AN IMPROVED IMPLIED COPULA MODEL AND ITS APPLICATION TO THE VALUATION OF BESPOKE CDO TRANCHES

John Hull and Alan White

Joseph L. Rotman School of Joseph L. Rotman School of Management
University of Toronto

First Draft: April 2008
This draft: June 2008

ABSTRACT

In Hull and White (2006) we showed how CDO quotes can be used to imply a probability distribution for the hazard rate over the life of the CDO. This is known as the "implied copula approach." In this paper we develop a parametric version of the implied copula approach and show how it can be used for valuing bespoke CDOs. Both homogeneous and heterogeneous versions of the model are presented and the differences between the results obtained from the two versions of the model are examined. Results are also presented for the situation where the model is extended so that hazard rates are driven by more than one factor.

AN IMPROVED IMPLIED COPULA MODEL AND ITS APPLICATION TO THE VALUATION OF BESPOKE CDO TRANCHES

John Hull and Alan White

A important activity for derivatives traders is using the market prices of actively traded instruments to estimate prices for similar less actively traded instruments. In most derivatives markets they have developed ways of doing this that are not heavily model dependent. Consider, for example, the problem of valuing an option whose strike price and time to maturity are different from those options for which market prices are available. Traders use available options in conjunction with Black-Scholes to identify points on the volatility surface. Interpolation (and when necessary extrapolation) procedures are then employed to estimate an implied volatility for the option of interest. This volatility is substituted into Black-Scholes to provide a price for the option. A model is used, but its main role is to facilitate the interpolation between market prices. The results obtained using the Black-Scholes model are similar to those that would be obtained using another model of stock price behavior.

In the case of portfolio credit derivatives, the actively traded instruments are the tranches of standard portfolios such as iTraxx Europe or CDX NA IG. Traders need to use the prices of these instruments to estimate the prices of tranches of nonstandard portfolios in the synthetic collateralized debt obligation (CDO) market. This is a similar type of problem to the one just mentioned faced by options traders, but much more complicated. This paper suggests one way of proceeding.

The model proposed is a parametric version of the implied copula approach in Hull and White (2006). A key advantage of the parametric version is that extending the model from the homogeneous to the heterogeneous case, or from one factor to two factors, is straightforward. The paper compares the model with some of the base correlation mapping procedures that have been proposed. It provides results on the difference between the heterogeneous and homogeneous version of the model, and on the effect of assuming that hazard rates are driven by more than one factor.

I. THE SYNTHETIC CDO MARKET

The most popular portfolio credit derivative is a collateralized debt obligation (CDO). In this a portfolio of obligors is defined and a number of tranches are specified. Each tranche is responsible for losses between U_1 % and U_2% of the total principal for some U_1 and U_2. As the market has developed, standard portfolios and standard tranches have been specified to facilitate trading. One example is the CDX NA IG portfolio. This is an equally weighted portfolio of 125 investment grade North American companies with the notional principal (size of the credit exposure) being the same for each company. The equity tranche is responsible for default losses in the range 0 to 3% of the total notional principal. The mezzanine tranche is responsible for default losses in the range 3 to 7% of the total notional principal. Other tranches are responsible for losses in the ranges 7 to 10%, 10 to 15%, 15 to 30%, and 30 to 100% of total principal. The buyer of protection pays a predetermined annual premium (known as a spread) on the outstanding tranche principal and is compensated for losses that are in the relevant range. (In the case of the equity tranche the arrangement is slightly different: the buyer of protection pays a certain percentage of the tranche principal upfront and then 500 basis points on the outstanding tranche principal per year.)

Several other standard portfolios and associated tranches have been defined. For example, iTraxx Europe is an equally weighted portfolio of 125 investment grade European credit exposures. The tranches for this portfolio are 0 to 3%, 3 to 6%, 6 to 9%, 9 to 12%, 12 to 22%, and 22 to 100%. The most popular life of a CDO is five years. However, 7-year, 10-year, and to a lesser extent 3-year CDOs now trade fairly actively.

Tranches of nonstandard portfolios are regularly traded. These are referred to as “bespokes.” Bespoke portfolios differ in the names that are included in the portfolio, the average CDS spread for the names in the portfolio, and in the dispersion of the CDS spreads. The approach to estimating tranche spreads for a bespoke depends on its characteristics. If the portfolio consists almost entirely of investment grade North American companies, it should be bench-marked to the market quotes for CDX NA IG. If it consists of more risky North American companies it should be bench-marked to the market quotes for CDX NA IG and CDX NA HY. When a portfolio is primarily European iTraxx quotes should be used; other portfolios that consist of both European and North American companies should be bench-marked to both CDX and iTraxx quotes.

The standard market model for valuing tranches of synthetic CDOs is a one-factor Gaussian copula model for time to default. This was proposed by Li (2000) and Gregory and Laurent (2005). Traders often imply what are termed base correlations from the model. The base correlation for a loss level of X% is the correlation which, when substituted into the Gaussian copula model, produces an expected loss for the 0 to X% tranche that is consistent with that calculated (again using the one-factor Gaussian copula model) from the market. Typically the base correlation is an increasing function of X.

As explained in Baheti and Morgan (2007), traders have tried various approaches for calculating base correlations for a bespoke portfolio from the base correlations for a standard portfolio. For example, three commonly used approaches (all using calculations that are based on the one-factor Gaussian copula model) are:

1. ATM Mapping: If the ratio of the standard portfolio expected loss to the bespoke portfolio expected loss is α, it is assumed that the 0 to X% tranche of the bespoke portfolio is valued with the same correlation as the 0 to αX% tranche of the standard portfolio.
2. Probability matching (PM): If the probability of losses exceeding X% for the bespoke portfolio is the same as the probability of losses exceeding αX% for the standard portfolio, it is assumed that the 0 to X% tranche of the bespoke portfolio can be valued with the same correlation as the 0 to αX% tranche of the standard portfolio. The base correlation for the standard portfolio corresponding to attachment point αX% is used in conjunction with the standard Gaussian copula to calculate the probabilities in both cases.
3. Proportional tranche loss matching (TPL): If the expected loss on a tranche of the bespoke portfolio with detachment point X% as a proportion of the bespoke portfolio expected loss is the same as the expected loss on a tranche of the standard portfolio with detachment point αX% as a proportion of the standard portfolio expected loss, it is assumed that the 0 to X% tranche of the bespoke portfolio can be valued with the same correlation as the 0 to αX% tranche of the standard portfolio. The base correlation for the

standard portfolio corresponding to attachment point αX% is used in conjunction with the standard Gaussian copula to calculate the expected tranche losses in both cases.

The problem with these approaches is that it is not easy to develop intuition about base correlation.[1] Also, interpolating base correlations is fraught with difficulties. As shown by Hull and White (2006), simple procedures for interpolating base correlations give poor results when used to value nonstandard tranches of standard portfolios. Presumably, the results they give for the nonstandard tranches of nonstandard portfolios are at least as bad as the results for standard portfolios. This point is now generally recognized by the market and the use of base correlation as an interpolation tool is not as popular as it was in the early 2000s.

II. THE IMPLIED COPULA APPROACH

Hull and White (2006) suggest what has become known as the implied copula approach. In the simplest version of the model the hazard rate, λ, over the life of a CDO is a constant and the same for each company. Hull and White show that defining a probability distribution for the hazard rate, λ, is equivalent to defining a one-factor copula model.[2] The hazard rate can be thought of as a variable defining the severity of the credit environment over the life of the CDO. It plays a similar role to the level of the underlying factor in the one-factor Gaussian copula model.

The procedure for deriving the probability distribution for λ in Hull and White (2006) is as follows:

1. Choose a set of representative hazard rates from the very low to the very high. In the homogeneous version of the model these hazard rates apply to all companies.
2. Search for probabilities to assign to the hazard rates so that the index and all tranche quotes are matched as closely as possible.
3. Include in the objective function a term that penalizes probability distributions that are not smooth.

