

Modelling Dependent Defaults: Asset Correlations Are Not Enough!

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

March 9, 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. Of particular interest in the normal mixture family are distributions which possess tail dependence, such as the t distribution. In contrast to the multivariate normal, a distribution with tail dependence has a much greater tendency to generate simultaneous extreme values, even when correlation matrices are held constant (Embrechts, McNeil, and Straumann 1999). This effect is highly important in latent variable models, since simultaneous low values of asset values will tend to cause joint defaults of many obligors and this is precisely the phenomenon that leads to extreme risk in large credit portfolios. The large credit portfolio losses that occurred in

particular in the early 90's show 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. For large portfolios of tens of thousands of counterparties there remains considerable model risk. Risk managers who employ the latent variable methodology must be aware of this.

2 Latent Variable Models

Consider a portfolio of m obligors 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 (X_1, \dots, X_m) be an m -dimensional random vector with continuous marginal distributions representing the latent variables at time T and let (D_1, \dots, 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 (Y_1, \dots, Y_m) if the following relationship holds:

$$Y_i = 1 \iff X_i \leq D_i. \quad (1)$$

The total number of defaults in the portfolio will be denoted by the random variable $M = \sum_{i=1}^m Y_i$.

In the KMV model the latent variables are interpreted as relative changes in the firm's asset value (so-called asset returns) and are assumed to be multivariate normal. The asset returns are considered to follow a factor model and estimation of the corresponding covariance matrix is achieved by regressing the asset returns on a small set of observable, macro-economic variables. For determining the thresholds D_i , KMV uses historical data on firm asset values combined with an option pricing technique.

The two-state latent variable model that we have described can easily be generalized to allow more non-default states, such as a graduated rating system. We simply partition the range of the latent variables X_i using a larger set of cut-off levels and assign obligors to rating classes or default according to the value of X_i at time T . CreditMetrics is usually presented as a multi-state latent variable model. Cut-off levels are chosen to agree with a rating state transition matrix, which is usually estimated from historical rating data.

For estimation purposes and to link the credit risk model to market data, KMV and CreditMetrics assume the asset returns X_i are influenced by macro-economic factors through the following p -dimensional ($p \ll m$) factor model:

$$X_i = \sum_{j=1}^p a_{ij} \Theta_j + \sigma_i \epsilon_i. \quad (2)$$

In this model $\Theta = (\Theta_1, \dots, \Theta_p)$ is multivariate normal: $\Theta \sim N_p(\mathbf{0}, \Omega)$. The ϵ_i are independent standard normal and Θ and ϵ_i are independent for all $1 \leq i \leq m$. Θ_j represents a macro-economic (explanatory) variable and ϵ_i represents firm-specific (idiosyncratic) risk.

3 Default Correlation vs. Asset Correlation

Before we focus on alternative models we emphasise the difference between default correlation and asset return correlation (or henceforth simply asset correlation). To do so we need some further notation. The marginal default probability of obligor i we denote by

$p_i = P(Y_i = 1)$ and the joint probability that obligor i and j both default by time T we denote by $p_{ij} = P(Y_i = 1, Y_j = 1)$. Clearly

$$\begin{aligned} p_i &= P(X_i \leq D_i), \\ p_{ij} &= P(X_i \leq D_i, X_j \leq D_j), \end{aligned} \tag{3}$$

and simple calculations give

$$\rho_{ij}^Y := \text{Corr}(Y_i, Y_j) = \frac{p_{ij} - p_i p_j}{\sqrt{p_i(1-p_i)p_j(1-p_j)}}$$

We refer to ρ_{ij}^Y as the default correlation.

