20

Extreme events and multi-name credit derivatives[1]

Roy Mashal, Marco Naldi and Assaf Zeevi

Lehman Brothers; Columbia University

The dependence structure of asset returns lies at the heart of a class of models widely employed for the valuation of multi-name credit derivatives as we saw in the last chapter. In this chapter, we study the dependence structure of asset returns using copula functions. Employing a statistical methodology that relies on a minimal amount of distributional assumptions, we first investigate whether the popular tenet of Normal dependence between asset returns is supported on the basis of empirical observations. We also compare the dependence structures of asset and equity returns to provide some insight into the common practice of estimating the former using equity data. Our results show that the presence of joint extreme events in the data is incompatible with the assumption of Normal dependence, and support the use of equity returns as proxies for asset returns. Furthermore, we present evidence that the likelihood of joint extreme events does not diminish as we decrease the sampling frequency of our observations. Building on our empirical findings, we then describe how to capture the effects of joint extreme events by means of a simple and computationally efficient time-to-default simulation. Using a t-copula model, we analyse the impact of extreme events on the fair values and risk measures of popular multi-name credit derivatives such as nth-to-default baskets and synthetic loss tranches.

INTRODUCTION

The valuation of default-contingent instruments calls for the modelling of default mechanisms. A well-known dichotomy in credit models distinguishes between a "structural approach", where default is triggered by the market value of the borrower's assets falling below its liabilities, and a "reduced-form approach", where the default event is directly modelled as an unexpected arrival.

Currently, a major challenge facing credit models is represented by the rapid growth of multi-name instruments, whose valuation entails modelling the joint default behaviour of a set of reference names. Although both the structural and the reduced-form approaches can, in principle, be extended to the multivariate case, the calibration of the parameters governing the likelihood of joint defaults poses a number of problems.

If we think of defaults as being generated by asset values falling below a given boundary, then the probabilities of joint defaults over a specified horizon must follow from the joint dynamics of asset values. Consistent with their descriptive approach of the default mechanism, multivariate structural models rely on the dependence of asset returns in order to generate dependent default events. In the first part of this chapter, we focus on the empirical properties of the dependence structure – also known as the *copula function* – of asset returns. Roughly speaking, the copula function summarises the dependence structure of a multivariate distribution by "factoring out" the marginals (see the section on *Dependence Structure Modelling*).

The last chapter described a multivariate Normal model for generating the joint default distributions. Indeed several well-known multivariate models assume a joint normal distribution for asset returns. For example, Hull and White (2001) generate default dependence by simulating correlated Brownian motions that are supposed to mimic the asset values dynamics. Similarly, two of the most commercially successful multi-name models, developed by KMV and CreditMetrics, rely on the joint normality of the default-triggering variables.[2] The widespread use of the multivariate Normal distribution is certainly related to the simplicity of its dependence structure, which is fully characterised by the correlation matrix.

A number of recent studies have shown that the joint behaviour of *equity* returns is not consistent with the correlation-based Gaussian modelling paradigm. In particular, extreme co-movements between equities tend to be more adequately described by a fat-tailed dependence structure – eg, the one derived from the Student-t multivariate distribution (the so called *t*-copula). To that end, empirical evidence suggests that correlation, and therefore the Gaussian dependence structure (aka Normal-copula), are not sufficient to appropriately characterise the dependencies between equity returns.[3]

This observation is particularly relevant to the world of default-contingent instruments. The evidence in favour of a fat-tailed dependence structure implies that the likelihood that equity values exhibit large co-movements is higher than correlation-based models would predict. Since the pay-offs of credit instruments are triggered by defaults, and defaults are typically modelled as tail realisations, the increased likelihood of such *extremal events* has substantial implications on the analysis and valuation of these instruments.

Of course, to substantiate the discussion above, one needs to first verify whether the joint behaviour of asset returns is similar to that found in equity returns, as – according to the "structural approach" – it is the former that governs the risk profile of default-contingent instruments. Given that asset returns are not directly observable, it has become customary to proxy asset dependence with equity dependence, and to estimate the parameters governing the joint behaviour of asset returns from equity return series. Fitch Ratings (2003), for example, have recently published a special report describing their methodology for constructing portfolio loss distributions: it is based on a Gaussian copula parameterised by equity correlations.

The use of equity returns to infer the joint behaviour of asset returns is often criticised on the grounds of the different leverage of assets and equity. Even those who accept it as a valid approximation for high-grade issuers are often critical of this approach when applied to low-grade borrowers. This is because a high-quality issuer has a relatively low probability of default, and every variation in the market value of its assets translates almost dollar-by-dollar into a variation of its market capitalisation. On the other hand, high-yield borrowers are closer to the default threshold, and a variation in their asset value can also produce a significant variation in the market value of their debt. This "leverage effect" may generate significant differences in the joint dynamics of equity and asset values.

This chapter has three main objectives.

1. We test whether evidence of extreme events in asset return data is statistically significant, as it was found to be in equity return data. Building on dependence concepts, we find that this question can be answered affirmatively; our statistical study also sheds some light on the magnitude of the error induced by using equity data as a proxy for actual asset returns.
2. Based on these findings, we then illustrate how one can construct simple models that support extremal dependencies in defaults times. We require that such models be easily calibrated to empirical data and be computationally tractable. This culminates in a simple simulation algorithm based on a *t*-copula model of correlated defaults.
3. Finally, based on these models for correlated default times, we illustrate the impact of extreme event modelling on the practical issue of measuring and pricing the risk of two popular multi-name credit products, namely *n*th-to-default baskets and synthetic loss tranches. We close with comments on our results and offer some concluding remarks.

DEPENDENCE STRUCTURE MODELLING

At this point we will briefly describe some dependence modelling concepts to be used in the sequel. The key ingredient in modelling and testing dependencies is the observation that any d-dimensional multivariate distribution can be specified via a set of d marginal distributions that are "knitted" together using a *copula function*. Alternatively, a copula function can be viewed as "distilling" the dependencies that a multivariate distribution attempts to capture by factoring out the effect of the marginals. More formally, a copula function is a multivariate distribution defined over $[0,1]^d$, with uniform marginals. It is possible to confirm that copula functions capture the dependence structure of a multivariate distribution from the following fundamental result; this is known as Sklar's theorem and is adapted from Theorem 1.2 of Embrechts *et al* (2001).

Sklar's Theorem

Given a d-dimensional distribution function H with continuous marginal cumulative distributions $F_1, \ldots, F_d$, there exists a unique d-dimensional copula C: $[0,1]^d \rightarrow [0,1]$ such that

$$H(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)) \tag{1}$$

In particular, the copula is given by

$$C(u_1, \ldots, u_d) = H(F_1^{-1}(u_1), \ldots, F_d^{-1}(u_d)) \tag{2}$$

where $u_i \in [0,1]$ and $F_i^{-1}(\cdot)$ denotes the inverse of the cumulative distribution function F, for $i = 1, \ldots, d$.

