Comparing the Dependence Structure of Equity and Asset Returns[1]

Roy Mashal
212-526-8883
rmashal@lehman.com

Marco Naldi
212-526-1728
mnaldi@lehman.com

Assaf Zeevi
212-854-9678
assaf@gsb.columbia.edu

The dependence structure of asset returns lies at the heart of a class of models that is widely employed for the valuation of multi-name credit derivatives. Using a statistical methodology that relies on a minimal amount of distributional assumptions, we investigate whether the popular tenet of Normal dependence between asset returns is empirically motivated. 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 provide strong evidence against the assumption of Normal dependence, and support the use of equity returns as proxies for asset returns. A simple application to portfolio loss tranches shows that these findings have significant consequences on the pricing and rating of multi-name credit instruments.

1. INTRODUCTION

The valuation of default-contingent instruments calls for the modeling 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 modeled as an unexpected arrival.

Currently, a major challenge facing credit models is represented by the rapid growth of multi-name instruments, whose valuation entails modeling the joint default behavior 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 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 this paper, we focus on the empirical properties of the dependence structure – also known as the *copula function* – of asset returns.

Several well-known multivariate models assume a joint normal distribution for asset returns. Hull and White (2001), for example, 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[2], rely on the joint normality of the default-triggering variables. The widespread use of the multivariate Normal distribution is certainly related to the simplicity of its dependence structure, which is fully characterized by a correlation matrix. It remains to be seen, however, whether this assumption is supported by empirical evidence.

[1] *We would like to thank Dominic O'Kane, Stuart Turnbull and two anonymous referees for comments and suggestions.*

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

A number of recent studies have shown that the joint behavior of *equity* returns is better described by a "fat-tailed" *t*-copula than by a Normal copula, and that correlations are therefore not sufficient to appropriately characterize their dependence structure.[3] The first goal of this article is to apply the same kind of analysis to *asset* returns, and test the null hypothesis of Gaussian dependence versus the alternative of "joint fat tails."

Given the low liquidity of multi-name instruments, it is not yet possible to use their market prices to back-out implied values for the dependence parameters. Instead, practitioners generally estimate the copula of asset returns from historical data. From a valuation perspective, this amounts to the assumption that the dependence structure of asset returns remains unchanged when we move from the objective probability measure to the pricing (risk-neutral) distribution. Rosenberg (2001) identifies general conditions under which this equivalence holds.

Even if we are willing to rely on this invariance, we still face a major obstacle when attempting to estimate the dependence structure from historical data: asset returns are not directly observable. In fact, the use of unobservable underlying processes is one of several criticisms that the structural approach has received over the years. Given the lack of observable asset returns, it has become customary to proxy the asset dependence with equity dependence, and to estimate the parameters governing the joint behavior 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 parameterized by equity correlations.

The use of equity returns to infer the joint behavior of asset returns is often criticized on the grounds of the different leverage of assets and equity. Even those who accept it as a valid approximation for high-grade issuers, often criticize this approach when it is 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 capitalization. On the other hand, high-yield borrowers are closer to the default threshold, and a variation in their asset value can potentially produce a significant variation in the market value of their debt as well. This "leverage effect" may generate significant differences in the joint dynamics of equity and asset values. The second goal of this article is to shed some light on the magnitude of the error induced by using equity data as a proxy for actual asset returns.

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, – i.e., recognizing 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 and 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.[4] In this article, we use the asset value series generated by KMV's model to study the dependence properties of asset returns.

[3] *See, for example, Mashal and Naldi (2002) and Mashal and Zeevi (2002).*

[4] *Copyright © 2000-2002 KMV LLC. All rights reserved. KMV and the KMV logo are registered trademarks of KMV LLC. CreditEdge and EDF are trademarks of KMV LLC.*

In Section 2, we 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. This section also describes a related test statistic that can be used to evaluate the statistical relevance of our point estimates. Section 3 presents the results of our empirical investigation, while Section 4 applies our findings to the practical issue of measuring and pricing the risk of a popular multi-name credit product. Finally, Section 5 comments on our results and offers some concluding remarks.