We illustrate the relation between asset correlation and default correlation by a numerical example. As in the standard models assume that (X_1, \dots, X_m) follow a multivariate normal distribution. Consequently, the variables (X_i, X_j) are bivariate normal and without loss of generality assume the marginals are standard normal. Let $f_\alpha(u, v)$ denote the standard bivariate normal density function with correlation coefficient α and let ρ_{ij} denote the asset correlation between asset values X_i and X_j . By the distributional assumption on (X_i, X_j) and (3), the joint default probability of company i and j is then

$$p_{ij} = \int_{-\infty}^{D_i} \int_{-\infty}^{D_j} f_{\rho_{ij}}(u, v) du dv.$$

Hence, the asset correlation ρ_{ij} influences the corresponding default correlation by entering in the joint default probability p_{ij} . In the following table we calculated three asset and corresponding default correlations for preset marginal default probabilities $p_i = p_j = 1\%$. Note that default correlation is much lower than the corresponding asset correlation in all cases.

Asset Corr.	Default Corr.
10%	0.94%
20%	2.41%
30%	4.61%

Table 1: Asset correlation vs. default correlation

4 Non-Gaussian Latent Variables

In this Section we replace the assumption of multivariate normal latent variables with the assumption that latent variables have a distribution in the family of multivariate normal variance-mixture models. We go on to look at the special case where $\mathbf{X} = (X_1, \dots, X_m)$ has a multivariate t distribution.

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 $\mathbf{Z} = (Z_1, \dots, Z_m)$. That is

$$\mathbf{X} \stackrel{d}{=} S \cdot \mathbf{Z}, \tag{4}$$

where $\mathbf{Z} \sim \mathcal{N}_m(\mathbf{0}, \Sigma)$ and S is positive, independent of \mathbf{Z} and has a finite second moment. Well known members of this family are the class of symmetric hyperbolic distributions, often applied in finance for modelling stock returns, or the Student t distribution, which

will provide our main example. Normal mixture distributions inherit the correlation matrix of the multivariate normal distribution of \mathbf{Z}

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

which means essentially that the correlation matrices of these models are as easy to calibrate as that of the Gaussian model.

We can build factor models and interpret them in the following way. Assume (4) holds and that the multivariate normal vector \mathbf{Z} has the factor structure given in (2). Then

$$X_i = \sum_{j=1}^p a_{ij}(S\Theta_j) + \sigma_i(S\epsilon_i).$$

Clearly \mathbf{X} also has a factor structure. The macro-economic variables are now $S\Theta_j$ and the idiosyncratic effects $S\epsilon_i$. (The idiosyncratic effects are mutually uncorrelated and are also uncorrelated with the macro-economic variables, which is sufficient to define a factor model.) The additional stochastic term S can be thought of as a shock which affects both idiosyncratic and macro-economic variables.

\mathbf{X} is said to have an m -dimensional Student t distribution with ν degrees of freedom if

$$S = \sqrt{\frac{\nu}{W}} \quad (5)$$

where W has a chi-squared distribution with μ degrees of freedom. Note that \mathbf{X} has mean $\mathbf{0}$ and covariance matrix $\frac{\nu}{\nu-2}\Sigma$. This is usually denoted by $\mathbf{X} \sim t_m(\nu, \mathbf{0}, \Sigma)$.

We choose the t distribution for our analysis for two reasons. First, it converges to the Gaussian distribution as the degree of freedom parameter $\nu \rightarrow \infty$. This enables us to start with an approximately normal model and move away from this model gradually by choosing smaller and smaller values of ν . Second, the t distribution is flexible enough to generate simultaneous extreme events with reasonably high probabilities, i.e. it provides a feature called extremal or tail dependence. This is important in our context, as this leads to higher probabilities of joint defaults.

5 Tail Dependence

Consider a bivariate random vector (X_1, X_2) with marginal distributions F_1, F_2 and quantile functions $\text{VaR}_u(X_1)$ and $\text{VaR}_u(X_2)$ respectively. $\text{VaR}_u(X_1)$ is the u -quantile of X_1 where $u \in (0, 1)$; it may be thought of as the Value-at-Risk of asset X_1 at the $100 * (1 - u)\%$ level. X_1 and X_2 are said to have lower tail dependence if

$$\lambda := \lim_{u \searrow 0} P(X_2 < \text{VaR}_u(X_2) | X_1 < \text{VaR}_u(X_1)) > 0 \quad (\text{assuming } \lambda \text{ exists}).$$

λ is called the coefficient of lower tail dependence. (A coefficient of upper tail dependence may be defined in a similar way.)