By plugging in various multivariate distributions – eg, the Gaussian distribution – for H, one produces the various copulas that underlie these distributions and are derived from them. This study focuses on a natural generalisation of the Gaussian dependence structure, namely the Student-*t* copula. To this end, let t_ν denote the (standard) univariate Student-*t* cumulative distribution function with ν degrees-of-freedom, namely,

$$t_\nu(x) = \int_{-\infty}^{x} \frac{\Gamma((\nu + 1)/2)}{\Gamma((\nu/2)(\nu\pi)^{1/2}} (1 + y_2 / \nu)^{-(\nu + 1)/2} dy \tag{3}$$

Then, for $u = (u_1, \ldots, u_d) \in [0,1]^d$

$$C(u_1, \ldots, u_d; \nu, \Sigma) = \int_{-\infty}^{t_\nu^{-1}(u)} \frac{\Gamma((\nu + d)/2)}{|\Sigma|^{1/2}\Gamma((\nu/2)(\nu\pi)^{1/2}} (1 + y^T\Sigma^{-1} y / \nu)^{-(\nu + d)/2} dy \tag{4}$$

is the Student-*t* copula parameterised by (ν, Σ), where Σ is the correlation matrix, and

$$y = t_\nu^{-1}(u) := (t_\nu^{-1}(u_1), \ldots, t_\nu^{-1}(u_d))$$

The density of the *t*-copula, $c(\cdot; \nu, \Sigma)$, is obtained by differentiating the *t*-copula w.r.t. $u_1, \ldots, u_d$ (see, for example, Embrechts *et al*, 2001).

While the Gaussian copula lies at the heart of most financial models and builds on the concept of correlation, a number of alternative dependence models have been proposed in the literature: the Clayton, Frank, Gumbel and Student-*t* are probably the most notable examples. In this chapter, we focus on the dependence structure underlying the multivariate Student-*t* distribution for a number of reasons.

First, Mashal and Zeevi (2003) compare the different copulas mentioned above in the context of modelling the joint behaviour of financial returns. Applying a formal test, they find that the *t*-copula provides a better fit than the others (see also the study by Bremann *et al* (2003), which also finds empirical support for the *t*-copula using a different statistical test).

Second, the *t*-copula retains the notion of correlation while adding an extra parameter into the mix: the degrees-of-freedom (DoF). The latter plays a crucial role in modelling and explaining extreme co-movements of financial variables, and is of paramount importance for the valuation and risk management of default-sensitive instruments. Moreover, it is well known that the Student-*t* distribution is very "close" to the Gaussian when the DoF is sufficiently large (say, greater than 30); thus, the Gaussian model is nested within the *t* family. The same statement holds for the underlying dependence structures, and the DoF parameter effectively serves to distinguish the two models. This suggests how empirical studies might test whether the ubiquitous Gaussian hypothesis is valid or not. In particular, these studies target the dependence structure rather than the distributions themselves, thereby eliminating the effect of marginal returns that would "contaminate" the estimation problem in the latter case.

To summarise, the *t*-dependence structure constitutes an important and quite plausible generalisation of the Gaussian modelling paradigm, which is our main motivation for focusing on it in this study. For a further discussion of copulas, and the many important characteristics that make them a central concept in the study of joint dependencies, see, eg, the recent survey paper by Embrechts *et al* (2001).

THE DEPENDENCE STRUCTURE OF ASSET RETURNS: EMPIRICAL EVIDENCE AND MODELLING IMPLICATIONS

We continue with a discussion of how asset return data is "backed out" from observable equity return values via the Merton model. We will go on to describe a semi-parametric methodology that allows for the estimation of the dependence structure of a set of returns without imposing any parametric restriction on their marginal distributions. After introducing a test statistic that can be used to evaluate the significance of our point estimates, we apply this methodology to study the dependence structure of asset returns, comparing it with that of equity returns. Next, we present a fully-parametric model with a *t*-copula and *t*-marginals; the empirical results show that, since the univariate *t*-distribution accurately fits the time-series of equity returns, a fully-parametric *t*-model produces estimates of the dependence parameters that are extremely close to the one produced by the semi-parametric methodology. We then provide some concluding remarks and a summary of our empirical results.

The Merton model and implied asset values

Obviously, the main obstacle when attempting to estimate the dependence structure from historical data is that asset returns are not directly observable. In fact, the use of unobservable underlying processes is one of several criticisms directed at the structural approach over the years. To provide a plausible answer to these questions, we first need to "back out" asset values from observable data. One way to estimate the market value of a company's assets is to implement a univariate structural model. Using Merton's (1974) approach – ie, recognising the identity between a long-equity position and the payoff of a call option written on the asset value process – one can apply standard option pricing arguments to derive two conditions that can be simultaneously solved for the asset value of the company and its volatility. This procedure is at the heart of KMV's CreditEdge, a popular credit tool that first computes a measure of distance-to-default and then maps it into a default probability (EDF) by means of a historical analysis of default frequencies. The empirical analysis that follows employs the asset value series generated by KMV's model to study the dependence properties of asset returns.

A semi-parametric estimation procedure and a test for fat-tailed dependence

We now describe an estimation procedure that is used to calibrate a certain class of dependence structures to the equity and asset returns data. Bearing our earlier discussion in mind (see the section on *Dependence Structure Modelling*), the key question that we now face is how to estimate the parameters of the dependence structure. In particular, consider a basket of d names, each following an arbitrary marginal F_i $i = 1, \ldots, d$, and having a joint distribution H with underlying *t*-dependence structure, which is denoted by $C(\cdot; \nu, \Sigma)$ see Equation (4) for the precise parametric form. Here, Σ denotes the correlation matrix and ν the DoF parameter. Suppose we have n observations $\{x_i\}_{i=1}^{n}$ on these d names, where the returns $x_i = (x_{i1}, \ldots, x_{id})$ are assumed to be mutually independent and distributed according to H.

If the marginal distributions were known, then we could use Sklar's Theorem to conclude that $U := F(X_i) \sim C(\cdot; \nu, \Sigma)$, where $F(X_i) := (F_1(X_{i1}), \ldots, F_d(X_{id}))$ is the vector of marginal distributions, the symbol ":=" reads "defined as", and the symbol "~" reads "distributed according to".

Since the structure of the marginals is arbitrary and unknown to us, we propose to use the empirical distribution function as a surrogate, that is,

$$\hat{F}_j(\cdot) := \frac{1}{n}\sum_{i=1}^{n} I\{X_{ij} \leq \cdot\},\ j = 1, \ldots, d \tag{5}$$

where $I\{\cdot\}$ is the indicator function, ie,

$$I(A) := \begin{cases} 1, \text{ if } A \text{ occurs,} \\ 0, \text{ otherwise} \end{cases}$$

We then work with the *pseudo-sample* observations:[4]

$$\hat{U}_i = \left(\hat{F}_1(X_{i1}), \ldots, \hat{F}_d(X_{id})\right),\ i = 1, \ldots, n$$

Focusing on the t-dependence structure $C(\cdot;\nu,\Sigma)$ – formally given by the t-copula in Equation (4) – let us denote by

$$\Theta = \left[(\nu, \Sigma): \nu \in (2, \infty)\right],\ \Sigma \in R^{d\times d} \text{ is symmetric and positive definite}\}$$

the feasible parameter space, and set $\theta := (\nu, \Sigma)$.

