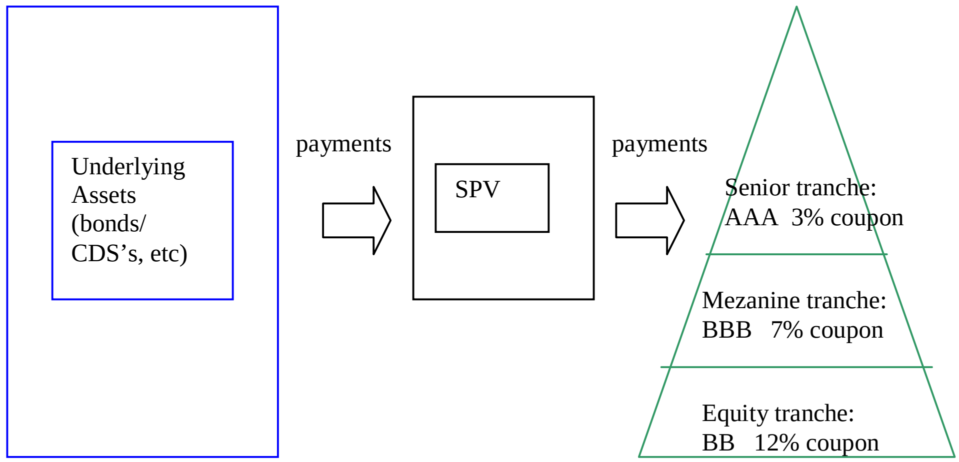

Figure 6.1 A schematic view of a CDO transaction

So a CDO is a method of taking a large number of high risk investments, and creating a several artificial structures (the tranches) with varying levels of risk.

A credit default swap (CDS) is an insurance contract on a corporate debt: one party pays a regular premium, while the other party agrees to pay an amount if a credit event occurs – such as the default of a company on its debt. The premium is set such that the (risk-neutral or market implied) expected value of payment on default is equal to the value of the premium payments.

Let s be the premium rate paid, and R be the "recovery value" (the amount which bond holders will receive on their investment following a default). Then we can write the equation of value (in continuous time) as follows:

$$\int_0^t s e^{-ru} P_u \, du = \int_0^t (1-R) e^{-ru} \, dP_u$$

where P_u is the probability of survival to time s. This equation is more familiar to actuaries in the form:

$$s\bar{a}_{x:\overline{n}|} = \bar{A}^{1}_{x:\overline{n}|}(1-R)$$

Which is the equation of value for a regular premium term assurance. However, for the CDS we are interested in the life time of a bond, rather than the lifetime of a policyholder. In order to price the CDS (ie to find the value of s which solves the above equations) we need to derive or observe the default probability function P_u. In addition this equation can be used in the opposite direction – we use s for different terms to find P_u

The pricing problem

Pricing a CDO, comes down to calculating the (joint) probabilities of default of the underlying instruments. The problem is that we are attempting to price a basket of credit instruments which may have dependent risks of default.

Additionally, the individual risks of default are traded directly in the corporate bond and credit derivatives market. So we need our CDO tranche prices to be consistent with the individual bond prices and CDS rates if we are to avoid arbitrage opportunities.

Note that, as in all pricing problems, it is not the real-world probabilities that are needed, but the risk-neutral probabilities [28].

A Copula based solution.

If we think of the prices of individual bonds and CDS's as reflecting the *marginal* risks of default, and the price of a CDO tranche as reflecting the *joint* risk of defaults, we see immediately that a solution involving copulas is indicated.

Given that the marginal prices (and hence probabilities of default) are observable, we could assume a copula and then either: by observing the price of the CDO, infer the relevant correlation structure, or, by estimating the correlation structure exogenously, calculate the fair price of the CDO.

The copulas most used by market practitioners are the standard Gaussian, the one-factor Gaussian, and sometimes the Clayton copula. We describe in more detail how and why they are used below.

Standard Gaussian copula method

The Gaussian copula is used to generate Monte Carlo simulations of the defaults of the underlying instruments, which are then used to price the CDO [29]. The correlation structure used is the pair-wise asset correlation as used, for example, in CreditMetrics. Often equity correlation, as derived from historical time series, is used as a proxy for asset correlation.