If (X_1, X_2) are standard bivariate normal with linear correlation $\rho \in (-1, 1)$, then $\lambda = 0$ and X_1 and X_2 are said to be asymptotically independent. However, if (X_1, X_2) have a standard bivariate t distribution with linear correlation ρ , then

$$\lambda = 2t_{\nu+1} \left(-\sqrt{(\nu+1) \left(\frac{1-\rho}{1+\rho} \right)} \right) > 0, \quad (6)$$

where $t_{\nu+1}$ stands for the distribution function of a univariate t distribution with $\nu + 1$ degrees of freedom.

According to (6) λ is increasing in ρ and decreasing in ν . Hence, even when holding the linear correlation coefficient ρ constant we can adjust the amount of extremal dependence by varying ν . Furthermore, even for a linear correlation of $\rho = 0$ the t distribution shows strictly positive tail dependence. As a consequence, the components X_1 and X_2 cannot be independent.

It can be shown that the extremal dependence property of a bivariate distribution (measured by λ) does not depend on the marginal distributions at all. It only depends on the dependence structure, or the so-called copula of (X_1, X_2) (Embrechts, McNeil, and Straumann 2001). In fact, we can take data from a bivariate t distribution and transform these data so that marginally they have normal distributions. However the transformed data will show a completely different dependence structure to that of a bivariate normal, and in particular they will show extremal dependence.

Figure 1 gives an example. The left hand plot shows 5000 points from a standard bivariate normal distribution; the right-hand plot shows 5000 points from a t distribution with $\nu = 3$ degrees of freedom, which have been transformed to have normal marginal distributions. 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 model 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 $\rho > 0$; in this setup the correlation matrix R of \mathbf{X} is an equicorrelation matrix and \mathbf{X} has effectively a one-factor structure. The models are:

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

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 . Thus in the Gaussian model we set $D_i = \Phi^{-1}(\pi)$ and in

the t model $D_i = t_\nu^{-1}(\pi)$. 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 2: 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 2. 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 3: 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 2 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 , 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 3. 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.




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.

7 Extensions and Conclusions

It is clear that the impact of different latent variable models 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 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. Too much model risk remains. An assumption of multivariate normality may not adequately model the potential extreme risk in the portfolio. Models incorporating extremal dependence of asset values show that much more worrying scenarios are possible.

We are of the opinion that a consideration of economic factor models may be a very useful exercise for identifying groups of counterparties that are susceptible to similar factors, but that this is only one part of the model building exercise. Ideally we should also attempt to use historical default data to estimate true default correlations and incorporate this information in the calibration of models. For further information on the ideas in this paper we refer the reader to Frey and McNeil (2000).

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.
- EMBRECHTS, P., A. MCNEIL, and D. STRAUMANN (2001): “Correlation and dependency in risk management: properties and pitfalls,” in *Risk Management: Value at Risk and Beyond*, ed. by M. Dempster, and H. Moffatt. Cambridge University Press, To appear. Preprint (postscript and pdf versions) available from <http://www.math.ethz.ch/~mcneil>.
- FREY, R., and A. MCNEIL (2000): “Modelling dependent defaults,” Preprint, ETH Zürich, available from <http://www.math.ethz.ch/~mcneil>.
- 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.
- MERTON, R. (1974): “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates,” *Journal of Finance*, 29, 449–470.
- NYFELE, M. (2000): “Modelling Dependencies in Credit Risk Management,” Diploma Thesis, ETH Zürich, available from <http://www.risklab.ch/Papers.html>.