Then, for a given pseudo-sample $\{U_i\}_{i=1}^{n}$ we set the *pseudo log-likelihood* function to be

$$L_n(\theta) = \sum_{i=1}^{n} \log c(\hat{U}_i; \theta)$$

where $c(\cdot;\theta)$ is the t-copula density function associated with C. Now, let $\theta := (\hat{\nu}, \hat{\Sigma})$ denote the *maximum likelihood* (ML) estimator of the DoF and correlations, ie, the value of $\theta \in \Theta$ maximising $L_n(\theta)$. This maximisation is generally very involved, and a naive numerical search is likely to fail because of the high dimensionality, $(1 + d(d-1)/2)$, of the parameter space. A simpler way to search for a maximum in this large parameter space is to estimate the correlation matrix using Kendall's Tau (see Lindskog, 2000) and then maximise the likelihood over the DoF parameter. The results reported in this study refer to estimates obtained in this manner. We note in passing that the estimator based on Kendall's Tau is efficient, in the sense that it achieves the minimal asymptotic variance in this context of the t-family – see, for example, Embrechts *et al*, (2001).

As we mentioned earlier, the DoF parameter ν controls the tendency to exhibit extreme co-movements, and also measures the extent of departure from the Gaussian dependence structure. Given its pivotal role, we subsequently focus on the accuracy of the DoF estimates in a more detailed manner.

Specifically, we use a *likelihood-ratio* formulation to test whether empirical evidence supports or rejects a given value of ν. To begin with, we fix a value of the DoF parameter ν_0 and consider the hypotheses H_0: $\theta \in \Theta_0$ vs. H_1: $\theta \in \Theta$ where $\Theta_0 = \{\theta \in \Theta: \nu = \nu_0\} \subset \Theta$.

Then, we set the *likelihood-ratio test statistic* to be

$$\Lambda_n(\hat{\nu} \mid \nu_0) = -2 \log \frac{\sup_{\theta \in \Theta_0} \prod_{i=1}^{n} c(\hat{U}_i; \theta)}{\prod_{i=1}^{n} c(\hat{U}_i; \hat{\theta})} \tag{6}$$

To determine the adequacy of each value of ν_0, we need to characterise the distribution of the statistic $\Lambda_n(\hat{\nu} \mid \nu_0)$. Since this distribution is not tractable, the standard approach is to derive the asymptotic distribution and use that as an approximation. Specifically, Mashal and Zeevi (2002) arrive at the approximation

$$\Lambda_n(\hat{\nu} \mid \nu_0) \approx (1 + \gamma)\,\chi_1^2 \tag{7}$$

where $\gamma > 0$ is a constant that depends on the null hypothesis, χ_1^2 denotes a random variable distributed according to a Chi-squared law with one degree-of-freedom, and "≈" reads "approximately distributed as" (for large values of n).[5] Thus, we can calculate approximate p-values as a function of ν_0 as follows

$$p - \text{value}(\nu_0) \approx P\left(\chi_1^2 \geq \frac{\Lambda_n(\hat{\nu} \mid \nu_0)}{(1 + \gamma)}\right)$$

By letting the *null hypothesis* ν_0 vary over the parameter space, we can compute the corresponding p-values and detect the range of degrees-of-freedom that are supported (or rejected) by the observed return data. Notice that *large* values of the test statistic correspond to *small* p-values, thus indicating that the hypothesised ν_0 is not plausible on the basis of the empirical observations. This hypothesis testing formulation illustrates the "sharpness" of the estimation results in a much stronger manner than if we had just focused on the associated confidence intervals for the parameter estimates.

Empirical evidence

We now apply the methodology outlined above to study the dependence structure of asset returns and compare it with that of equity returns. For consistency, asset and equity values are both obtained from KMV's database. The reader should keep in mind, however, that equity values are observable, while asset values have been "backed out" by means of KMV's CreditEdge implementation of a univariate Merton model. We use daily data covering the period from December 31, 2000 to November 8, 2002, and focus our attention on two portfolios, the 30-name Dow Jones Industrial Average and a 20-name high-yield portfolio.

DJIA Portfolio

Following the semi-parametric methodology described in the previous section, we estimate the number of DoF of a t-copula without imposing any structure on the marginal distributions. Using the test statistic we introduced, Figure 1 presents a sensitivity analysis for various null hypotheses of "joint-tail fatness", as captured by the DoF parameter. The two horizontal lines represent significance levels of 99% and 99.99%; a value of the test statistic falling below these lines corresponds to a value of DoF that is not rejected at the respective significance levels.

The minimal value of the test statistic is achieved at 12 DoF (ν =12) for asset returns and at 13 DoF (ν =13) for equity returns. In both cases, we can reject any value of the DoF parameter outside the range [10,16] with 99% confidence; in particular, the null assumption of a Gaussian copula ($\nu = \infty$) can be rejected with an infinitesimal probability of error.

Finally, notice that the point estimate of the asset returns' DoF lies within the non-rejected interval for the equity returns' DoF, and vice versa, indicating that the two are essentially indistinguishable from a statistical significance viewpoint. Moreover, the difference between the joint tail behaviour of a 12- and a 13-DoF *t*-copula is negligible in terms of any practical application.

Figure 2 reports the point estimates of the DoF for asset and equity returns in the DJIA basket, as well as for three subsets consisting of the first, middle, and last 10 names (in alphabetical order). The similarities

Figure 1 DJIA portfolio: asset and equity returns test statistics

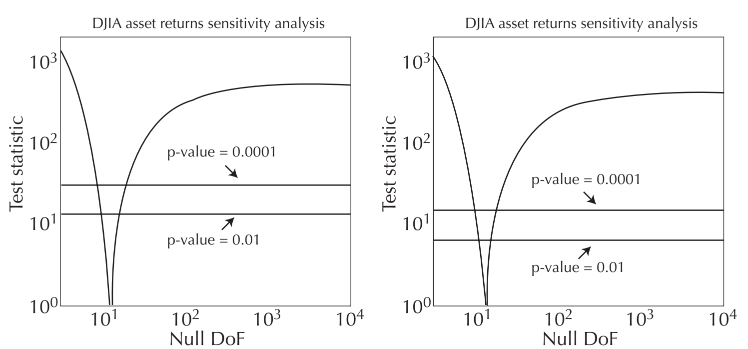

Figure 2 Maximum likelihood estimates of DoF for DJIA portfolios

Portfolio	Asset Returns DoF	Equity Returns DoF
30-Name DJIA	12	13
First 10 Names	8	9
Middle 10 Names	10	10
Last 10 Names	9	9

between the joint-tail behaviour (as measured by the DoF) of asset and equity returns are quite striking.[6]

Next, we compare the remaining parameters that define a *t*-copula, ie, the correlation coefficients. Using robust estimates based on Kendall's rank statistic, we compute the two 30×30 correlation matrices from asset and equity returns. The maximum absolute difference (element-by-element) is 4.6%, and the mean absolute difference is 1.1%, providing further evidence of the similarity of the two dependence structures.

To summarise, the empirical evidence strongly supports the widespread practice of using equity return series to estimate underlying dependencies between asset returns.

High-yield portfolio

We now investigate whether the similarities between the dependence structures of asset and equity returns persist when we restrict our attention to lower-quality, higher-leverage issuers. Figure 3 shows the constituents of a 20-name portfolio that we have randomly selected from the universe of publicly traded, high-yield companies covered by KMV.

Figure 3 The constituents of the high-yield portfolios

Names 1–5	Names 6–10	Names 11–15	Names 16–20
AES	Atlas Air Worldwide Holding Inc.	MGM Mirage	Safeway Inc.
Adaptec Inc.	Echostar Communication Corp.	Navistar International	Saks Inc.
Airgas Inc.	Gap Inc.	Nextel Communications	Service Corporation International
AK Steel Holding Inc.	Georgia-Pacific Corp.	Northwest Airlines Corp.	Solectron Corp.
Alaska Air Group Inc.	L-3 Communications Holdings Inc.	Royal Caribbean Cruises Ltd.	Sovereign Bancorp Inc.

