Copulas and credit models

Rüdiger Frey
Swiss Banking Institute
University of Zurich
freyr@isb.unizh.ch

Alexander J. McNeil
Department of Mathematics
ETH Zurich
mcneil@math.ethz.ch

Mark A. Nyfeler
Investment Office RTC
UBS Zurich
mark.nyfeler@ubs.com

October 2001

1 Introduction

In this article we focus on the latent variable approach to modelling credit portfolio losses. This methodology underlies all models that descend from Merton's firm-value model (Merton 1974). In particular, it underlies the most important industry models, such as the model proposed by the KMV corporation and CreditMetrics.

In these models default of an obligor occurs if a latent variable, often interpreted as the value of the obligor's assets, falls below some threshold, often interpreted as the value of the obligor's liabilities. Dependence between default events is caused by dependence between the latent variables. The correlation matrix of the latent variables is often calibrated by developing factor models that relate changes in asset value to changes in a small number of economic factors. For further reading see papers by Koyluoglu and Hickman (1998), Gordy (2000) and Crouhy, Galai, and Mark (2000).

A core assumption of the KMV and CreditMetrics models is the multivariate normality of the latent variables. However there is no compelling reason for choosing a multivariate normal (Gaussian) distribution for asset values. The aim of this article is to show that the aggregate portfolio loss distribution is often very sensitive to the exact nature of the multivariate distribution of the latent variables.

This is not simply a question of asset correlation. Even when individual default probabilities of obligors and the matrix of latent variable correlations are held fixed, it is still possible to develop alternative models which lead to much heavier-tailed loss distributions. A useful source of alternative models is the family of multivariate normal mixture distributions, which includes Student's t distribution and the generalized hyperbolic distribution. In most cases it is as easy to base latent variable models on these mixture distributions as it is to base them on the multivariate normal distribution.

An elegant way of understanding how a multivariate latent variable distribution determines the distribution of the number of defaults in a portfolio is to use the concept of *copulas*. In this article we show that it is the copula (or dependence structure) of the latent variables that determines the higher order joint default probabilities for groups of obligors, and thus determines the extreme risk that there are many defaults in the portfolio.

If we choose alternative latent variable distributions in the normal mixture family then we implicitly work with alternative copulas which often differ markedly from the copula of a Gaussian distribution. Some of these copulas, such as the t copula, possess *tail dependence* and, in contrast to the multivariate normal, they have a much greater tendency to generate simultaneous extreme values (Embrechts, McNeil, and Straumann 1999). This effect is highly important in latent variable models, since simultaneous low asset values will lead to many joint defaults and past experience shows that realistic credit risk models need to be able to give sufficient weight to scenarios where many joint defaults occur.

This article may be understood as a model risk study in the context of latent variable models. Individual default probabilities and asset correlations are insufficient to determine the portfolio loss distribution, since they do not fix the copula of the latent variables. For large portfolios of tens of thousands of counterparties there remains considerable model risk. Risk managers who employ the latent variable methodology should be aware of this.

2 Latent Variable Models

Consider a portfolio of m obligers and fix some time horizon T, typically one year. For $1 \leq i \leq m$, let the random variable Y_i be the default indicator for obligor i at time T, taking values in $\{0, 1\}$. We interpret the value 1 as default and 0 as non-default. At time $t = 0$ all obligors are assumed to be in a non-default state.

Let $\mathbf{X} = (X_1, \ldots, X_m)'$ be an m-dimensional random vector with continuous marginal distributions representing the latent variables at time T and let $(D_1, \ldots, D_m)$ be a vector of deterministic cut-off levels. We call $(X_i, D_i)_{1 \leq i \leq m}$ a latent variable model for the binary random vector $\mathbf{Y} = (Y_1, \ldots, Y_m)'$ if the following relationship holds:

$$Y_i = 1 \Longleftrightarrow X_i \leq D_i. \tag{1}$$

In the KMV model the latent variables X_i are assumed to be multivariate Gaussian and are interpreted as relative changes in the firm's asset value (so-called asset returns). For determining the thresholds D_i an option pricing technique based on historical firm value data is used. The asset return correlations are calibrated by assuming that asset returns follow a factor model, where the underlying factors are interpreted as a set of macro-economic variables.