The observed marginal default distributions, together with the correlations and the Gaussian copula, define the joint default distribution. Pricing can then be done by simulating from the joint distribution, and assessing the payouts in each simulation, as in normal Monte Carlo pricing methods.

We summarise by figure 6.2:

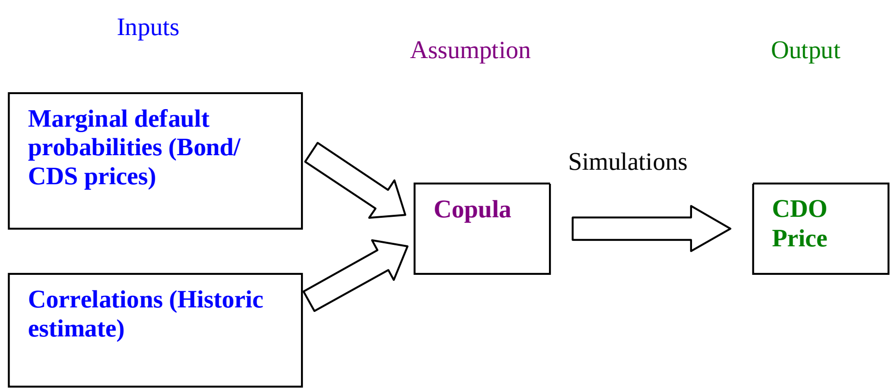

Figure 6.2 Pricing a CDO using a Gaussian Copula

This method is useful since it can be used to price CDO's accurately, given the appropriate marginal distributions and the asset correlation parameters. The disadvantages are that it requires a large number of inputs (the correlation parameters – CDO's can contain hundreds of names) and the computation time required for the Monte Carlo simulation can be onerous. Finally, given a market price for the CDO tranche, we are unable to "invert" this price to find the implied correlation parameters.

One Factor Gaussian and Clayton

Credit risk can also be calculated using a one-factor model [30]. These models postulate that defaults are dependent through a single random factor, often identified with the state of the economy, or some other macro-economic variable. Conditional on the value of that variable, defaults are held to be independent. So during a recession we have more defaults than during a boom, but if we know we are in a recession, company A and company B will default independently (but both with a higher probability than in a boom).

This model further assumes that the underlying portfolio consists of a large number of homogenous risks.

Default occurs in this model for name i if $X_i < c_i$ where X_i is a random variable and c_i is some threshold for the i^{th} asset. This is a simple proxy to the structural models of Merton [31].

We then assume that

$$X_i = r_i \cdot Z + \sqrt{1 - r_i^2} \cdot e_i$$

Where Z is our driving factor and the e_i are mutually independent and independent of Z. Clearly conditional on Z, the X are independent, and the relationship between the X_i's is multivariate normal, with a common correlation ρ.

Because the defaults are conditionally assumed to be independent, the conditional joint density function factorises. The unconditional joint density function can then be found by performing a one-dimensional integration over the possible values of the factor. This integration must usually be done numerically, but the computation time and effort is far smaller than for the Monte Carlo method above.

In this model the marginal default probabilities are correlated with the factor variable by a common amount. Since we no have a single correlation factor, we can invert market (observed) prices for CDO tranches, and solve for the "implied correlation".

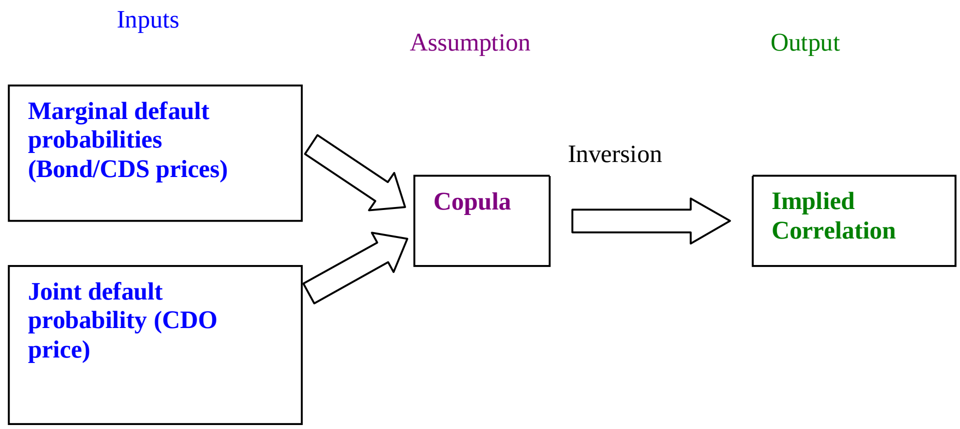

Figure 6.3 Using a one factor model to derive a market implied correlation.