Figure 4 Maximum likelihood estimates of DoF for high-yield portfolios

Portfolio	Asset Returns DoF	Equity Returns DoF
20-Name Portfolio	15	16
Names 1–5	15	13
Names 6–10	14	12
Names 11–15	10	10
Names 16–20	13	15

Figure 4 reports the ML estimates for the DoF of asset and equity returns for this 20-name high-yield portfolio, as well as for the four 5-name sub-portfolios. Once again, the estimated DoF for asset and equity returns are very close. When analysing correlations, a similar behaviour is also observed: specifically, the maximum absolute difference in the correlation coefficients is 6.7% and the mean absolute difference is 1.6%.

A fully-parametric *t*-model

In order to estimate the dependence structure of equity returns, one needs to estimate the whole multivariate distribution. Namely, even though the marginal distributions of returns can be estimated without any assumption on the dependence structure, the inference of the dependence structure cannot be done without estimating the marginal distributions.

Earlier, we used the empirical marginals to estimate a *t*-dependence for returns and test the null hypothesis of Gaussian dependence. In this section, we show how to estimate the copula of equity returns while modelling the marginals as shifted, scaled *t*-distributions.[7] The numerical examples we now utilise show that, since the univariate *t*-distribution generally represents a good probability model for unconditional equity returns, the estimates of the parameters of the *t*-copula are not very sensitive to the choice between the two methods.

There is plethora of evidence that the *t*-distribution accurately fits univariate equity returns (see Praetz, 1972; Blattberg and Gonedes, 1974; Glasserman *et al*, 2002). As an example, in Figure 5 we use a quantile-to-quantile (Q-Q) plot to show that a *t*-distribution with four degrees of freedom offers a much better fit for the daily equity return series of Boeing than the Normal distribution.

Figure 5 Q-Q plots of BA return vs. normal and t_4 distributions

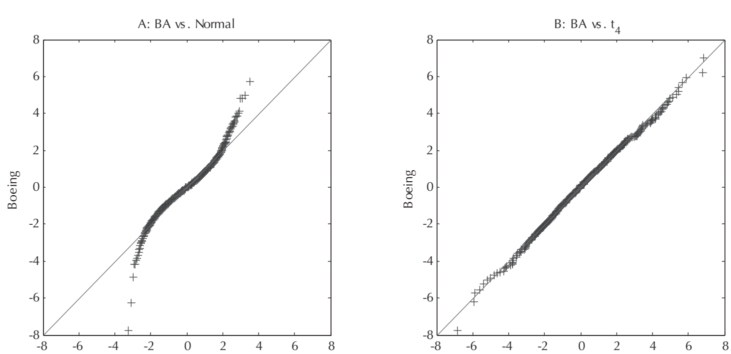

As before, we take advantage of Sklar's Theorem to split the estimation of the multivariate distribution into the following two steps:

1. estimating each marginal by itself, as the marginals are independent from each other and from the dependence structure, and
2. estimating the dependence structure based on the estimated marginals and the copula representation.

This procedure, which separates the estimation of the marginals from the joint dependence parameters, is sometimes referred to as IFM (see the monograph by Joe, 1997). To accomplish the first step, we now fit univariate *t*-distributions to the individual time-series of returns rather than using the empirical distributions as we did earlier. As for the second step, we follow exactly the same procedure described earlier.

The density of the standard *t*-distribution is given by the derivative of Equation (3):

$$f_\nu(x) = \frac{\Gamma((\nu + 1)/2)}{\Gamma((\nu/2)\sqrt{\pi\nu}} (1 + x^2 / \nu)^{-(\nu+1)/2}, \; x \in R \quad (8)$$

where ν represents the degrees of freedom (DoF) parameter. The variance of the *t*-distribution is equal to $\nu/(\nu - 2)$, $\nu > 2$. We model each equity return series by means of a shifted and scaled *t*-distribution (see Raiffa and Schlaifer, 1961, Section 7.9). Given a univariate return sample $\{X_i\}_{i=1,\ldots,n}$, where n is the number of observations, we assume that $\tilde{X}_i := (X_i - m) / \sqrt{H}$ is distributed according to a standard *t*-distribution, where m denotes the shift parameter and H denotes the scale parameter. Next we define $\theta^u = (m, H, \nu)$,

$$\Theta^U = \{\theta^U \colon m \in R, H > 0, \nu > 2\}$$

and the maximum likelihood estimator

$$\theta^U = (\hat{m}, \hat{H}, \hat{\nu}) = \underset{\theta^U \in \Theta^U}{\arg\max} \prod_{i=1}^{n} f_\nu(\tilde{X}_i)$$

which can be found by means of a simple numerical search.

The purpose of the univariate estimation is to enable the transformation of equity returns into the domain of the copula. This is achieved by computing

$$\hat{U}_i := t_{\hat{\nu}}^{-1} \left[(X_i - \hat{m}) / \sqrt{\hat{H}} \right]$$

Performing the above procedure for every return series we get

$$\{\hat{U}_{ij}\}, \; i = 1 \ldots n, \; j = 1 \ldots d$$

and we can now estimate the copula following the steps described already earlier.

Numerical examples and the effect of the sampling frequency

We will now use our two *t*-models to estimate the multivariate distribution of several 5- and 10-name baskets belonging to the Dow Jones Industrial Average. We use both daily and monthly equity return data, ranging from January 1991 to December 2001, for a total of 2,526 daily and 119 monthly observations.

Estimating the t-marginals

To implement the fully parametric methodology, we first need to estimate the *t*-marginals. Figure 6 presents the estimates of the shift, scale and DoF parameters for the 30 DJIA stocks (membership as of February 2002).