[1] Two base correlations are necessary to value a given tranche. The value of the U_1% to U_2% tranche depends on the base correlation for the 0 to U_1% and 0 to U_2% tranches.

[2] Conversely defining a one-factor copula model is equivalent to defining a "probability distribution for hazard rate term structures," that is a set of hazard rate term structures and their associated probabilities.

Inglis and Lipton (2007) have proposed a version of the implied copula model where there are only four different hazard rates. The lowest hazard rate is zero and the highest is infinite. The other two hazard rates and the probabilities assigned to the hazard rates are chosen to fit market data. The model has five free parameters (two hazard rates and three probabilities) and can fit market data well. However, it is a "highly discrete" representation of possible outcomes. As the authors point out, in the limit of a large homogeneous portfolio losses are concentrated at 0%, 100%, and just two intermediate points.

III. A PARAMETRIC IMPLIED COPULA: HOMOGENEOUS VERSION

We now present a parametric version of the implied copula model that we have found to be useful in applications of the model. We will refer to this as the "log-t implied copula". It assumes that the variable

$$\frac{\ln\lambda-\mu}{\sigma}=t_\nu \tag{1}$$

has a Student *t* distribution with ν degrees of freedom. This means that three free parameters, μ, σ, and ν, describe the probability distribution of λ. In many applications of the Student *t* distribution ν is an integer, but the distribution can be generalized so that ν is any positive number. We use the generalized version of the distribution. The probability density of t_ν is

$$f(t_\nu)=\frac{1}{\sqrt{\nu\pi}\;B(1/2,\nu/2)}\left(1+\frac{t_\nu^2}{\nu}\right)^{-(\nu+1)/2}$$

where B is the beta function. The cumulative probability distribution can be calculated from the incomplete beta function.[3]

From equation (1)

$$\lambda=\exp(t_\nu\sigma+\mu)$$

In calibrating the model to tranche quotes and index spreads the variable μ is primarily influenced by the level of the index. If ν and σ remain the same while μ increases (decreases),

[3] See, for example, Press et al (1991).

the distribution is stretched out (compressed) but retains its original shape. Suppose that μ changes from μ_1 to μ_2. After the change, a hazard rate of $k\lambda$ has the same probability density as λ did before the change where $k = \exp(\mu_2 - \mu_1)$. The coefficient of variation of the distribution remains the same.

As shown in Hull and White (2006), correlation in the implied copula model is governed by the dispersion of the hazard rate distribution. Both σ and ν therefore influence correlation. In general correlation increases as σ increases or ν decreases. The variable ν determines the heaviness of the tail of the distribution and has most impact on the pricing of senior tranches relative to other tranches.

The homogeneous one-factor version of the log-*t* implied copula is implemented as follows. A total of *n* hazard rates, λ_k $(1 \le k \le n)$, are chosen. (We typically choose $n = 100$.) These apply to all companies for the whole life of the CDO tranches. The lowest hazard rate, λ_1, is set so that there is virtually no chance of any default (we use $\lambda_1 = 10^{-8}$). The highest hazard rate, λ_n, is set to a value where all companies default almost immediately (we use $\lambda_n = 100$). The intermediate hazard rates are chosen so that the $\ln \lambda_k$ are equally spaced.

The present values of payments (including accrual payments), A_k, and payoffs, C_k, are calculated for each hazard rate, λ_k, for each tranche of the CDO assuming a principal of \$1. Trial values are chosen for μ, σ, and ν. These determine the probability, π_k, that applies to hazard rate, λ_k. The probabilities are $\pi_1 = F(q_1)$, $\pi_n = 1 - F(q_{n-1})$, and

$$\pi_k = F(q_k) - F(q_{k-1}) \qquad \text{for } 2 \le k \le n-1$$

where $q_k = 0.5(\lambda_k + \lambda_{k+1})$ and F is the cumulative probability distribution function for the current parameter values. These probabilities enable expected payoffs, $C = \sum \pi_k C_k$, and expected payments, $A = \sum \pi_k A_k$, to be calculated. For most tranches the "model quote" is C/A. For tranches involving an upfront payment and subsequent payments at a rate of r per year the model quote is $C - rA$. An iterative search procedure is used to find the values of μ, σ, and ν that minimize the sum of the squared differences between the model quotes and the market quotes.[4] (This is very fast because the λ_k's do not change and so the present value of payments and present value of payoffs for each λ_k have to be calculated only once.)

Numerical Example

To illustrate the approach with a simple example, we assume that 5-year quotes for the tranches of iTraxx and the index are as shown in the second column of Table 1. We assume that tranches last exactly five years and the zero curve is flat at 4%. A total of 100 hazard rates, chosen as described above, are used and the recovery rate is assumed to be 40%. Notional recoveries are assumed to pay down the super senior (22 to 100%) tranche.

The results of the calibration are shown in the third and fourth column in Table 1. The best fit values of μ, σ, and ν are −5.5190, 0.4977, and 1.8159, respectively. The root mean square error is 0.26.

Results from Fitting iTraxx and CDX Quotes

We have fitted the model to iTraxx and CDX NA IG 5-year and 10-year quotes from September 28, 2005 to February 29, 2008. The recovery rate was assumed to be 40%. The results are shown in Figures 1 to 4.[5]

Figure 1 shows that the parameter μ follows a similar pattern for the four data sets. As spreads increase (e.g., when moving from 5-year iTraxx to 5-year CDX NA IG) μ increases. This is as one might expect. As mentioned earlier, the calibrated μ parameter is primarily influenced by the level of the index and the CDX NA IG indices were higher than the corresponding iTraxx indices for the whole of the period considered.

Figure 2 shows that the parameter, σ, for the four sets of quotes are remarkably similar on any given day. Figure 3 shows that the same is true of ν. As mentioned earlier the impact of increasing μ while keeping σ and ν the same is to stretch out the hazard rate distribution in such

---

[4] There are a number of alternative objective functions. Proportional error rather than absolute errors can be considered; different weights can be assigned to different errors. Instead of minimizing squared errors one can minimize the sum of squared values of the tranches and the index, as suggested by Hull and White (2006)

[5] In producing Figures 1 to 4 we constrained day-to-day changes in ν to be less than 0.1. This had very little effect on the goodness of fit.

a way that it retains its shape. The charts therefore suggest that the shape of the implied hazard rate distribution for CDX IG NA is approximately a "stretched out version" of the implied hazard rate distribution for iTraxx Europe. This led us to carry out a fifth calibration where all four data sets are fitted simultaneously using six parameters. Each of the four data sets has its own value of μ while σ and ν are common to all four data sets. The results labeled "All" in Figures 2 and 3 show the values of σ and ν obtained from this calibration.

Figure 4 shows that for most of the period considered the root mean square error in the five-year quotes was less than three basis points and the RMSE for the 10-year tranches was about 5 basis points.[6] During periods of market stress the model fits less well. For example, the root mean square error rose at the start of the subprime mortgage crisis at the end of July 2007. It recovered somewhat between mid August and the end of November, but was high in December, January and February. One reason for our results may be that quotes are less reliable during periods of market stress. Another may be that tranches are priced more erratically during these periods.

A Two-Parameter Version

Figures 2 and 3 show that σ and ν tend to move together. This is not surprising. Both are measures of the dispersion of the distribution. A high σ combined with a high ν produces a similar result to a low σ combined with a low ν.

This suggests that the model can be simplified if σ and ν are replaced by a single parameter. After some experimentation we chose to set $\nu = 2.5$. This produces a more stable model. (Most of the time the three-parameter version of the model works well, but we did encounter some situations when the optimizer tried to set both σ and ν very high or very low.) Figures 5 and 6 show the best fit values of μ and σ given by the model when $\nu = 2.5$. Figure 6 shows that the implied σ parameters are remarkably similar across the four data sets. The implied σ's for 5-year iTraxx and 5-year CDX IG are almost identical; the same is true for 10-year iTraxx and 10-year CDX IG. The root mean square errors are in Figure 7. As one would expect, they are greater than those in Figure 4, but still quite reasonable for much of the period considered.