CreditMetrics is usually presented as a multi-state latent variable model. The X_i are again assumed to be multivariate Gaussian and their range is partitioned to represent a series of rating classes of decreasing creditworthiness, culminating in default. The cut-off levels which define these classes are chosen so that default and rating state transition probabilities agree with historical data; latent variable correlations are again determined by assuming a factor model structure.

The differences between KMV and CreditMetrics are really differences of presentation rather than differences of substance. The terms defining the model $(X_i, D_i)_{1 \leq i \leq m}$ may be interpreted and calibrated in slightly different ways, but in assuming multivariate Gaussianity of the latent variables the models turn out to be *structurally equivalent*. To understand this assertion we review the concept of copulas and state a simple proposition which is the basis of comparing existing models and defining new and structurally different models.

3 Copulas

Copulas are simply the joint distribution functions of random vectors with standard uniform marginal distributions. Their value in statistics is that they provide a way of understanding how marginal distributions of single risks are coupled together to form joint distributions of groups of risks; that is, they provide a way of understanding the idea of statistical dependence.

There are two principal ways of using the copula idea. We can *extract* copulas from well-known multivariate distribution functions. We can also *create new* multivariate distribution functions by joining arbitrary marginal distributions together with copulas. These ideas are summarised in the following proposition, known as Sklar's Theorem; see Nelsen (1999) for proof.

Proposition 1. *Let F be a joint distribution function with continuous margins $F_1, \dots, F_m$. Then there exists a unique copula $C : [0,1]^m \to [0,1]$ such that*

$$F(x_1, \dots, x_m) = C(F_1(x_1), \dots, F_m(x_m)), \tag{2}$$

holds. Conversely, if C is a copula and $F_1, \dots, F_m$ are distribution functions, then the function F given by (2) is a joint distribution function with margins $F_1, \dots, F_,$.

We extract a unique copula C from a multivariate distribution function F with continuous margins $F_1, \dots, F_m$ by calculating

$$C(u_1, \dots, u_m) = F\left(F_1^{-1}(u_1), \dots, F_m^{-1}(u_m)\right),$$

where $F_1^{-1}, \dots, F_m^{-1}$ are (generalised) inverses of $F_1, \dots, F_m$. We call C *the copula of* F, or of any random vector with distribution function F. The copula of a random vector remains invariant under strictly increasing componentwise transformations of the vector, an appealing property which is not shared by the correlation matrix.

Returning to the credit application, if we assume that the latent variables $\mathbf{X}$ have a multivariate Gaussian distribution with correlation matrix R then the copula of $\mathbf{X}$ may be represented by

$$C_R^{\mathrm{Ga}}(u_1, \dots, u_m) = \mathbf{\Phi}_R\left(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_m)\right),$$

where $\mathbf{\Phi}_R$ denotes the joint distribution function of a standard d-dimensional normal random vector with correlation matrix R, and Φ is the distribution function of univariate standard normal. C_R^{Ga} is known as the Gaussian copula, and this is the latent variable dependence structure which implicitly underlies *all* standard industry models.

In Section 5 we will consider building latent variable models with copulas other than the Gaussian. We conclude this section by noting that it is possible to build latent variable models with the Gaussian copula, but with marginal distributions other than univariate normal and an alternative latent variable model proposed by Li (1999) uses this idea. In this model $X_1, \dots, X_m$ are interpreted as times-to-default for each of the obligors and the thresholds $D_1, \dots, D_m$ are all set to take the value T, the time horizon. Each X_i is assumed to have an exponential distribution with parameter λ_i and the multivariate distribution function F of $\mathbf{X}$ is constructed by using the converse of Sklar's Theorem to join the exponential margins together with a Gaussian copula. This yields the distribution function $F(x_1, \dots, x_m) = C_R^{\mathrm{Ga}}(1 - \exp(-\lambda_1 x_1), \dots, 1 - \exp(-\lambda_m x_m))$.