Figure 6 MLE for shift, scale and DoF parameters for DJIA stocks

Ticker	Daily			Monthly		
	$\hat{m}$	$\hat{H}$	$\hat{v}$	$\hat{m}$	$\hat{H}$	$\hat{v}$
AA	$-1.4 \cdot 10^{-4}$	4119	4.8	$1.1 \cdot 10^{-2}$	231	4.9
AXP	$6.2 \cdot 10^{-4}$	3455	5.6	$2.6 \cdot 10^{-2}$	285	4.6
T	$-1.2 \cdot 10^{-4}$	6304	3.4	$4.6 \cdot 10^{-3}$	242	4.6
BA	$1.2 \cdot 10^{-4}$	5432	4.0	$1.1 \cdot 10^{-2}$	245	8.1
CAT	$1.4 \cdot 10^{-4}$	4027	4.8	$1.2 \cdot 10^{-2}$	180	10.2
C	$1.0 \cdot 10^{-3}$	3021	6.2	$2.9 \cdot 10^{-2}$	162	7.9
KO	$4.2 \cdot 10^{-4}$	5702	5.2	$1.9 \cdot 10^{-2}$	339	4.2
DD	$1.4 \cdot 10^{-4}$	4910	5.1	$8.2 \cdot 10^{-3}$	206	>100
EK	$-1.1 \cdot 10^{-5}$	6338	4.0	$4.5 \cdot 10^{-3}$	301	5.3
XOM	$3.5 \cdot 10^{-4}$	7768	5.8	$1.0 \cdot 10^{-2}$	612	27.4
GE	$8.1 \cdot 10^{-4}$	6427	6.1	$1.8 \cdot 10^{-2}$	281	144.2
GM	$-1.3 \cdot 10^{-5}$	3409	7.5	$5.3 \cdot 10^{-3}$	149	11.1
HD	$1.1 \cdot 10^{-3}$	3658	5.2	$2.3 \cdot 10^{-2}$	182	18.3
HON	$5.1 \cdot 10^{-4}$	5118	3.9	$1.8 \cdot 10^{-2}$	312	3.2
HWP	$8.6 \cdot 10^{-4}$	2667	4.7	$2.0 \cdot 10^{-2}$	106	10.1
IBM	$1.8 \cdot 10^{-4}$	4536	4.0	$8.9 \cdot 10^{-3}$	134	31.1
INTC	$1.5 \cdot 10^{-3}$	2214	5.3	$3.2 \cdot 10^{-2}$	97	7.7
IP	$-2.1 \cdot 10^{-4}$	4534	4.7	$3.1 \cdot 10^{-3}$	229	4.7
JPM	$1.5 \cdot 10^{-4}$	5410	4.0	$1.2 \cdot 10^{-2}$	273	4.7
JNJ	$5.5 \cdot 10^{-4}$	5137	7.1	$1.4 \cdot 10^{-2}$	206	>100
MCD	$3.1 \cdot 10^{-4}$	5431	5.5	$1.3 \cdot 10^{-2}$	236	>100
MRK	$5.6 \cdot 10^{-4}$	4649	6.5	$1.6 \cdot 10^{-2}$	188	22.2
MSFT	$1.1 \cdot 10^{-3}$	2870	5.9	$2.7 \cdot 10^{-2}$	129	6.0
MMM	$1.8 \cdot 10^{-4}$	8776	3.7	$8.3 \cdot 10^{-3}$	541	4.5
MO	$5.2 \cdot 10^{-4}$	5845	3.4	$1.2 \cdot 10^{-2}$	193	6.1
PG	$4.9 \cdot 10^{-4}$	6144	4.8	$1.6 \cdot 10^{-2}$	388	4.1
SBC	$4.5 \cdot 10^{-4}$	5911	4.4	$1.0 \cdot 10^{-2}$	295	10.9
UTX	$4.7 \cdot 10^{-4}$	5974	4.4	$2.0 \cdot 10^{-2}$	336	4.6
WMT	$5.1 \cdot 10^{-4}$	3942	4.9	$1.6 \cdot 10^{-2}$	185	14.8
DIS	$2.5 \cdot 10^{-4}$	5039	4.5	$9.7 \cdot 10^{-3}$	214	9.3

The DoF estimates are quite low, confirming the well-documented non-normality of equity returns. It is interesting to note that the DoF of most names increases as we decrease the sampling frequency from daily to monthly. This might be due to aggregation (similar to the type of behaviour observed in the Central Limit Theorem, which states that sums of random variables converge to a Gaussian distribution). However, this effect varies considerably across tickers: some names show a far more "stable" additivity than others.

It is also clear that the estimated DoF differ significantly across names, independently of the frequency. This confirms that in order to estimate the copula correctly, one should allow for different DoF in the marginals rather than fitting a multivariate *t*-distribution to the data.

Estimating the t-copula

To provide numerical examples for the estimation of the *t*-dependence structure, we choose nine baskets consisting of members of the DJIA. The semi-parametric model employs the empirical marginals to map each equity return series into the unit interval, while the fully parametric method utilises the estimates described in the previous section. In both cases, the *t*-copula is then estimated using Kendall's Tau transform for the correlation matrix and maximising the likelihood function over the DoF parameter.

Figure 7 shows the estimated DoF of the *t*-copula for the nine baskets (the term "Emp-marginals" refers to the non-parametric assumption on the distribution of the marginals). First, notice that the two methodologies produce very similar results, confirming that the univariate *t*-distribution generally represents a good probability model for the univariate equity return series. Second, note that while the estimates of the marginal DoF appear to decrease systematically as we increase the frequency of the observations, the estimates of the copulas' DoF do not change significantly. In other words, while the marginal fat tails get thinner due to potential aggregation effects, there is no evidence that a similar phenomenon is driving the behaviour of the joint tails.

Our results echo the recent work of Breymann, Dias and Embrechts (2003) on the joint behaviour of FX financial series. Breymann *et al* employ a *t*-copula to model the dependence structure of a set of exchange rates, and show that the DoF parameter is almost independent of the sampling window. Moreover, this parameter is small (of order 4–6), indicating that the *t*-copula provides a more accurate description of the data than the Gaussian counterpart. Their study also compares the *t*-copula to various competing models and finds further empirical support for the former.

Summary

Our empirical investigation of the dependence structure of asset returns sheds some light on two main issues that were raised in the introduction.

Figure 7 Estimates of DoF of the *t*-copula for different DJIA baskets

Basket	Tickers	*t*-marginals		Emp-marginals	
		Daily	Monthly	Daily	Monthly
1	AA, AXP, T, BA, CAT	8	7	8	6
2	C, KO, DD, EK, XOM	10	12	9	10
3	GE, GM, HD, HON, HWP	9	7	8	6
4	IBM, INTC, IP, JPM, JNJ	8	8	7	5
5	MCD, MRK, MSFT, MMM, MO	9	6	8	5
6	PG, SBC, UTX, WMT, DIS	8	7	7	5
7	Baskets 1+2	10	10	10	8
8	Baskets 3+4	10	11	10	8
9	Baskets 5+6	9	8	9	6

DoF estimates for the *t*-copula of different DJIA baskets under a fully parametric and a semi-parametric model. The semi-parametric specification uses the empirical distributions to estimate the marginals. The estimates are provided for two different sampling frequencies (daily and monthly)

First, the assumption of Gaussian dependence between asset returns can be rejected with extremely high confidence in favour of an alternative fat-tailed dependence. Second, the dependence structures of asset and equity returns appear to be strikingly similar. The KMV algorithm that produces the asset values used in our analysis is nothing more than a sophisticated way of de-leveraging the equity to get to the value of a company's assets. Therefore, the popular conjecture that the different leverage of assets and equity will necessarily create significant differences in their joint dependence seems to be empirically unfounded, even when we analyse the value processes of low-grade issuers. Instead, our results suggest that the differences in leverage are mostly reflected in the marginal distributions of returns. From a practical point of view, these results represent good news for practitioners who only have access to equity data for the estimation of the dependence parameters of their credit models.

Our analysis also shows that imposing *t*-marginals on the equity return series, rather than using their empirical distributions, does not have significant consequences for the estimation of the dependence parameters. Thus, the *t*-copula with *t*-marginals provides a simple and parsimonious model that can be used in various financial applications in a much more straightforward manner than the semi-parametric model based on the *t*-copula. Finally, the presence of fat-tailed dependence does not seem to diminish as we decrease the sampling frequency of our data. In other words, there is no evidence that the dependence structure of equity returns approaches the Gaussian dependence structure as we increase the measurement intervals and allow for aggregation.