[6] In the case of the equity tranche the error is the error in the percentage upfront payments, not basis points.

Determinants of σ

The common behavior of σ across all indices suggests that it is a fundamental factor in pricing credit risk. We explored the relationship of the implied σ for the iTraxx 5-year quotes to the level and slope of the term structure[7], the S&P500, the VIX index, and the 5-year iTraxx index by regressing the proportional changes in σ on the proportional changes and lagged proportional changes in the explanatory variables. The regression was done using the full sample and then using two sub-samples, the period up to the end of June 2007 before the credit crisis, and the period after the end of June 2007. Only the iTraxx index was statistically significant in all cases. Lagged changes in the S&P500 and the level of interest rates were significant in some of the regressions.

Many authors (for example Campbell and Taksler (2003), Schaefer and Strebulaev (2004), Schneider, Sögner, and Veža, (2007), and Ahn, Dieckmann, and Perez, (2008)) have observed that equity volatility is closely related to credit spreads. However, our results show that when the level of credit spreads is controlled for equity volatility as reflected in the VIX index had no power to explain default correlations.

The results for regressions that include only variables that are sometimes significant are shown in Table 2. The signs of the coefficients are the same in all three regressions. Increases in the contemporaneous credit spread (the iTraxx index), declines in the S&P500 yesterday, and declines in interest rates yesterday all increase the implied σ.

Relation to Base Correlations

In the two-parameter log-*t* implied copula σ plays a role similar to the copula correlation, ρ, in the one-factor Gaussian copula model. In this section we explore numerically the relationship between σ and ρ.

We use the two-parameter log-*t* implied copula model to generate tranche spreads for different levels of σ and for different index spreads. The model tranche spreads are generated for values of σ between 0.2 and 0.8 in steps of 0.1 and index spreads of 20 to 100 basis points in increments of 10 basis points. In all cases $\nu = 2$. The one-factor Gaussian copula model is then used to imply

[7] We used the 5-year swap rate to represent the level of the term structure and the 10-year swap rate less the one-year swap rate to represent the slope of the term structure.

the base correlations that are consistent with the tranche spreads. Every value of σ and index spread produces a term structure of base correlations. Increasing the value of σ for a given index spread results in higher base correlations while increasing the index spread for a given σ results in lower base correlations. To simplify the reporting of the relation between index spread and σ on base correlation we consider the effect on the average base correlation and the slope of the base correlation curve.[8]

Table 3 shows the results from regressing the average implied base correlation against the values of σ and the index spread. The relationship is highly significant and approximately linear. Increasing σ by one percentage point increases the average base correlation increases by about 0.67 percentage points. Increasing the index spread by 1 basis point reduces the average base correlation by about 0.32 percentage points. The relationship between the log-*t* model parameters and the slope of the base correlation curve is more complex but in general is small. When the index spread is low (20 basis points), increasing σ reduces the slope of the base correlation curve. This curve flattening effect is reduced for higher index spreads. Increasing the index spread while holding σ constant increases the slope of the base correlation curve for high values of σ and decreases it for low values of σ.

IV. HETEROGENEOUS MODEL

The model we have described so far is a homogeneous model in the sense that all companies are assumed to have the same hazard rate probability distribution. One of the attractive features of the log-*t* implied copula model is the ease with which it can be converted into a heterogeneous model.

The model is made heterogeneous by allowing the hazard rate distribution for company i to be log Student t with parameters μ_i and σ_i. The parameter ν is common to all companies. (The analogue of the two-parameter model mentioned earlier is obtained by setting $\nu = 2.5$.) Roughly speaking, the parameter μ_i reflects the company's credit risk and the parameter σ_i is a measure of its default correlation with other companies. Low values of the parameter σ_i correspond to firms

[8] The average base correlation is the arithmetic average of the base correlations for detachment points 3%, 6%, 9%, 12% and 22%. The slope is the difference between the 22% base correlation and the 3% base correlation.

whose credit risk is not highly correlated with that of other firms. In practice, it is likely that equity correlations will be used as a guide to determining the σ_i. The average σ_i is determined by tranche quotes (in the same way that σ is determined by tranche quotes in the homogeneous model). It is therefore assumed that equity correlations, or other data, is used to determine the σ_i only to within an arbitrary multiplicative constant. This means that $\sigma_i = \alpha\sigma_i^*$ where σ_i^* is an input to the model and α is a free parameter used to calibrate the model to market data.

To implement the heterogeneous model total of n points are selected from the t_ν distribution, $t_{\nu,k}$ $(1 \le k \le n)$, with $t_{\nu,k+1} > t_{\nu,k}$ for all k. (We typically choose n = 100.) Analogously to the homogeneous case, each $t_{\nu,k}$ has a probability π_k where $\pi_1 = F(q_1)$, $\pi_n = 1 - F(q_{n-1})$, and

$$\pi_k = F(q_k) - F(q_{k-1}) \qquad \text{for } 2 \le k \le n-1$$

$q_k = 0.5(t_{\nu k}+t_{\nu k+1})$ and F is the cumulative probability distribution function for the t-distribution with ν degrees of freedom. The values of $t_{\nu,k}$ are chosen so that the π_k's are roughly equal.

The hazard rate for the kth company is $\lambda_{ki} = \exp(t_{\nu,k}\sigma_i+\mu_i)$. As in the case of the homogeneous model, one factor determines the hazard rates of all different companies.

The procedure for implementing the model is as follows:

1. Choose trial values of α and ν
2. Find, for each company i, the value of μ_i that matches its CDS spread.
3. Determine hazard rates λ_{ki} $(1 \le k \le n)$ for the ith firm using $\lambda_{ki}=\exp(t_{\nu,k}\alpha\sigma_i^*+\mu_i)$,
4. Use procedures in Andersen et al (2003) and Hull and White (2004) to value CDO tranches.
5. Search for values of α and ν that minimize the sum of squared differences between the model quote and the market quote.

How Important is it to Use a Heterogeneous Model?

To test the difference between the prices given by the homogeneous and heterogeneous model we consider the model spreads for tranches of a five-year CDO for homogeneous and heterogeneous portfolios of 125 companies. The attachment and detachment points are the same as for iTraxx Europe.

Heterogeneity is introduced by allowing each firm to have a different μ and a different σ. Changing these firm characteristics changes the credit risk and correlation for the firms and may change the average level of portfolio credit risk and default correlation. These are in turn liable to have an effect on tranche spreads that is unrelated to the heterogeneity. To control for these effects we ensure that the portfolio index spread always equals a predetermined level, S_P, and that the equity tranche spread remains constant as heterogeneity is introduced.

Initially we consider only heterogeneity in credit risk, setting $\nu = 2$. The CDS spreads for the 125 companies are drawn from a lognormal distribution where m_1 and s_1 are the mean and standard deviation of the logarithm of the credit spread. The level of m_1 determines the average credit spread and the value of s_1 determines the variability of credit spreads. When $s_1 = 0$ (the homogeneous model), σ is set equal to 0.5.

Five different values of s_1 (0, 0.25, 0.5, 0.75, and 1) are considered. For each s_1 the value of m_1 is chosen so that the 125-name model index spread equals S_P. For all values of s_1 greater than zero σ is chosen to produce the same equity tranche spread as the s_1=0, σ =0.5 case. (In practice the value of σ when s_1>0 is close to but not equal to 0.5.) Two different values of the index spread are considered, 50 and 200 basis points. The 50 basis point spread is representative of iTraxx or CDX index spreads that are seen in practice and the 200 basis point spread is similar to spreads seen in some more credit-risky bespoke CDOs.

The results are shown in Tables 4 and 5. The dispersion of the spreads in CDX NA IG and iTraxx Europe correspond to a value of s_1 of about 0.25. However, the dispersion of spreads in many bespoke portfolios corresponds to a value of s_1 of about 0.75. The tables therefore show that the impact of moving from a homogeneous to a model that is heterogeneous in credit risk is fairly small for CDX NA IG or iTraxx Europe, but can be quite large for bespoke portfolios.