This "inversion" has become market practise, with some talk of implied correlation becoming the credit market equivalent of implied volatility. Indeed, this correlation varies by tranche of a particular CDO, where the model tells us it should be constant, in a way some see as being analogous to the implied volatility smile. Others merely think that this shows that better models are needed.

The Clayton copula is used in a similar way, but is technically more convenient for calculations.

The copulae used as standard in the pricing of CDO's are used primarily for convenience and tractability, rather than for any underlying theoretical reason.

Interestingly, in our research for this section, we noticed that much of credit risk valuation involves the analysis of survival probabilities – an area which should be a speciality for actuaries. Further, the structure of a CDO is similar in many ways to a reinsurance contract, where the underlying risks may be correlated. The pricing of reinsurance contracts is beyond the author's field of knowledge, but we are sure that the profession has something to contribute to this field.

7 Summary and conclusions

When we set out to produce this paper together we expected copulas to be a very new and novel concept. However our first surprise was that copulas have been in practice since the 1970’s. However they have only recently broken into the mainstream, this seems a long slow journey for such a useful technique to progress. Einstein didn’t have to wait that long for the bomb.

Our second surprise was the sheer number of applications for these tools, and their potential to revolutionise regression, garch and artificial intelligence applications. A copula develops an internal representation of associations between many variables at the same time. Conceivably more work on discontinuous copulas, spatial copulas, and feedback could lead to artificial thought that will surpass previous modelling using neural networks. The copula captures the internal workings of thought rather than the physical connections of thought that you see in a neural network.

Some of the material on copulas can be both mystifying and intimidating. However we hope that the reader is left with a better understanding of the principles of copulas.

Perhaps one day actuaries will be replaced by thinking computers powered by copula engines, but who will own them?

About the authors

Martyn Dorey (FIA)

Martyn is a senior consultant at PSolve Asset Solutions. He has10 years experience of research and development within the institutional investment industry. He provides consulting, and licenses techniques on econometric modelling, portfolio construction and asset allocation to investment institutions.

He has a BSc Hons degree in Applied Statistics from the University of Greenwich, and qualified as an actuary in 2004. Prior to joining PSolve in 2003, Martyn worked at a number of leading international investment consultancies, including Head of Quantitative research at Aon Investment Consulting and Quantitative Analyst at Bacon & Woodrow.
martyn.dorey@psolve.com

PSolve Asset Solutions, 126 Jermyn Street, London SW1Y 4UJ..

Phillipe Joubert (FFA)

Phil Joubert was formerly a manager at Deloitte Actuarial and Insurance Solutions, where he has worked extensively with life insurers on risk based capital, market consistent valuation, and other projects.

He is currently working for Ixis CIB in Equity Derivatives and Arbitrage.

He holds a BSc (Hons) in Pure Mathematics from the University of the Witwatersrand, South Africa, and will shortly be completing an MSc in Mathematics and Finance at Imperial College. He qualified as an FFA in 2003.

pj104@imperial.ac.uk

The authors offer their thanks to Coomaren P. Vencatasawmy for his help and contributions toward this paper.

(Coomaren.Vencatasawmy@aviva.com)

Bibliography

[1] Chatterjee, S.K. 2003. *Statistical Thought: A Perspective and History*. Oxford University Press. Oxford.

[2] Kronmal. R.A. 1993. Spurious Correlation and the Fallacy of the Ratio Standard Revisited. *Journal of the Royal Statistical Society A*, **156**, 379-392.

[3] Halley.E. 1686. An Historical Account of the Trade Winds, and Monsoons, Observable in the Seas between and near the Tropicks; with an Attempt to Assign the Physical cause of Said Winds. *Philosophical Transactions*, **183**, 153-168.