EXTREME EVENTS AND CORRELATED DEFAULTS

A number of different frameworks have been proposed in the literature for modelling correlated defaults and pricing multi-name credit derivatives. Hull and White (2001) generate dependent default times by diffusing correlated latent variables and calibrating default thresholds to replicate a set of given marginal default probabilities. Multi-period extensions of the one-period CreditMetrics paradigm are also commonly used, even though they produce the undesirable serial independence of the realised default rate. A computationally more expensive approach is that based on the implementation of stochastic intensity models, as proposed by Duffie and Singleton (1998) and Duffie and Garleanu (1998). Finger (2000) offers an excellent comparison of several multivariate models in terms of the default distributions they generate over time when calibrated to the same marginals and first-period joint default probabilities. While most multi-name models require simulation, the need for accurate and fast computation of greeks has pushed researchers to look for modelling alternatives. Finger (1999) and, more recently, Gregory and Lambert (2003) show how to exploit a low-dimensional factor structure and conditional independence to obtain semi-analytical solutions.

In an influential paper, Li (2000) presents a simple and computationally inexpensive algorithm for simulating correlated defaults. His methodology builds on the implicit assumption that the multivariate distribution of default times and the multivariate distribution of asset returns share the same copula, which he assumes to be Gaussian. The approach described in the previous chapter is based on this model. The results of our earlier statistical analysis suggest that the dependence structure of asset returns is very similar to the dependence structure of equity returns, and that the dependence structure of equity returns is better described by a *t*-dependence than by a Gaussian copula.

Therefore, it seems natural to modify Li's methodology to account for the likelihood of extreme joint realisations, and simulate correlated default times under the assumption that their dependence structure is the same as that of the associated equity returns.

Simulating default times with fat-tailed dependence

To construct the multivariate distribution of default times under the objective probability measure, we first need to estimate the marginal distributions, which we will denote with F_1, F_2, ..., F_d. These can be derived from univariate structural models (such as KMV's EDF) or simply estimated using observed default frequencies within relatively homogeneous peer groups (such as Moody's default frequencies by rating). We then join these marginals with a *t*-copula, and estimate the dependence parameters (correlation matrix Σ and DoF ν) from equity returns using either one of the two procedures described earlier.

For valuation purposes, we need to construct the multivariate distribution of default times under the risk-neutral probability measure. In this case, it is common practice to back out the marginals $F_1, F_2, \ldots, F_d$ from single-name defaultable instruments (such as credit default swaps). Given the low liquidity of multi-name instruments, it is not yet possible to use their market prices to obtain implied values for the dependence parameters. Instead, practitioners generally estimate the copula using historical data, implicitly relying on the extra assumption that the dependence structure of default times remains unchanged when we move from the objective to the pricing probability measure.

Simulating default times from this multivariate distribution is straightforward:

1. Commence by simulating a multivariate-t random vector $X \in R^d$ with correlation Σ and ν DoF.
2. Next, transform the vector into the unit hyper-cube using $U = (t_\nu(x_1), t_\nu(x_2), \ldots, t_\nu(x_d))$
3. Translate U into the corresponding default times vector τ using the inverse of the marginal distributions: $\tau = (\tau_1, \tau_2, \ldots, \tau_d) = (F_1^{-1}(u_1), F_2^{-1}(u_2), \ldots, F_d^{-1}(u_d))$.

The simulation algorithm is illustrated in Figure 8.

It is easy to verify that τ has the given marginals and that its dependence structure is given by a t-copula with correlation Σ and ν DoF. The logic of the proof can be established via the following.

- X has a $t_{\Sigma,\nu}$ copula by definition.
- Since copulas are preserved by strictly monotonic transformations of the variables and the univariate t-distribution is strictly increasing, U also possesses a $t_{\Sigma,\nu}$ copula.

Figure 8 Mapping a student-t random variable to a default time

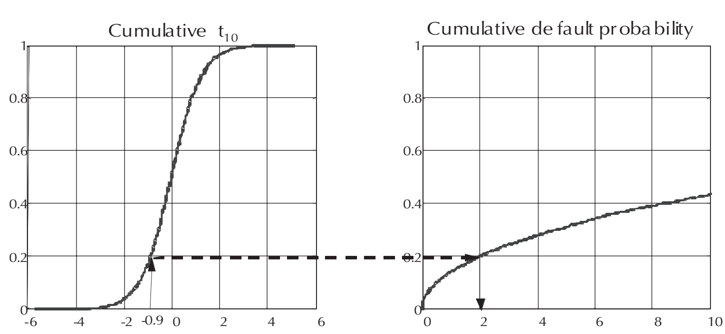

- ❏ Since copulas are preserved by strictly monotonic transformations of the variables, while the marginal distributions of default times are strictly monotonic, τ also has a $t_{\Sigma,\nu}$ copula.
- ❏ Since $t_{\Sigma,\nu}$ has marginals, U has uniform marginals.
- ❏ Since U has uniform marginals, τ has F_1, F_2, ..., F_d marginals.

Asset correlation, default time correlation and default event correlation

The simulation algorithm discussed above is based on the assumption that asset returns and default times share the same copula, and consequently share the same correlation matrix. To understand better the impact of different dependence assumptions on the valuation and risk measures of default-contingent instruments, it is useful to introduce the concept of "default event correlation".

Default event correlation measures the tendency of two credits to default jointly within a specified horizon. Formally, it is defined as the correlation between two binary random variables that indicate defaults, ie,

$$\rho_D = \frac{p_{AB} - p_A p_B}{\sqrt{p_A(1 - p_A)}\sqrt{p_B(1 - p_B)}} \tag{9}$$

where p_A and p_B are the marginal default probabilities for credits A and B, and p_{AB} is the joint default probability. Of course, p_A, p_B and p_{AB} all refer to a specific horizon. Notice that default event correlation increases linearly with the joint probability of default, and is equal to zero if – and only if – the two default events are independent.

Default event correlations are the fundamental drivers in the valuation of multi-name credit derivatives. Unfortunately, the scarcity of default data makes joint default probabilities, and therefore default event correlations, very hard to estimate directly. As a result, researchers rely on alternative methods to calibrate the frequency of joint defaults within their models: the method we described above solves this problem by assuming that rarely observable default times and frequently observable equity returns share the same copula.

It is interesting to see how the DoF parameter – which regulates tail dependence and the likelihood of joint extreme events – influences default event correlations. Using a 5-year horizon and two credits whose default times are exponentially distributed with hazard rates of 1%, Figure 9 compares a normal copula and a t-copula with three and 10 degrees of freedom. Tail dependence increases default event correlation for any value of asset correlation. In particular, notice that even when asset returns are uncorrelated (ie, they are linearly independent), tail dependence can produce a significant amount of default event correlation.

Figure 9 Asset correlation vs. default event correlation

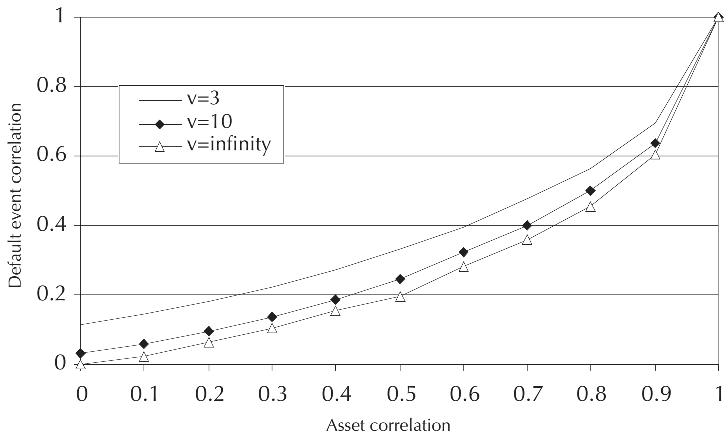

MULTI-NAME CREDIT DERIVATIVES

In the previous section, we proposed an algorithm to simulate tail-dependent defaults under either the objective or the pricing measures. Fair values and risk metrics for any multi-name credit derivative can now be easily computed. In this section, we analyse the consequences of making different assumptions with respect to the dependence structure of default times. Specifically, we focus on two popular multi-name credit instruments, namely *n*th-to-default baskets and portfolio loss tranches.