We now consider the case in which there is no heterogeneity in credit risk (all the firms have the same CDS spread, $s_1 = 0$) but the σ's differ from company to company. As before $\nu = 2$ and we consider index spreads of 50 and 200 basis points. The scaling parameter, α, is set to one and the σ_i^* are drawn from a lognormal distribution with where the mean and standard deviation of the logarithm of the variable are m_2 and s_2. Five values of s_2 are considered 0.0, 0.1, 0.2, 0.3, 0.4, and 0.5. The $s_2 = 0$ case corresponds to the homogeneous model in which case m_2 is chosen so that $\sigma = 0.50$. This is the case reported in the top lefthand corner of Tables 4 and 5. When s_2 is not zero, m_2 is chosen so that the spread for the equity tranche remains unchanged. This ensures that the default correlation for the equity tranche remains constant.

In all cases the average value of σ is about 0.50, the standard deviation of the σ's is approximately one-half of s_2 and the 95% confidence interval for the σ's is about twice the value of s_2. For example, in the case in which $s_2 = 0.20$ the 95% confidence interval for the firm σ's is from about 0.30 to about 0.70. Relating this to the base correlation regression results reported in Table 3 this roughly corresponds to a range of base correlations of about 25%.[9] The results are shown in Tables 6 and 7 which show that non-homogeneity in default correlation is much less important than non-homogeneity in spreads.

The final case considered is that in which both CDS spreads and σ_i^* are drawn from correlated log-normal distributions with parameters m_1, s_1, m_2 and s_2. The correlation between the distributions is β. The procedure followed is to choose β, s_1, and s_2. The variable m_1 is then set to a level that produces an index spread of 50 basis points and the variable m_2 is set to a level that produces an equity tranche spread of 28.0%. The target equity tranche spread is the same as for the case in which s_1 and s_2 are zero and σ is 0.50 (that is, the case reported in the first column of Table 6).[10]

Two cases are considered. In the first case $s_1 = 0.25$ and $s_2 = 0.20$. This is roughly consistent with the level of non-homogeneity seen in investment-grade CDX or iTraxx portfolios. The range of CDS spreads in the portfolio is between 20 and 80 basis points and individual firm σ's are between 0.30 and 0.80. The second case has $s_1 = 0.50$ and $s_2 = 0.40$. This is similar to the level of non-homogeneity seen in some bespoke portfolios. In both cases correlations between –0.50 and +0.50 are considered.

The results for the two cases are shown in Tables 8 and 9. The results in Table 8 should be compared with the results in the second column of Table 4 ($s_1 = 0.25$) and the results in the third

---

[9] If all the σ's were 0.70 the average base correlation would be about 25% higher than it would be if all the σ's were 0.30.

[10] In practice, due to interactions between the two distributions m_1 and m_2 are chosen simultaneously.

column of Table 6 (s_2 = 0.20). The results in Table 9 should be compared with the results in the third column of Table 4 (s_1 = 0.50) and the results in the fifth column of Table 6 (s_2 = 0.40). Table 8 shows that when s_1 = 0.25 and s_2 = 0.20 the correlation between the credit spread and σ, the default correlation, does not play an important role. The tranche spreads are close to the average of the corresponding spreads in Tables 4 and 6. Table 9 shows that when s_1 = 0.50 and s_2 = 0.40 the correlation between the credit spread and σ is more important and there is more variability in the tranche spreads. This supports our previous observation that non-homogeneity does not play an important role for investment-grade portfolios but may be important in bespoke portfolios where credit spreads are higher.

V. BESPOKE VALUATION

As with a nonstandard derivative, the valuation of a bespoke CDO depends on the market data available. In this section we consider a few alternative situations and suggest ways of proceeding.

Let us start with the simplest possible case, a bespoke portfolio for which there is a clearly defined reference portfolio. The bespoke portfolio is assumed to differ from the reference portfolio only in the average level of default risk. Suppose, for example, the companies underlying the bespoke CDO are all European. We would then choose the iTraxx portfolio as the reference portfolio. One approach is to fit the two- (or three-) parameter homogeneous model described previously to the index and the tranches of iTraxx Europe and then assume that the σ (or σ and ν) estimated for iTraxx Europe apply to the bespoke. The parameter μ is chosen to match the average spread for the companies underlying the bespoke. Alternatively, a model where credit spreads are heterogeneous can be fitted to the tranches of iTraxx Europe. Again it is assumed that the σ (or σ and ν) estimated for iTraxx Europe applies to the bespoke. In this case a different μ_i is calculated for each company underlying the bespoke.

To explore how this approach relates to the base correlation mapping schemes described in Section I we used the homogeneous version of the log-*t* implied copula to generate tranche quotes for the iTraxx standard tranches. The model parameters were ν = 2 and σ = 0.5. Five different values of μ were used, consistent with index spreads of 50, 75, 100, 150, and 200 basis points. The model tranche spreads for the 50 basis point index were considered to be the quotes for the reference portfolio. The other four cases were considered to be bespoke portfolios with higher levels of default risk.

The model tranche spreads for the reference portfolio were used in conjunction with the Gaussian copula to determine the base correlations for the reference portfolio. The three different base correlation mapping approaches described in Section I were then used to determine the appropriate base correlation for each tranche of the four bespoke portfolios. These were then compared with the base correlations calculated from the log-*t* implied copula model quotes. In all cases the interpolation of the base correlations for the reference portfolio was done using a cubic spline. Since all of the bespoke portfolios have higher default risk, the 3% attachment point for a bespoke portfolio is mapped into an attachment point less than 3% for the reference portfolio, the lowest base known correlation. Base correlations for the reference portfolio then have to be extrapolated rather than interpolated.

The average difference across all four bespokes and all five detachment points in each bespoke (3%, 6%, 9%, 12% and 22%) between the base correlations calculated using the ATM mapping and the base correlations consistent with the log-*t* version of the implied copula is –0.2%. The maximum and minimum differences are 2.1% and –3.6% respectively. The mapping tends to overstate the base correlation relative to the log-*t* implied copula model for low attachment points and understate it for high attachment points. Using the probability matching (PM) mapping the corresponding results are –0.8%, 1.7% and –2.2%. Again the mapping tends to overstate the base correlation for low attachment point and understate it for high attachment points. The tranche loss mapping (TPL) produced results that were most different from the log-*t* implied copula. It overestimated the average base correlation by about 6% with a wide range of errors.

An alternative test is to use the base correlation mapping scheme to generate base correlations for the bespokes and then use these base correlations to determine the breakeven tranche spreads for the CDO tranches. These are then compared with the model spreads generated by the log-*t* version of the implied copula. There is a wide range of tranche spreads from a low of about 50 basis points for the 12% to 22% tranche when the index spread is 75 basis points to a high of about 2200 basis points for the 3% to 6% tranche when the index is 200 basis points. To deal with this large range we considered the proportional difference in tranche spreads.[11] The proportional difference is defined as the spread based on the log-*t* model less the spread based on the base correlation mapping divided by the spread based on the log-*t* model.

Tranche spreads that are generated by interpolating base correlations are sensitive to the exact interpolation scheme. As a result the differences in tranche spreads are a result of both the probability mapping procedure and the procedure that was used to interpolate the base correlations. For the ATM base correlation mapping the average proportional spread difference is –5.3%. That is the ATM base correlation mapping produces tranche spreads that are lower than the log-*t* implied copula by about 5% of the correct spread. The standard deviation of the proportional spread difference is 6.3%. The results for the probability mapping (PM) are average proportional errors of –1.8% and a standard deviation of proportional errors of 7.8%. Overall the two mapping schemes produce similar results indicating that they are roughly consistent with the log-*t* implied copula model with constant ν and σ.