2. METHODOLOGY

This section describes an estimation procedure that is used to calibrate a certain class of dependence structures to the equity and asset returns data. We first digress briefly to discuss some dependence-related concepts, and then proceed to describe the estimation methodology and an associated testing procedure.

The key ingredient in modeling 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. Copulas have many important characteristics that make them a central concept in the study of joint dependencies, see, e.g., the recent survey paper by Embrechts *et al.* (2001).

A particular copula that plays a crucial role in our study is given by the dependence structure underlying the multivariate Student t distribution. While the Gaussian distribution lies at the heart of most financial models and builds on the concept of correlation, the Student t retains the notion of correlation but adds an extra parameter into the mix, namely, the *degrees-of-freedom* (DoF). The latter plays a crucial role in modeling and explaining extreme co-movements in the underlyings.

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 would target the dependence structure rather than the distributions themselves, thus eliminating the effect of marginal returns that would "contaminate" the estimation problem in the latter case. To summarize, the t-dependence structure constitutes an important and quite plausible generalization of the Gaussian modeling paradigm, which is our main motivation for focusing on it in this study.

With this in mind, 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,…,d, and having a joint distribution H with underlying t-dependence structure, which is denoted by

$$C(\cdot;\nu,\Sigma).$$

Here, Σ denotes the correlation matrix and *v* the DoF parameter.[5] Suppose we have *n* observations $\{X_i\}_{i=1}^n$ on these *d* names, where the returns $X_i = (X_{i1},\dots,X_{id})$ are assumed to be mutually independent and distributed according to *H*.

If the marginal distributions were known, then we could use the representation (1) in the Appendix to conclude that

$$U := F(X_i) \sim C(\,\cdot\,;\nu,\Sigma),$$

where $F(X_i) := (F_1(X_{i1}),\dots,F_d(X_{id}))$ 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,

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

where $I\{\,\cdot\,\}$ is the indicator function, i.e.,

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

We then work with the *pseudo-sample* observations [6]

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

Focusing on the *t*-dependence structure $C(\,\cdot\,;\nu,\Sigma)$ [formally given by the *t*-copula (2) in the Appendix], let us denote by

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

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* (see the Appendix). Now, let

$$\hat{\theta} := (\hat{\nu},\hat{\Sigma})$$

5 For details on the relation between the joint distribution H and the copula C, see a version of Sklar's Theorem in the Appendix.

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

denote the *maximum likelihood* (ML) estimator of the DoF and correlations, i.e., the value of $\theta \in \Theta$ maximizing $L_n(\theta)$. The results reported in the next section refer to estimates obtained in this manner.

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, in the sequel we 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

$$H_0 : \theta \in \Theta_0 \quad \text{vs.} \quad H_1 : \theta \in \Theta ,$$

where

$$\Theta_0 = \left\{ \theta \in \Theta : \nu = \nu_0 \right\} \subset \Theta .$$

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

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

To determine the adequacy of each value of ν_0, we need to characterize 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

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

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).[7] Thus, we can calculate approximate *p*-values as a function of ν_0 as follows

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

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 (respectively, rejected) by the observed return data. Notice that *large* values of the test statistic correspond to *small p*-values, thus indicating that the hypothesized ν is not plausible on the basis of the empirical observations. This hypothesis testing formulation illustrates the

[7] *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.*

"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.

3. EMPIRICAL EVIDENCE

In this section, we 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 12/31/00 to 11/8/02.

In the following sub-sections, we focus our attention on two portfolios, the 30-name Dow Jones Industrial Average and a 20-name high-yield portfolio.

3.1 DJIA Portfolio

Following the semi-parametric methodology described in Section 2, we estimate the number of degrees-of-freedom (DoF) of a *t*-copula without imposing any structure on the marginal distributions. Using the test statistic introduced earlier, Figure 1 presents a sensitivity analysis for various null hypotheses of the underlying tail dependence, 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 (v =12) for asset returns and at 13 DoF (v =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 (v =∞) can be rejected with an infinitesimal probability of error.