*n*th-to-default baskets

In an *n*th-to-default basket swap, two counterparties agree on a maturity and a set of reference assets, and enter into a contract whereby the protection seller periodically receives a premium (also called "basket spread") from the protection buyer. In exchange, the protection seller stands ready to pay the protection buyer par minus recovery of the nth referenced defaulter in the event that the *n*th default occurs before the agreed-upon maturity. First- and second-to-default swaps are the most popular orders of protection.

Taking extreme events into account has significant consequences on the estimate of the Expected Discounted Loss (EDL) of a basket swap (see the previous chapter for the definition of the EDL). This is because, other things equal, simulating defaults by means of a fat-tailed copula increases the probability of joint defaults, and thus default event correlation (see Equation 9). We focus on EDL because this measure relates both to agency rating methodologies (when computed under real-world probabilities)

Figure 10 *n*th-to-default basket swap

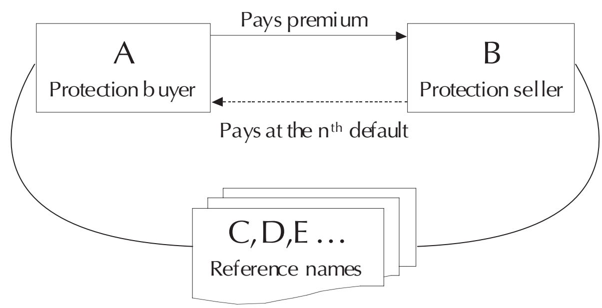

and to the fair compensation for the default risk exposure (when computed under risk-neutral probabilities).

The sign of the relation between EDL and default event correlations, however, depends on the order of the basket. The EDL of a first-to-default exposure is always monotonically decreasing in default event correlations. In terms of valuation, this means that allowing for joint extreme events makes first-to-default protection unambiguously cheaper.

The EDL of a second-to-default exposure is not necessarily monotonic in default event correlations. Rather, it generally increases up to a maximum, then it starts decreasing. The location of the turning point depends on all other parameters and, in particular, on the number of names in the basket. With a low number of names, the EDL of a second-to-default exposure is generally increasing in default event correlations over most of the domain. Intuitively, with only a handful of names in the portfolio, the event that at least two of them default becomes more likely as we increase their tendency to default together.

These qualitative relations are consistent with the results reported in Figure 11, where we compare the EDL of five-year first-, second-, and third-to-default exposure on a 5-name basket using both a Normal copula and a *t*-copula with 12 DoF. In both cases, the marginal distributions of default times are assumed to be defined by a constant yearly hazard rate equal to 1%, recovery rates are known and equal to 40%, and the risk-free discounting curve is flat at 2%. We compute EDL for three different levels of asset correlations, namely 0%, 20%, and 50%. The standard errors of the Monte Carlo estimators are reported in brackets and represent the simulation standard error for the EDL as a percentage of the EDL. As one would expect, the percentage difference between the Gaussian copula and the *t*-dependence is higher when the triggering event is less likely.

Figure 11 Expected discounted loss of *n*th-to-default baskets

Asset Correlation	Copula	1st-TD	2nd-TD	3rd-TD
0%	Normal	0.1265 (0.059%)	0.0121 (0.206%)	0.0006 (0.898%)
	t12	0.1207 (0.059%)	0.0167 (0.169%)	0.0017 (0.566%)
	% Difference	–5%	38%	185%
20%	Normal	0.1151 (0.066%)	0.0205 (0.155%)	0.0033 (0.380%)
	t12	0.1094 (0.063%)	0.0239 (0.137%)	0.0051 (0.155%)
	% Difference	–5%	17%	55%
50%	Normal	0.0934 (0.072%)	0.0305 (0.131%)	0.011 (0.231%)
	t12	0.0888 (0.071%)	0.0318 (0.129%)	0.0127 (0.197%)
	% Difference	–5%	4%	15%

Normal vs. Student *t*-copula with DoF=12, 10M- path Monte Carlo simulation, standard errors in parenthesis

Portfolio loss tranches

According to a recent survey published in Risk Magazine (February 2003), portfolio loss tranches have become one of the most common types of multi-name credit exposures traded in the market.[8] In a typical portfolio loss tranche, a protection buyer pays a periodic premium to a protection seller, who, in exchange, stands ready to compensate the buyer for a pre-specified slice (tranche) of the losses affecting a set of reference obligations. The reference portfolio is generally composed of dozens (and sometimes hundreds) of credits, and each name is represented in the portfolio according to a given notional amount.

Here we consider a portfolio of 100 names, each with US$1 million notional. A tranche exposure is defined by a lower and an upper percentile of the total notional. For example, the seller of protection on the 5%–10% tranche of our 100-name portfolio will be responsible for covering losses exceeding US$5 million and up to US$10 million (US$5 million exposure). Losses are defined as the notional amount of defaulted credits times the associated loss given default (LGD). In our example, we assume uniform recovery rates of 35%, ie, 65% LGD for every credit in the reference portfolio.

We first consider a valuation exercise using the following parameters:

- ❏ 1% risk-neutral hazard rate for each reference name;
- ❏ 20% asset correlation between every pair of credits;
- ❏ flat risk-free curve at 2%; and
- ❏ a 5-year maturity deal.

Figure 12 compares the risk-neutral EDL for several tranches under the two alternative assumptions of Gaussian dependence and *t*-dependence with 12 DoF. The results show the significant impact that the (empirically motivated) consideration of tail dependence has on the distribution of losses across the capital structure: expected losses are clearly redistributed from the junior to the senior tranches, as a consequence of the increased default event correlations. This implies that the Gaussian assumption underestimates the fair compensation for senior exposures and overestimates the fair compensation for junior risk.

Even larger differences can be observed if one compares higher moments or tail measures of the tranches' loss distributions. Let us now assume that each of the 100 reference names has an objective default

Figure 12 Expected discounted loss of portfolio tranches

Tranche	Normal Copula EDL (Std Err)	*t*-Copula DoF=12 EDL (Std Err)	% Difference
0% – 5%	2,256,300 (0.14%)	2,012,200 (0.23%)	–11%
5% – 10%	533,020 (0.63%)	601,630 (0.66%)	13%
10% – 15%	146,160 (1.37%)	221,120 (1.06%)	51%
15% – 20%	41,645 (1.70%)	90,231 (1.62%)	117%
20% – 100%	16,188 (4.94%)	59,042 (2.79%)	265%

Normal vs. Student *t*-copula with DoF=12, 100K-path Monte Carlo simulation, standard errors in parenthesis

Figure 13 Value-at-risk and expected shortfall at the 95th percentile

Tranche	Copula	$VaR_{95\%}$	$ES_{95\%}$
0% – 5%	Normal	5,000,000	5,000,000
	t12	5,000,000	5,000,000
	% Difference	0%	0%
5% – 10%	Normal	850,000	3,119,812
	t12	2,150,000	4,278,209
	% Difference	153%	37%
10% – 15%	Normal	0	600,480
	t12	0	1,583,187
	% Difference	0%	164%
15% – 20%	Normal	0	124,750
	t12	0	584,986
	% Difference	0%	369%
20% – 100%	Normal	0	32,747
	t12	0	339,124
	% Difference	0%	936%

Normal vs. Student *t*-copula with DoF=12, 100K-path Monte Carlo simulation intensity equal to 0.5%. The remaining parameters are unchanged. Figure 13 compares the two dependence assumptions in terms of the 95% Value-at-Risk (VAR) and expected shortfall that they produce for a number of loss tranches.