In the preceding examples there was only one reference portfolio. As a result ν and σ are held constant and only μ changes when valuing a bespoke. In some case two reference portfolios may be available. For example, if the underlying portfolio consists entirely of North American companies, both CDX NA IG and CDX NA HY can act as reference portfolios. Calibrating the log-*t* implied copula to CDX NA IG and CDX NA HY separately can reveal how σ (or ν and σ) is related to the level of credit risk.[12] The values of σ and ν that are appropriate for the bespoke are then determined by interpolation. Suppose that the CDX NA IG index is x_{IG}, the CDX NA HY index is x_{HY} and the average CDS spread for the companies underlying the bespoke is x_{BE}. Suppose further that σ_{IG} and ν_{IG} are the parameters of the log Student-*t* distribution that is fitted to CDX NA IG while σ_{HY} and ν_{HY} are the corresponding parameters for CDX NA HY. The σ and ν parameters for the bespoke would be chosen as

$$\sigma_{BE} = \frac{(x_{HY} - x_{BE})\sigma_{IG} + (x_{BE} - x_{IG})\sigma_{HY}}{x_{HY} - x_{IG}} \qquad \nu_{BE} = \frac{(x_{HY} - x_{BE})\nu_{IG} + (x_{BE} - x_{IG})\nu_{HY}}{x_{HY} - x_{IG}}$$

[11] Considering spread differences tends to overemphasize errors for high spread tranches. Considering proportional spread differences tends to overemphasize errors for low spread tranches.
[12] A small amount of data was available to calibrate the log-*t* implied copula to CDX NA HY. We found that in the two-parameter version of the model σ was higher for CDX NA HY than for CDX NA IG, but appeared to follow a similar pattern through time.

A refinement of this would be to do a separate interpolation between IG and HY values for each name underlying the bespoke.

A Two-Factor Model

The model we have proposed can be extended so that hazard rates are driven by more than one factor. To provide a concrete example, we consider the case of a bespoke portfolio consisting of both European and North American names. For the European names the reference portfolios is iTraxx Europe and for the North American names the reference portfolio is CDX NA IG. The hazard rates of European names have log-*t* distributions and are perfectly correlated; the hazard rates of North American names have log-*t* distributions and are perfectly correlated. However, North American and European hazard rates are less than perfectly correlated.

Consider first the case where there is homogeneity within Europe and within North America. The procedure is as follows:

1. Fit each of iTraxx Europe and CDX NA IG to the two-parameter (or three-parameter) version of the homogeneous (or heterogeneous) version of the model. Define σ_A and σ_E as the North America and European σ values and ν_A and ν_E as the North American and European ν values.

2. Select *n* representative hazard rates for North American companies and *n* representative hazard rates for European companies with their associated marginal probabilities as described in Section III.

3. A Gaussian (or some other) copula is used to define the joint probability distribution of North American and European hazard rates. This results in an *n* by *n* table. Each cell in the table represents a hazard rate for European companies, a hazard rate for North American companies, and the probability that these two hazard rates will occur together.[13]

[13] In practice a few of the companies in some bespoke portfolios may be neither North American nor European. Our (somewhat arbitrary) approach for handling this is to assume that the hazard rates for these companies are an average of what they would be if they were North American and European. As there are typically not many names in this category the way they are treated does not make a great deal of difference to the spread calculated for the bespoke.

4. The present value of expected payoffs and payments are determined for each cell in the table. These are multiplied by the probabilities applicable to the cells and summed to determine the global present values of expected payoffs and payments. These are used to determine a spread for the bespoke.

When there is heterogeneity within Europe and North America the procedure is similar except that the table that is created represents the joint distribution of t's. Each pair of t's is used to generate firm specific hazard rates as described in Section IV.

The approach can be extended so that hazard rates are assumed to be driven by more than two factors. However, the calculations in a multi-factor model are more time consuming than those in a one-factor model. In a one-factor model if $n = 100$ the tranche value must be calculated 100 times. In the two factor model the tranche value must be calculated 10,000 times (possibly involving the Andersen et al (2003) or Hull and White (2004) procedure). For a three-factor version one million calculations would be required.

As an example we consider the case where there are 50 European and 50 North American investment grade names. The model parameters we assume for iTraxx and CDX are shown in Table 10. These parameters are roughly consistent with market conditions in December 2007. The tranche quotes that are consistent with the assumed model parameters are shown in Table 11. The log of the CDS spreads for the European names are drawn from a normal distribution with mean –5.4 and standard deviation 0.40. The resulting portfolio is similar to but slightly more risky than the iTraxx portfolio. The lowest CDS spread is about 18 basis points, the highest is about 115 basis points and the average spread is about 49 basis points. The 50 North American names are sampled from a similar normal distribution with mean –4.9 and standard deviation 0.60. The resulting portfolio is similar to but slightly more risky than the CDX NA IG portfolio. The lowest CDS spread is about 18 basis points, the highest is about 301 basis points and the average spread is about 89 basis points.

The joint hazard rate distributions for European and American investment grade firms are generated using a Gaussian copula. To avoid any confusion with the standard market Gaussian copula model, we will refer to this as the "Eur-Am hazard rate copula." The joint distribution hazard rate distribution is used to calculate the spreads for the index and tranches 0 to 3%, 3 to 5%, 5 to 10%, 10 to 15%, 15 to 30% and 30 to 100%. The tranche spreads are reported in Table

12 for Eur-Am hazard rate copula correlations of 0.0, 0.5, 0.9, 0.95, and 0.99. (Setting the correlation equal to 0.99 creates a situation close to the one-factor model.) For comparison purposes the table also shows the breakeven spreads for a 100-name portfolio that is completely European and one that is completely American. The spreads for the names in the two comparison portfolios are chosen using the same procedures used to select the spreads for the 50-name portfolios.

First consider the results for the single risk-source portfolios (all American or all European names). When the tranche spreads for these portfolios are compared with those for iTraxx and CDX shown in Table 11 we notice some significant differences. Some of this is a result of the fact that the bespoke portfolios are a little more risky on average than the comparable index portfolios. Some of the difference is a result of the somewhat different attachment and detachment point for the tranches, and some is due to the fact that as portfolio size is reduced risky tranches tend to become more risky and senior tranches tend to become less risky.[14]

Now consider the results for the bespoke portfolio. As the correlation between European and North American hazard rates increases, the breakeven spread for junior tranches (tranches with low attachment and detachment points) declines while the breakeven spread for senior tranches rises. This effect is similar to, but much less extreme than, the changes observed when the copula correlation is increased in the market-standard Gaussian copula model because the change in correlation affects only the diversification between European and North American default risk. The equity tranche spread is always higher than the average of the single risk-source spreads and all other spreads are always lower than the average. Figure 11 shows the breakeven spread for each bespoke tranche as a proportion of the breakeven spread for that tranche when the Eur-Am hazard rate copula correlation is 0.99 (i.e. when there is effectively only one factor). The equity and 3% to 5% tranche spreads decline uniformly as correlation rises. The super senior, 30% to 100%, tranche spread rises uniformly as correlation rises. For all other tranches spreads initially rise and then decline as correlation rises. The impact of the correlation between factors is most marked for the super senior tranche.

[14] See Hull and White (2006) for a discussion of this phenomenon.

Hedging

Based on the results in this section a reasonable approach to managing a portfolio of bespoke tranches would seem to be to use the one-factor or two-factor model where credit spreads, but not credit correlations, are assumed to be heterogeneous. As shown in Section III very little is sacrificed by setting ν=2.5 so that σ is the only free correlation parameter. The relevant Greek letters can be calculated by perturbing each credit spread and the value of σ. Our research suggests that for an equally weighted portfolio, a reasonable estimate of the impact of increasing any given credit spread by, say, 10 basis points can be obtained by increasing all credit spreads by 10 basis points and then dividing by the number of names

VI. CONCLUSIONS

We have presented a new version of the implied copula model. It has a number of advantages over the previous version. It is based on a small number of parameters and is more robust. Furthermore, the transition from a homogeneous to a heterogeneous is much easier.

The two parameter homogeneous version of the model provides a single implied correlation measure, σ, on any given day. We have attempted to relate changes in σ to a number of macroeconomic variables. The one with the most explanatory power is the level of credit spreads.