4 The Role of Copulas in Latent Variable Models

To understand that the use of the Gaussian copula leads to models that are structurally equivalent we introduce a formal definition of equivalence for latent variable models and present a simple new result.

Definition 1. Let $(X_i, D_i)_{1 \le i \le m}$ and $(\widetilde{X}_i, \widetilde{D}_i)_{1 \le i \le m}$ be two latent variable models generating default indicator vectors $\mathbf{Y}$ and $\widetilde{\mathbf{Y}}$. The models are called equivalent if $\mathbf{Y} \stackrel{\mathrm{d}}{=} \widetilde{\mathbf{Y}}$.

Thus two models are equivalent if they give rise to exactly the same default indicator distribution, which means of course that the distribution of the number of defaults in the portfolio will be the same.

A sufficient condition for two latent variable models to be equivalent is that individual default probabilities are the same in both models and the copulas of the latent variables are the same. Formally we have the following, which is proved in Frey and McNeil (2001).

Proposition 2. *Consider two latent variable models* $(X_i, D_i)_{1 \le i \le m}$ *and* $(\widetilde{X}_i, \widetilde{D}_i)_{1 \le i \le m}$ *with default indicator vectors* $\mathbf{Y}$ *and* $\widetilde{\mathbf{Y}}$*. The models are equivalent if*

1. $P(X_i \le D_i) = P(\widetilde{X}_i \le \widetilde{D}_i),\ i \in \{1, \ldots, m\}$,
2. $\mathbf{X}$ *and* $\widetilde{\mathbf{X}}$ *have the same copula.*

Thus KMV, CreditMetrics and the approach off Li (1999) can all be thought of as essentially equivalent approaches. If they are calibrated in consistent ways they will lead to very similar results.

We underline the importance of latent variable copulas in credit risk models by noting that higher order joint default probabilities can be written in terms of copulas and individual obligor default probabilities. Consider an arbitrary subset of k obligors $\{i_1, \ldots, i_k\} \subset \{1, \ldots, m\}$, with individual default probabilities $p_{i_1}, \ldots, p_{i_k}$. Then the joint default probability of all k obligors is given by

$$P(Y_{i_1} = 1, \ldots, Y_{i_k} = 1) = P(X_{i_1} \le D_{i_1}, \ldots, X_{i_k} \le D_{i_k}) = C_{i_1,\ldots,i_k}(p_{i_1}, \ldots, p_{i_k}),$$

where $C_{i_1,\ldots,i_k}$ is a k-dimensional marginal distribution of the copula C of $\mathbf{X}$ (and thus is itself a copula). If we are looking for alternative copulas which lead to higher extreme risk of many joint defaults than the Gaussian, then we should look for copulas which tend to give large values of $P(Y_{i_1} = 1, \ldots, Y_{i_k} = 1)$ for small values of $p_{i_1}, \ldots, p_{i_k}$.

5 Alternative Latent Variable Copulas

There are many alternative copulas to the Gaussian. We choose to work with the copulas which are implicit in the kinds of multivariate distributions that might be considered natural alternative models for asset values and asset returns. It would also be possible to work with families of simple closed-form parametric copulas such as the Archimedean family(Nelsen 1999).

A popular family of distributions for modelling financial market returns is the family of multivariate normal mixture models. When relaxing the assumption of multivariate normality for asset returns it seems natural to look at this family, which contains such distributions as the multivariate t and the hyperbolic.

A member of the m-dimensional family of variance mixtures of normal distributions is equal in distribution to the product of a scalar random variable S and a normal random vector $\mathbf{Z} = (Z_1, \ldots, Z_m)$. That is

$$\mathbf{X} \overset{d}{=} S \cdot \mathbf{Z}, \tag{3}$$

where $\mathbf{Z}$ is multivariate normal with mean vector $\mathbf{0}$ and covariance matrix Σ, and S is positive, independent of $\mathbf{Z}$ and has a finite second moment. Normal variance mixture distributions inherit the correlation matrix of the multivariate normal distribution of $\mathbf{Z}$

$$\mathrm{Corr}(X_i, X_j) = \mathrm{Corr}(Z_i, Z_j),$$

which means essentially that the correlation matrices of these models can be calibrated in exactly the same way as that of the Gaussian model.