Finally, notice that the point estimates 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 behavior of a 12- and a 13-DoF *t*-copula is negligible in terms of any practical application.

Figure 1. DJIA Portfolio: Asset and Equity Returns Test Statistics as Functions of Null Hypothesis for DoF

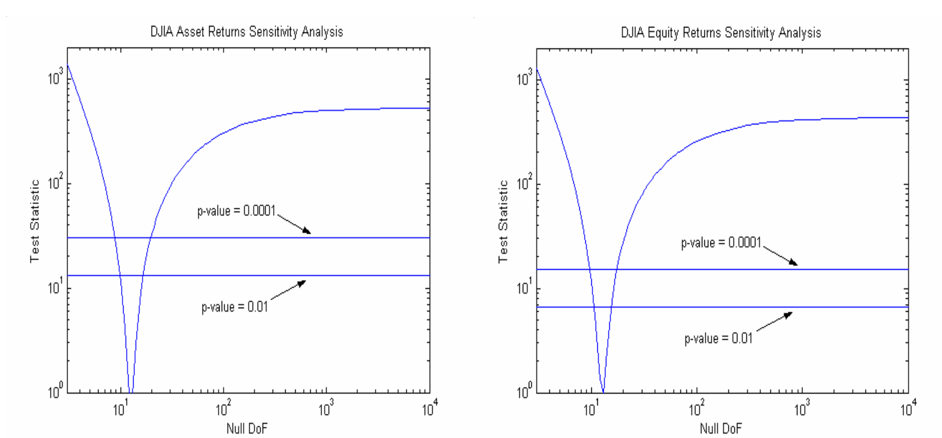

Figure 2 reports the point estimates of the DoF for asset and equity returns in the DJIA basket, as well as for three subsets consisting of the first, middle, and last 10 names (in alphabetical order). The similarities between the joint tail dependence (as measured by the DoF) of asset and equity returns are quite striking. [8]

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

Next, we compare the remaining parameters that define a *t*-copula, i.e., the correlation coefficients. Using a robust estimator based on Kendall's rank statistic[9], we compute the two 30x30 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.

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

3.2 High-Yield Portfolio

In this section we 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. High-Yield Portfolio

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.

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

[9] *See Lindskog (2000).*

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 shown in Figure 3. Once again, the estimated DoF for asset and equity returns are very close. When analyzing correlations, a similar behavior is also observed, specifically, the maximum absolute difference in the correlation coefficients is 6.7% and the mean absolute difference is 1.6%.

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

4. APPLICATION: PORTFOLIO LOSS TRANCHE

In this section, we analyze the potential consequences of making different assumptions with respect to the dependence structure of asset returns. To illustrate this, we compute the expected discounted losses (EDL) of portfolio loss tranches. We focus on EDL because this measure relates both to the agency rating (when computed under real-world probabilities) and to the fair compensation for the credit exposure (when computed under risk-neutral probabilities).

According to a recent survey published in Risk Magazine (February 2003 issue), portfolio loss tranches have become one of the most common type of multi-name credit exposures traded in the market. 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 $1MM 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 $5MM and up to $10MM ($5MM 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%, i.e., 65% LGD for every credit in the reference portfolio. The additional parameters are:

1. 1% yearly hazard rate for each reference name,
2. 20% asset correlation between every pair of credits,
3. flat risk-free curve at 2%,
4. 5-year maturity deal.

Figure 5 compares the expected discounted losses 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 volatility of the overall portfolio loss distribution. Of course, even larger differences can be observed if one compares higher moments or tail measures of the tranches' loss distributions.