The model provides a natural approach to valuing bespokes. We have compared the results from the homogeneous model with three of the base correlation mapping procedures that have been proposed. The model gives similar results to two of them. The advantage of the model over using a base correlation mapping procedure is that the assumptions being made are transparent and the extension of the model to incorporate heterogeneity and multiple factors is straightforward.

The model provides a way of testing the impact of moving from a homogeneous model to a heterogeneous model. There are two types of heterogeneity. One relates to credit spreads; the other to credit correlation. We find that credit-spread heterogeneity has only a small effect when the variability of spreads is similar to that observed for companies in the iTraxx Europe (or CDX NA IG) portfolio. Credit-spread heterogeneity does have an appreciable effect on pricing when spread variability is higher (as it is in many bespokes). The impact of credit-correlation heterogeneity is smaller than that of credit-spread heterogeneity. This last result should be welcomed by market participants as credit correlation parameters are much more difficult to estimate than credit spread parameters.

The model also provides a way of testing the effect of more than one factor. As an example we considered the case where a portfolio contains European and North American companies and a Gaussian copula model defines the correlation between the hazard rates of the two types of companies. The impact of the correlation between European and North American hazard rates is different for different tranches. For example, as this correlation increases, the breakeven spread for the equity tranche declines modestly while the breakeven spread for the super senior tranche increases very fast.

REFERENCES

Ahn, S. C., Dieckmann, S., and Perez, M. P., "Exploring the Common Factors in the Term Structure of Credit Spreads," Arizona State University working paper, March 2008.

Andersen, L., J. Sidenius, and S. Basu, "All Your Hedges in one Basket," Risk, 16 (November 2003), 67–72.

Baheti, P. and S. Morgan, "Base Correlation Mapping," Lehman Brothers (2007), 3–20.

Campbell, Y., and G.B. Taksler, 2003, "Equity volatility and corporate bond yields", Journal of Finance, 58(6), 2321–2349.

Gregory, J., and J-P. Laurent. "Basket Default Swaps, CDOs, and Factor Copulas," Journal of Risk, 7 (2005), 8–23.

Hull, J. and A. White, "Valuation of a CDO and nth to default CDO without Monte Carlo Simulation," Journal of Derivatives, 12, 2 (Winter 2004), 10–23

Hull, J. and A. White, "Valuing Credit Derivatives Using an Implied Copula Approach, "Journal of Derivatives, 14, 2 (Winter 2006), 8–28.

Inglis, S. and A. Lipton, "Factor Models for Credit Correlation," Risk, 20, 12 (December 2007), 110–115.

Li, D.X. "On Default Correlation: A Copula Approach," Journal of Fixed Income, 9 (2000), 43–54.

Press, W.H., B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling, Numerical Recipes in C: The Art of Scientific Computing, Cambridge University Press, 1988.

Schaefer, S. M., and Strebulaev, I. A., "Structural Models of Credit Risk are Useful: Evidence from Hedge Ratios on Corporate Bonds," London Business School working paper, 2004.

Schneider, P., Sögner, L., and Veža, T., "The Economic Role of Jumps and Recovery Rates in the Market for Corporate Default Risk," Vienna University of Economics and Business Administration working paper, 2007.

Table 1
Data Used to Illustrate the Log-t Implied Copula Model

Quotes are in basis points except for the 0 to 3% tranche where the quote is the percentage of the tranche principal that must be paid upfront in addition to 500 basis points per year

	Market	Model	Error
0-3%	23.00	23.30	0.30
3-6%	160.00	159.98	-0.02
6-9%	80.00	80.18	0.18
9-12%	58.00	57.87	-0.13
12-22%	40.00	40.01	0.01
22-100%	10.00	10.41	0.41
Index	49.00	48.60	-0.40

Table 2
Regression of Proportional Changes in Implied σ for 5-Year iTraxx

Proportional changes in the implied σ for 5-year iTraxx are regressed against proportional changes in the iTraxx index, lagged proportional changes in the S&P 500 and level of interest rates. *t*-statistics appear in parentheses.

Period	Intercept	iTraxx 5	S&P lag1	TS lag 1	R^2	*n*
Sep 2005 to Feb 2008	0.002	0.377	–0.469	–0.221	21.6%	492
	(1.926)	(9.67***)	(–3.75***)	(–2.58**)		
Sep 2005 to Jun 2007	0.000	0.109	–0.028	–0.108	1.3%	343
	(0.441)	(2.33**)	(–0.285)	(–1.231)		
July 2007 to Feb 2008	0.005	0.421	–0.823	–0.198	26.8%	149
	(1.320)	(5.93***)	(–3.07***)	(–1.231)		

*** significant at 1%
** significant at 5%

Table 3

Regression of Average Implied Base Correlation on log-*t* Model Parameters

Model tranche spreads are generated using different values of σ and the index spread. The tranche spreads in conjunction with the market-standard Gaussian copula model are used to determine base correlations from which an average base correlation is calculated. The average base correlations are then regressed against σ and the index spread. In all cases ν = 2. *t*-statistics appear in parentheses.

Intercept	Sigma	Index (bp)	R^2	*n*
0.3011	0.6735	–0.0032	97.49%	63
(26.38***)	(41.98***)	(–25.50***)		

*** significant at 1%

Table 4
Model Quotes for a Portfolio where Credit Spreads are Heterogeneous and the Index Level is 50 basis points

The variable s_1 is the standard deviation of the logarithm of the CDS spread of the companies in the portfolio. When s_1=0, σ = 0.50. In all other cases σ is chosen to set the equity tranche spread to the level it has when s_1=0. For all companies ν = 2.

	$s_1 = 0$	$s_1 = 0.25$	$s_1 = 0.50$	$s_1 = 0.75$	$s_1 = 1.00$
0-3%	28.0	28.0	28.0	28.0	28.0
3-6%	183	184	188	193	201
6-9%	80.3	81.1	83.2	86.3	90.0
6-12%	53.7	54.2	55.4	57.2	59.1
12-22%	33.2	33.4	34.0	34.6	34.9
22-100%	8.1	8.0	7.7	7.2	6.8

Table 5
Model Quotes for a Portfolio where Credit Spreads are Heterogeneous and the Index Level is 200 basis points

The variable s_1 is the standard deviation of the logarithm of the CDS spread of the companies in the portfolio. When s_1=0, σ = 0.50. In all other cases σ is chosen to set the equity tranche spread to the level it has when s_1=0. For all companies ν = 2.

	$s_1 = 0$	$s_1 = 0.25$	$s_1 = 0.50$	$s_1 = 0.75$	$s_1 = 1.00$
0-3%	81.2	81.2	81.2	81.2	81.2
3-6%	2214	2221	2240	2273	2315
6-9%	1038	1042	1054	1072	1089
6-12%	528	531	538	549	557
12-22%	211	212	214	217	218
22-100%	25.4	25.1	24.2	23.1	22.1

Table 6

Model Quotes for a Portfolio where Correlations are Heterogeneous and the Index Level is 50 basis points

The variable s_2 is the standard deviation of the logarithm of the σ's. All firms have a CDS spread of 50 basis points, the equity tranche spread is 28.0%, and ν = 2.

	$s_2 = 0$	$s_2 = 0.10$	$s_2 = 0.20$	$s_2 = 0.30$	$s_2 = 0.40$
0-3%	28.0	28.0	28.0	28.0	28.0
3-6%	183	182	182	180	178
6-9%	80.3	80.6	81.6	82.9	84.2
6-12%	53.7	54.1	55.2	56.9	59.1
12-22%	33.2	33.6	35.0	37.0	39.3
22-100%	8.1	8.0	7.8	7.5	7.2

Table 7

Model Quotes for a Portfolio where Correlations are Heterogeneous and the Index Level is 200 basis points

The variable s_2 is the standard deviation of the logarithm of the σ's. All firms have a CDS spread of 200 basis points, the equity tranche spread is 81.2%, and ν = 2.