For a concrete example we consider the t distribution. $\mathbf{X}$ is said to have an m-dimensional Student t distribution with ν degrees of freedom (written $\mathbf{X} \sim t_m(\nu, \mathbf{0}, \Sigma)$) if

$$S = \sqrt{\frac{\nu}{W}} \tag{4}$$

where W has a chi-squared distribution with μ degrees of freedom.

We choose the t distribution for our analysis for two reasons. First, it converges to the Gaussian distribution as the degree of freedom parameter $\nu \to \infty$. This enables us to start with an approximately normal model and move away from this model gradually by choosing progressively smaller values of ν. Second the copula which is implicit in the multivariate t is very different to the Gaussian copula. It has the property of tail dependence, so that it tends to generate simultaneous extreme events with higher probabilities than the Gaussian copula. This is important in our context, as this leads to higher probabilities of joint defaults.

Figure 1 contrasts the lack of tail dependence of the Gaussian copula with the strong tail dependence of the copula of a t distribution with $\nu = 3$ degrees of freedom. The left hand plot shows 5000 points from a standard bivariate normal distribution; the right-hand plot shows 5000 points from a composite distribution with a t copula and standard normal margins. The linear correlation in both plots is 0.7. Clearly, in the lower left and upper right quadrants, the t dependence structure produces more joint extreme values close to the diagonal

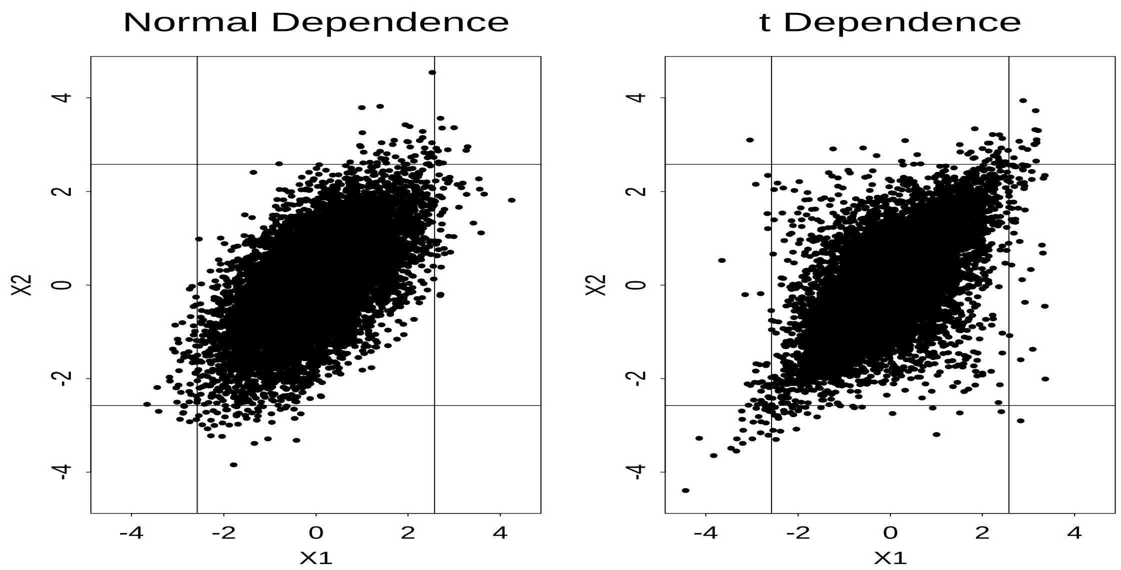

Figure 1: Gaussian dependence vs. t dependence. Vertical and horizontal lines at 99.5% and 0.5% quantiles of marginal distributions.