CONCLUSIONS

The empirical study presented in the first part of this chapter has two main findings. First, empirical evidence suggests that large joint movements of equity values occur with higher likelihood than what is predicted by correlation-based models. In particular, empirical evidence seems to favour a fat-tailed dependence structure instead of the widely-used Gaussian one. The second observation is that the dependence structures of asset and *equity* returns appear to be strikingly similar.

One interesting corollary to these empirical findings is the lack of support for the popular conjecture regarding the ways in which the different leverage of assets and equity create significant differences in their joint dependence. To this end, the KMV algorithm that we use to "back-out" asset values from observed equity data can be viewed simply as a means for de-leveraging the equity to arrive at the value of a company's assets. Our results suggest that the differences in leverage are mostly reflected in the marginal distributions of returns. A practical consequence is that (observed) equity data seem to provide a valid and consistent proxy for (unobserved) asset returns, at least for the purpose of calibrating the dependence structure.

We conclude our empirical study with a recommendation for a simple and parsimonious framework that can be used to model asset dependencies. (Such a model will also drive the analysis of multi-name credit derivatives.) Specifically, this is a fat-tailed multivariate distribution, having marginals that each follow a univariate t-distribution (with possibly different parameters), and with a dependence structure given by a t-copula. It is important to note that while the estimates of the parameters of the marginals (in particular, the degrees-of-freedom which dictates the fatness of the marginal tail) may depend on the sampling frequency of the data, the tail behaviour in the dependence structure seems to be quite insensitive in this regard. In particular, there is no empirical evidence that the dependence structure of equity returns approaches the Gaussian dependence structure as we increase the measurement intervals and allow for aggregation.

These empirical findings have significant bearing on multi-name credit derivatives models. In particular, the results described above indicate that asset returns exhibit non-negligible tail dependence; therefore, if one follows the "structural approach", default times seem to be more accurately modelled using a t-dependence structure rather than the widely used Gaussian one. The examples on the valuation of first- and second-to-default baskets illustrate the importance of the modelling choice for pricing purposes. In addition, the example that considers synthetic loss tranches suggests that multivariate Gaussian models will generally underestimate default correlations and thus overestimate the expected loss of junior positions and underestimate the expected loss of mezzanine and senior tranches.

1 We would like to thank Mark Broadie, Paul Glasserman, Mark Howard, Prafulla Nabar, Dominic O'Kane, Lutz Schloegl and Stuart Turnbull for comments and suggestions on earlier versions of this document. The usual disclaimer applies.

2 A description of these models can be found in Kealhofer and Bohn (2001) and Gupton, Finger and Bhatia (1997).

3 See, for example, Mashal and Naldi (2002), Mashal and Zeevi (2002) and Breymann *et al* (2003).

4 This approach follows the semi-parametric estimation framework developed in a more abstract context by Genest *et al* (1995).

5 A rigorous derivation and an explicit characterization of γ is given in Appendix A of Mashal and Zeevi (2002) who also validate this asymptotic numerically.

6 The range of accepted DoF is very narrow in each case, exhibiting similar behaviour to that displayed in Figure 1.

7 This should not be confused with a multivariate-t model, since we are not restricting the copula and all of the marginals in order to achieve the same DoF parameter.

8 See "Credit derivatives survey, flow business booms", by Navroz Patel, Risk Magazine, February 2003, pp. S20-S23.

BIBLIOGRAPHY

Blattberg, R. C. and N. J. Gonedes, 1974, "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices", *Journal of Business*, 47, pp. 244–80.

Breymann, W., A. Dias and P. Embrechts, 2003, "Dependence Structures for Multivariate High-Frequency Data in Finance", Working Paper, Department of Mathematics, ETH Zurich.

Duffie, D. and N. Garleanu, 2001, "Risk and Valuation of Collateralized Debt Obligations", *Financial Analyst Journal*, 57, pp. 41–59.

Duffie, D. and K. Singleton, 1998, "Simulating Correlated Defaults", Working Paper, Stanford University, Stanford.

Embrechts, P., F. Lindskog and A. J. McNeal, 2001, "Modelling Dependence with Copulas and Applications to Risk Management", Working Paper, Risklab, ETH, Zurich.

Finger, C., 1999, "Conditional Approaches for CreditMetrics Portfolio Distributions", RiskMetrics Group, New York.

Finger, C., 2000, "A Comparison of Stochastic Default Rate Models", RiskMetrics Group, New York.

Fitch Ratings, 2003, "Default Correlation and its Effect on Portfolios of Credit Risk", Structured Finance, Credit Products Special Report, February.

Genest, C., K. Ghoudi and L-P. Rivest, 1995, "A Semiparametric Estimation Procedure of Dependence Parameters in Multivariate Families of Distributions", Biometrika, 82, pp. 543–52.

Glasserman P., P. Heidelberger and P. Shahabuddin, 2002, "Portfolio Value-at-Risk with Heavy-Tailed Risk Factors", *Mathematical Finance*, 12, pp.239–70.

Gregory, J. and J.P. Laurent, 2003, "I Will Survive", *Risk*, June, pp. 103–107.

Gupton, G.M., C.C. Finger and M. Bhatia, 1997, "CreditMetrics – Technical Document", J.P. Morgan, New York.

Hull, J. and A. White, 2001, "Valuing Credit Default Swaps II: Modelling Default Correlations", *Journal of Derivatives*, 8, pp. 12–22.

Joe, H., 1997, *Multivariate Models and Dependence Concepts*, "Monographs on Statistics and Applied Probability", 37, Chapman and Hall, London.

Kealhofer, S. and J.R. Bohn, 2001, "Portfolio Management of Default Risk", KMV, San Francisco.

Li, D., 2000, "On Default Correlation: a Copula Function Approach", *Journal of Fixed Income*, 9, pp. 43–54.

Lindskog, F., 2000, "Linear Correlation Estimation", Working Paper, RiskLab, ETH Zurich.

Mashal, R. and M. Naldi, 2002, "Extreme Events and Default Baskets" *Risk*, June, pp.119–122.

Mashal, R. and A. Zeevi, 2002, "Beyond Correlation: Extreme Co-Movements between Financial Assets", Working Paper, Columbia University, New York.

Mashal, R. and A. Zeevi, 2003, "Inferring the Dependence Structure of Financial Assets: Empirical Evidence and Implications", Working Paper, Columbia University, New York.

Merton, R.C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates", *Journal of Finance*, 2, pp. 449–70.

Praetz, P.D., 1972, "The Distribution of Share Price Changes", *Journal of Business*, 45, 49–55.

Raiffa, H. and R. Schlaifer, 1961, *Applied Statistics Decision Theory*, Harvard University Press, Boston.