	$s_2 = 0$	$s_2 = 0.10$	$s_2 = 0.20$	$s_2 = 0.30$	$s_2 = 0.40$
0-3%	81.2	81.2	81.2	81.2	81.2
3-6%	2214	2209	2195	2172	2142
6-9%	1038	1035	1026	1011	992
6-12%	528	527	525	522	518
12-22%	212	212	215	219	225
22-100%	25.4	25.5	25.7	26.1	26.6

Table 8
Model Quotes for Situation Where Both Spreads and Correlations Exhibit Low Level of Heterogeneity

The standard deviation of the logarithm of the σ's is 0.20. The standard deviation of the logarithm of the CDS spreads is 0.25. The index spread is 50 basis points, the equity tranche spread is 28.0% and $\nu = 2$. β is the correlation between CDS spreads and firm σ's.

	$\beta = -0.50$	$\beta = -0.25$	$\beta = 0.0$	$\beta = 0.25$	$\beta = 0.50$
0-3%	28.0	28.0	28.0	28.0	28.0
3-6%	180	182	183	184	186
6-9%	80.4	81.3	82.2	83.0	83.7
6-12%	54.4	55.0	55.5	56.0	56.5
12-22%	34.4	34.7	35.1	35.3	35.5
22-100%	8.0	7.9	7.7	7.6	7.5

Table 9
Model Quotes for Situation Where Both Spreads and Correlations Exhibit High Level of Heterogeneity

The standard deviation of the logarithm of the firm σ's is 0.40. The standard deviation of the logarithm of the CDS spreads is 0.50. The index spread is 50 basis points, the equity tranche spread is 28.0% and $\nu = 2$. β is the correlation between CDS spreads and firm σ's.

	$\beta = -0.50$	$\beta = -0.25$	$\beta = 0.0$	$\beta = 0.25$	$\beta = 0.50$
0-3%	28.0	28.0	28.0	28.0	28.0
3-6%	174.0	178.8	183.6	188.2	192.7
6-9%	81.0	84.0	86.9	89.5	91.9
6-12%	56.2	58.2	60.2	61.8	63.3
12-22%	37.6	38.3	38.7	38.8	38.9
22-100%	7.7	7.3	6.9	6.6	6.3

Table 10
Model parameters used to calculate spreads for bespoke tranches

	iTraxx	CDX
Mu	–5.50	–4.80
Sigma	0.50	0.60
degrees of freedom	2.00	2.50

Table 11
Tranche Spreads Consistent with Model Parameters in Table 10

iTraxx			CDX		
Start	End	Spread	Start	End	Spread
0%	3%	23.8	0%	3%	50.2
3%	6%	151.7	3%	7%	431.2
6%	9%	71.0	7%	10%	152.5
9%	12%	49.5	10%	15%	85.1
12%	22%	33.0	15%	30%	41.4
22%	100%	8.1	30%	100%	7.5
Index		45.5	Index		80.8

Table 12
Breakeven Tranche Spreads for a Bespoke Portfolio

The portfolio contains 50 European and 50 North American investment-grade names. The results for comparable portfolios which are all European or all North American are also shown. Results are based on the parameters in Table 10.

Tranche		Eur-Am Hazard Rate Copula Correlation					Single Risk Portfolio		
Start	End	0.00	0.50	0.90	0.95	0.99	Eur 100	Am 100	Average
0%	3%	48.2	45.7	43.6	43.4	43.2	26.9	54.4	40.7
3%	5%	474.2	466.9	446.4	442.5	439.2	211.6	721.1	466.3
5%	10%	155.3	159.7	152.2	150.1	148.1	82.9	245.4	164.1
10%	15%	68.8	71.5	67.6	66.4	65.1	45.3	95.0	70.2
15%	30%	32.7	34.8	34.2	33.8	33.3	27.3	42.5	34.9
30%	100%	0.4	1.8	4.9	5.5	6.0	6.1	6.5	6.3
Index		67.90	67.9	67.9	67.9	67.9	67.9	48.7	87.6

Figure 1
Value of μ when the Three-Parameter Model is fitted to 5-year and 10-year iTraxx and CDX IG

Figure 2
Value of σ when the Three-Parameter Model is Fitted to
5-year and 10-year iTraxx and CDX IG.
The "All" result is from fitting all four data sets with a single value of σ and a Single value of ν

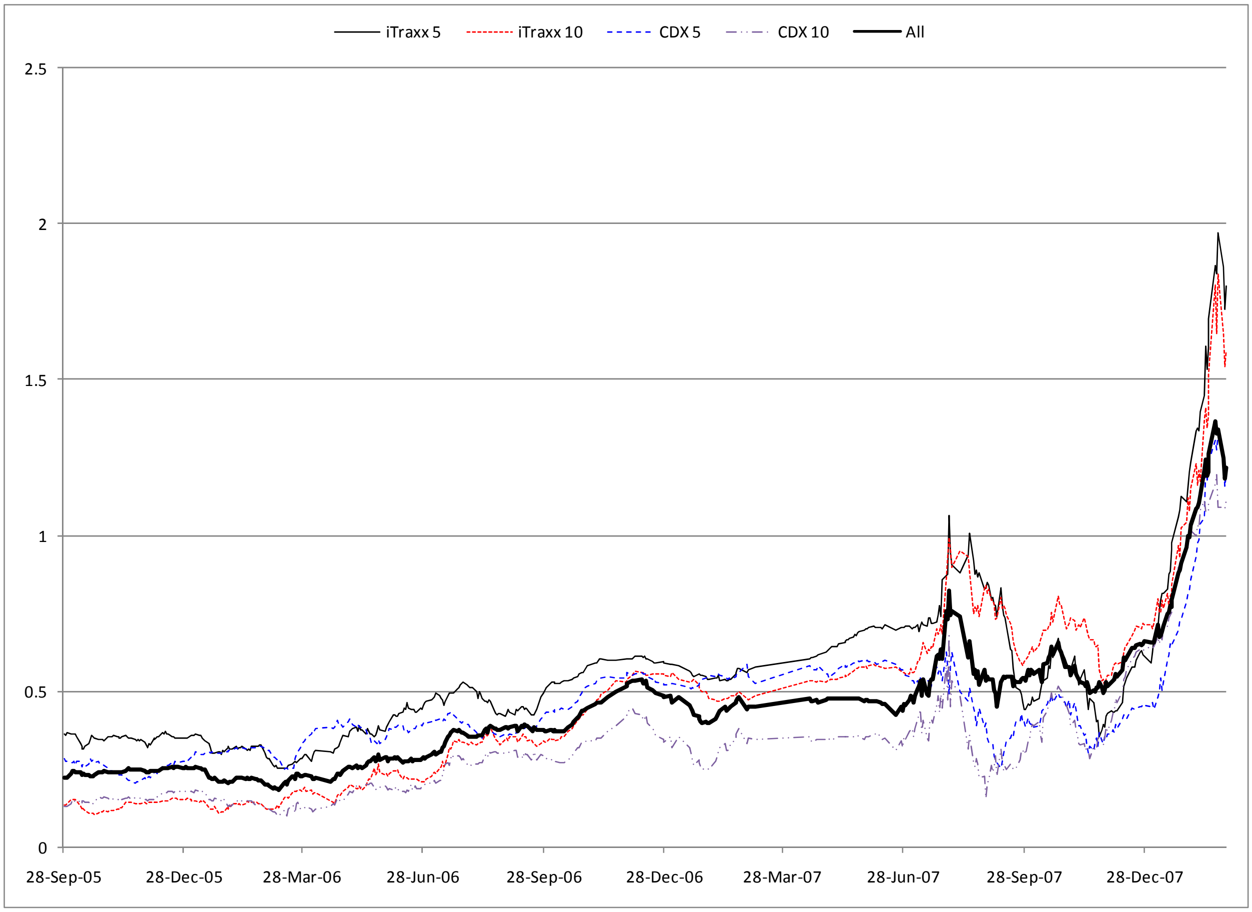

Figure 3
Value of ν when the Three-Parameter Model is Fitted to
5-year and 10-year iTraxx and CDX IG.