6 Comparison of the Models

For simplicity we compare the normal and the t copulas in the framework of homogeneous portfolios, where all default probabilities are identical and where the asset correlation of any two counterparties equals a given constant $\rho > 0$. The models are:

1. Gaussian latent variables. $\mathbf{X} \sim N_m(\mathbf{0}, R)$
2. Student t latent variables. $\mathbf{X} \sim t_m(\nu, \mathbf{0}, R)$,

where R is an *equicorrelation matrix* with off-diagonal element ρ.

In both cases we choose cut-off levels so that $P(Y_i = 1) = \pi$, $1 \leq i \leq m$, for some fixed default probability parameter π. Comparison of the models is performed by a simulation study where we vary the portfolio size m, the individual default probabilities π, the correlation of the latent variables ρ and the degrees of freedom parameter ν of the t latent variables.

Group	π	ρ
A	0.01%	2.58%
B	0.50%	3.80%
C	7.50%	9.21%

Table 1: Values of π (default probability) and ρ (asset correlation) for the three groups in the simulation study.

We define 3 groups of decreasing credit quality, which we label A, B, C. The groups are characterised by the parameter settings in Table 1. The π-values do not correspond exactly to the A, B and C rating categories used by any of the well-known rating agencies, but they are nonetheless realistic values for Gaussian latent variable models for real obligors and were chosen after discussions with UBS Switzerland.

m	Group	$m_{0.95}$				$m_{0.99}$			
		$\nu = \infty$	$\nu = 50$	$\nu = 10$	$\nu = 4$	$\nu = \infty$	$\nu = 50$	$\nu = 10$	$\nu = 4$
1000	A	2	3	3	0	3	6	13	12
1000	B	12	16	24	25	17	28	61	110
1000	C	163	173	209	261	222	241	306	396
10000	A	14	23	24	3	21	49	118	126
10000	B	109	153	239	250	157	261	589	1074
10000	C	1618	1723	2085	2587	2206	2400	3067	3916

Table 2: Results of Simulation study. Estimated 95th and 99th percentiles of the distribution of M, the number of defaulting obligors, in an exchangeable model. See Table 1 for the values of π and ρ corresponding to the 3 groups A, B and C.

In all simulations we generate 100'000 realisations of $M = \sum_{i=1}^{m} Y_i$, the total number of defaulting obligors. Of course $E(M) = m\pi$ in all cases, and it is easily confirmed that the empirical average number of defaults is always very close to $m\pi$. Of greater interest are high quantiles of the distribution of M which give a better indication of the extreme risk in the model and are consistent with the Value-at-Risk (VaR) approach to measuring risk. We denote the empirically estimated 95% and 99% quantiles of the distribution of the number of defaults M by $m_{0.95}$ and $m_{0.99}$ respectively and tabulate them in Table 2. In Figure 2 we plot the ratio of estimated quantiles for a Student t model with 10 degrees of freedom and a Gaussian model in the case of Group B and a portfolio of size 10'000.

Clearly ν has a massive influence on these risk measures, particularly for groups of poorer credit quality (B and C). If we only specify the latent variable correlation ρ and do not fix the degrees of freedom ν then our inference concerning extreme risk is subject to huge model risk. For instance, for the 10000 obligors in Group B, Figure 2 can be interpreted as saying that when we move from a Gaussian model to a t model with 10 degrees of freedom, our 95% VaR is inflated by a factor 2.2, our 99% VaR by a factor 4.0, our 99.5% VaR by a factor 4.8 and our 99.9% VaR by a factor 6.1.

7 Extensions and Conclusions

It is clear that the impact of different latent variable copulas on the tail of the distribution of M, the number of defaults, will carry over to the tail of the total loss distribution. Suppose we denote the credit exposure of obligor i to the lender by e_i and the loss given default by the random variable L_i with distribution on $[0, 1]$. The total loss will be given