[4] Student.1907. On the Error of Counting with a Haemacytometer. *Biometrika*, **5**, 351-360

[5] Fisher. R.A.1935. *The Design of Experiments*. Oliver and Boyd, Edinburgh

[6] Wang S.S. 1997. Aggregation of Correlated Risk Portfolios: Models & Algorithms. Working Paper. CAS

[7] Rebonato, R. 1999. Volatility and Correlation in the Pricing of Equity, FX and Interest Rate Options. John Wiley and Sons. Chichester

[8] Neifeld. M.R. 1927. A Study of Spurious Correlation. *Journal of the American Statistical Association*. **22**, 331-338.

[9] Aldrich, J. 1995. Correlations Genuine and Spurious in Pearson and Yale. *Statistical Science*. **10**, 364-376.

[10] Stigler, S.M. 1989. Francis Galton,s Account of the Invention of Correlation. *Statistical Science*, **4**, 73-79.

[11] Laloux, L., Cizeau, P., Bouchaud, J.P., and Potters M. 1999. Noise Dressing of Financial Correlation Matrices. *Physical Review Letters*, **83**, 1467-1470.

[12] Young G.A. 1988. A Note on Bootstrapping the Correlation Coefficient. *Biometrika*, **75**, 370-373

[13] Babour, A.D. and Costi, M. 1993. Correlation Tests for Non-Linear Alternatives. *Journal of the Royal Statistical Society B*, **55**, 541-548..

[14] Moran, P.A.P. 1948. Rank Correlation and Product Moment Correlation. *Biometrika*, **35**, 203-206

[15] Gideon, R.A. and Hollister, R.A. 1987. A Rank Correlation Coefficient Resistant to Outliers. *Journal of the American Statistical Association*. **82**, 656-666..

[16] Abdullah, M.B. 1990. On a Robust Correlation Coefficient. *The Statistician*, **39**, 455-460.

[17] Rebonato, R. and Jackel, P. 1999. The most general methodology to create a valid correlation matrix for risk management and option pricing purposes. Preprint

[18] Al-Subaibi, A.A. 2000. Simulating Correlated Multivariate Pseudorandom Numbers. Preprint.

[19] Sklar A 1959 Fonctions de repartition a n dimensions et leurs marges. Pul Inst Statist Univ Paris

[20] Clemen R.T. and Reilly, T. 1999. Correlations and Copulas for Decision and Risk Analysis. *Management Science*, **45**, 208-224.

[21] Venter, G.G. Tails of Copula. *Preprint. CAS*

[22] Nelsen R. R. 1999. **An Introduction to Copulas**. Springer, New-York.

[23] The t Copula and Related Copulas, Stefano Demarta & Alexander J. McNeil, Department of Mathematics May 2004 http://www.math.ethz.ch/~mcneil/ftp/tCopula.pdf

[24] MATLAB copula tutorial material. For details contact Steve Wilcockson, Financial Products Account Manager, The MathWorks Ltd, Matrix House,Cambridge Business Park, Cambridge, CB4 0HH Tel: +44 (0)1223 423200 x246 Fax: +44 (0)1223 423289 www.mathworks.co.uk

[25] Nystrom, K., Skoglund, J. (2002), "Univariate Extreme Value Theory, GARCH and Measures of Risk", Preprint, Swedbank.

[26] Venter, G.G. Quantifying Correlated Reinsurance Exposures with Copulas. *Preprint, CAS*

[27] McNeil, A. and Frey, R. (2000), "Estimation of Tail Related Risk Measure for Heteroscedastic Financial Time Series: An Extreme Value Approach". Journal of empirical finance 7, 271-300.

[28] Burtschell, x., Gregory, J., Laurent, J.-P., *A comparative analysis of CDO pricing models*, 2005

[29] Li, David X., *On Default Correlation: A Copula Function Approach*, RiskMetrics, 2000

[30] Friend, A., Rogge, E., *Correlation at First Sight,* Economic Notes: Review of Banking, Finance and Monetary Economics, 2004

[31] Merton, R. (1974). On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance*, 29: 449–470