Figure 5. Expected Discounted Loss (EDL), 100K-path Monte Carlo Simulation, Standard Errors in Parenthesis

Tranche	Normal Copula EDL (Std Err)	*t* Copula DoF=12 EDL (Std Err)	Pctg 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%

5. DISCUSSION AND CONCLUDING REMARKS

We now turn to some comments on the methodology and results described in this paper. We start by noting that the empirical evidence indicating the similarity between equity and asset return dependencies should be considered relative to the time scales of the financial data. This applies also to the evidence rejecting the Gaussian dependence structure in favor of the *t*-copula. For example, a common empirical observation is that data sampled at relatively low frequencies tends to exhibit "lighter tails" in the *marginals*, thus, closer to a Gaussian distribution. However, the recent work of Breymann, Dias and Embrechts (2003) that studies the joint behavior of FX financial series indicates that the DoF parameter in the *t*-copula (reflecting extreme tail dependence) is almost independent of the sampling window. Moreover, this parameter is small (of order 4-6) indicating that the *t*-copula provides a much 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.

The second comment concerns an intrinsic drawback of the copula paradigm, in particular, the fact that this approach only provides a *static* model. To this end, a *conditional copula* model may be more appropriate from a modeling standpoint, but this results in far more complicated specification and inference issues that restrict applicability. Given that we consider credit instruments with a fixed and given maturity over a long time horizon, we believe that this issue essentially does not restrict our analysis. (In general one can certainly consider dependence structures that differ according to the associated maturity date.)

To summarize, our empirical investigation of the dependence structure of asset returns sheds some light on the two main issues that were raised in the introduction. First, the assumption of Gaussian dependence between asset returns can be rejected with extremely high confidence in favor of an alternative "fat-tailed dependence." Multivariate structural models that employ correlated Gaussian processes for the diffusion of asset values will generally underestimate default correlations, and thus undervalue the junior risk and overvalue the mezzanine and senior risk of multi-name credit products. Fat-tailed increments of the joint value processes will produce more accurate joint default scenarios and more accurate valuations. 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 else 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 dynamics seems to be empirically unfounded, even when we analyze 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.

APPENDIX

Copula Functions and the *t*-Dependence Structure

A *copula function* is a multivariate distribution defined over $[0,1]^d$, with uniform marginals. The importance of the copula stems from the fact that it captures the dependence structure of a multivariate distribution. This can be seen more formally from the following fundamental result, known as Sklar's theorem, 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

$$H(x_1,\ldots,x_d) = C(F_1(x_1),\ldots,F_d(x_d)). \quad (1)$$

As indicated in the main body, this study focuses on a natural generalization 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,

$$t_\nu(x) = \int_{-\infty}^{x} \frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2)(\nu\pi)^{1/2}} (1 + y^2/\nu)^{-(\nu+1)/2}\, dy .$$

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

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

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

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

The density of the *t*-copula, $c(\cdot;\nu,\Sigma)$, is obtained by differentiating the *t*-copula w.r.t. $u_1,\ldots,u_d$ [for details, see, e.g., Section 3.1 in Mashal and Zeevi (2002)].

REFERENCES

- Breymann, W., A. Dias, and P. Embrechts, 2003, "Dependence Structures for Multivariate High-Frequency Data in Finance", Working Paper, Department of Mathematics, ETH Zurich.
- Embrechts, P., F. Lindskog, and A. J. McNeal, 2001, "Modeling Dependence with Copulas and Applications to Risk Management," Working paper, Risklab, ETH, Zurich.
- 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: 543–552.
- 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: Modeling Default Correlations," *Journal of Derivatives*, 8: 12-22.
- Kealhofer, S. and J.R. Bohn, 2001, "Portfolio Management of Default Risk," KMV, San Francisco.
- Lindskog, F., 2000, "Linear Correlation Estimation," Working paper, RiskLab, ETH Zurich.
- Mashal, R. and M. Naldi, 2002, "Generalizing Asset Dependency: Applications to Credit Derivatives," Working Paper, Lehman Brothers, New York.
- Mashal, R. and A. Zeevi, 2002, "Beyond Correlation: Extreme Co-Movements between Financial Assets," 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: 449:470.
- Rosenberg, J.V., 2001, "Non-Parametric Pricing of Multivariate Contingent Claims," Working Paper, Stern School of Business, New York University, New York.