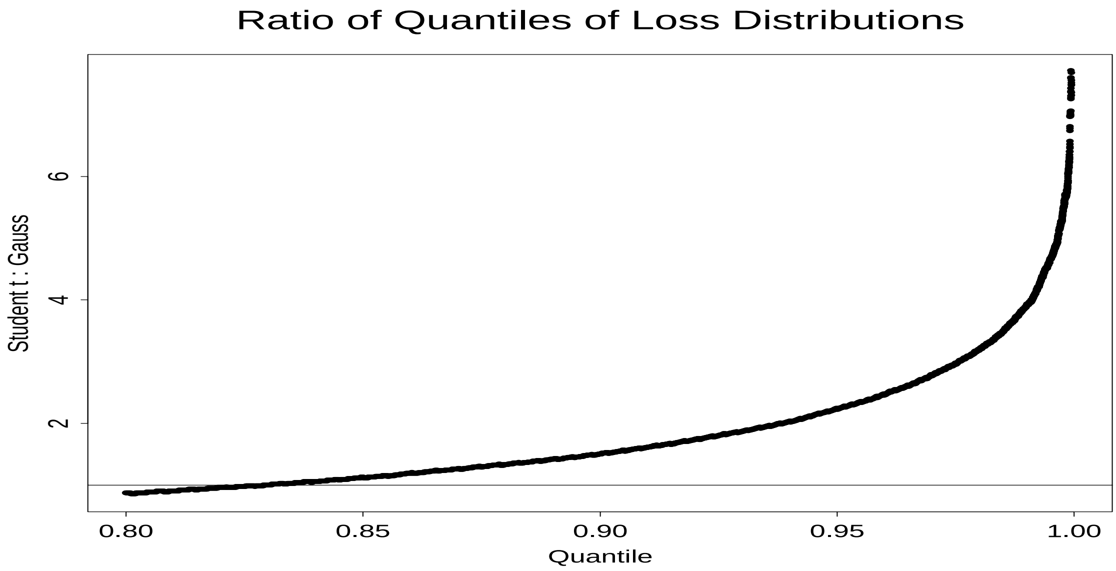

Figure 2: Ratio of estimated quantiles of distribution of M for Student t model with 10 degrees of freedom and Gaussian model in case of Group B with 10'000 obligors.

by

$$\text{Loss} = \sum_{i=1}^{m} Y_i \cdot L_i \cdot e_i.$$

The loss given defaults are usually taken to be independent of each other and independent of the default indicators Y_i. In a model of this kind Nyfeler (2000) has confirmed that the distribution of the latent variables has the anticipated effect on the tail of the total loss distribution.

Moreover, the phenomenon we have demonstrated in a homogeneous portfolio may also be observed in more heterogeneous portfolios consisting of counterparties with widely differing default probabilities and more complex asset correlation matrices. This has been confirmed by simulation studies at UBS Switzerland.

The basic message is that asset correlations are not enough to describe dependence between defaults. Asset correlations do not fully specify the copula of the latent variables and much model risk remains. The assumption of a Gaussian copula may not adequately model the potential extreme risk in the portfolio. Models allowing tail dependence of latent variables (such as the multivariate t copula) show that much more worrying scenarios are possible. Clearly, this finding is also very imprtant for the pricing of basket credit derivatives.

References

Crouhy, M., D. Galai, and R. Mark (2000): "A comparative analysis of current credit risk models," *Journal of Banking and Finance*, 24, 59–117.

Embrechts, P., A. McNeil, and D. Straumann (1999): "Correlation: pitfalls and alternatives," *RISK*, 12(5), 93–113.

Frey, R., and A. McNeil (2001): "Modelling dependent defaults," Preprint, ETH Zürich, available from `http://www.math.ethz.ch/~frey`.

Gordy, M. (2000): "A comparative anatomy of credit risk models," *Journal of Banking and Finance*, 24, 119–149.

Koyluoglu, U., and A. Hickman (1998): "Reconciling the Differences," *RISK*, 11(10), 56–62.

Li, D. (1999): "On Default Correlation: A Copula Function Approach," working paper, RiskMetrics Group, New York.

Merton, R. (1974): "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29, 449–470.

Nelsen, R. B. (1999): *An Introduction to Copulas*. Springer, New York.

Nyfeler, M. (2000): "Modelling Dependencies in Credit Risk Management," Diploma Thesis, ETH Zürich, available from `http://www.Risklab.ch/Papers.html`.