The "All" result is from fitting all four data sets with a single value of σ and a Single value of ν

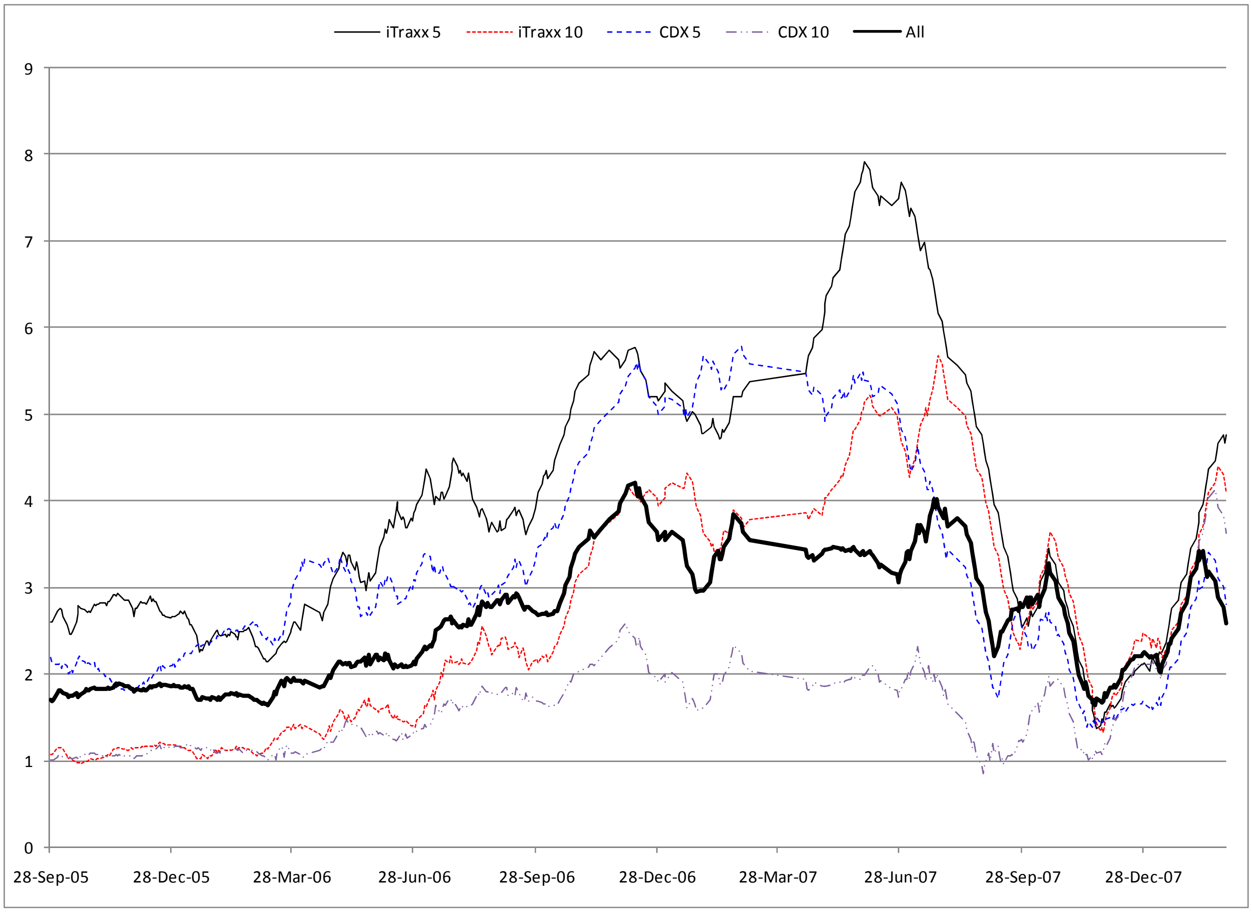

Figure 4
Value of Root Mean Square Error when the Three-Parameter Model is Fitted to 5-year and 10-year iTraxx and CDX IG.

The "All" result is from fitting all four data sets with a single value of σ and a Single value of ν

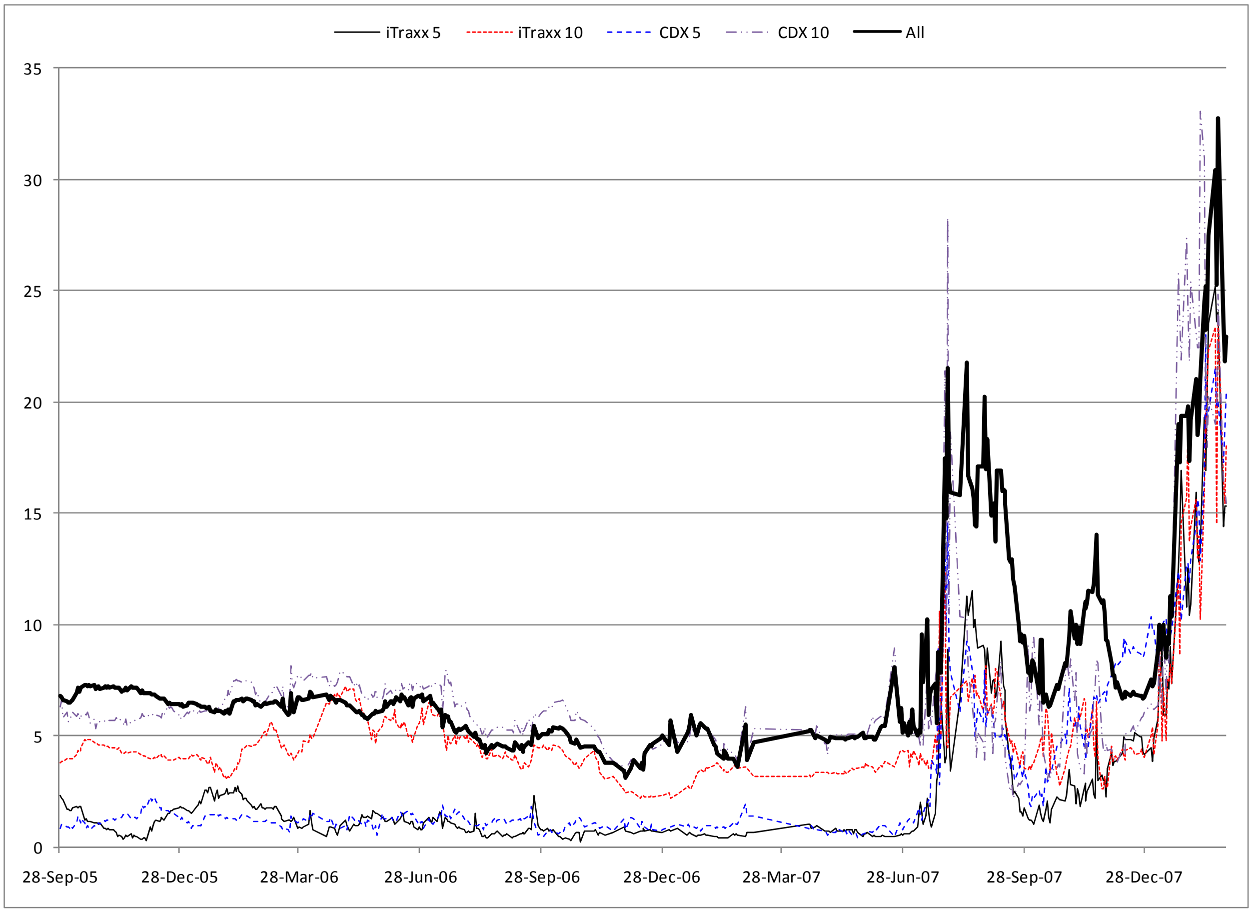

Figure 5
Value of μ when the Two-Parameter Model is Fitted to
5-year and 10-year iTraxx and CDX IG; ν=2.5

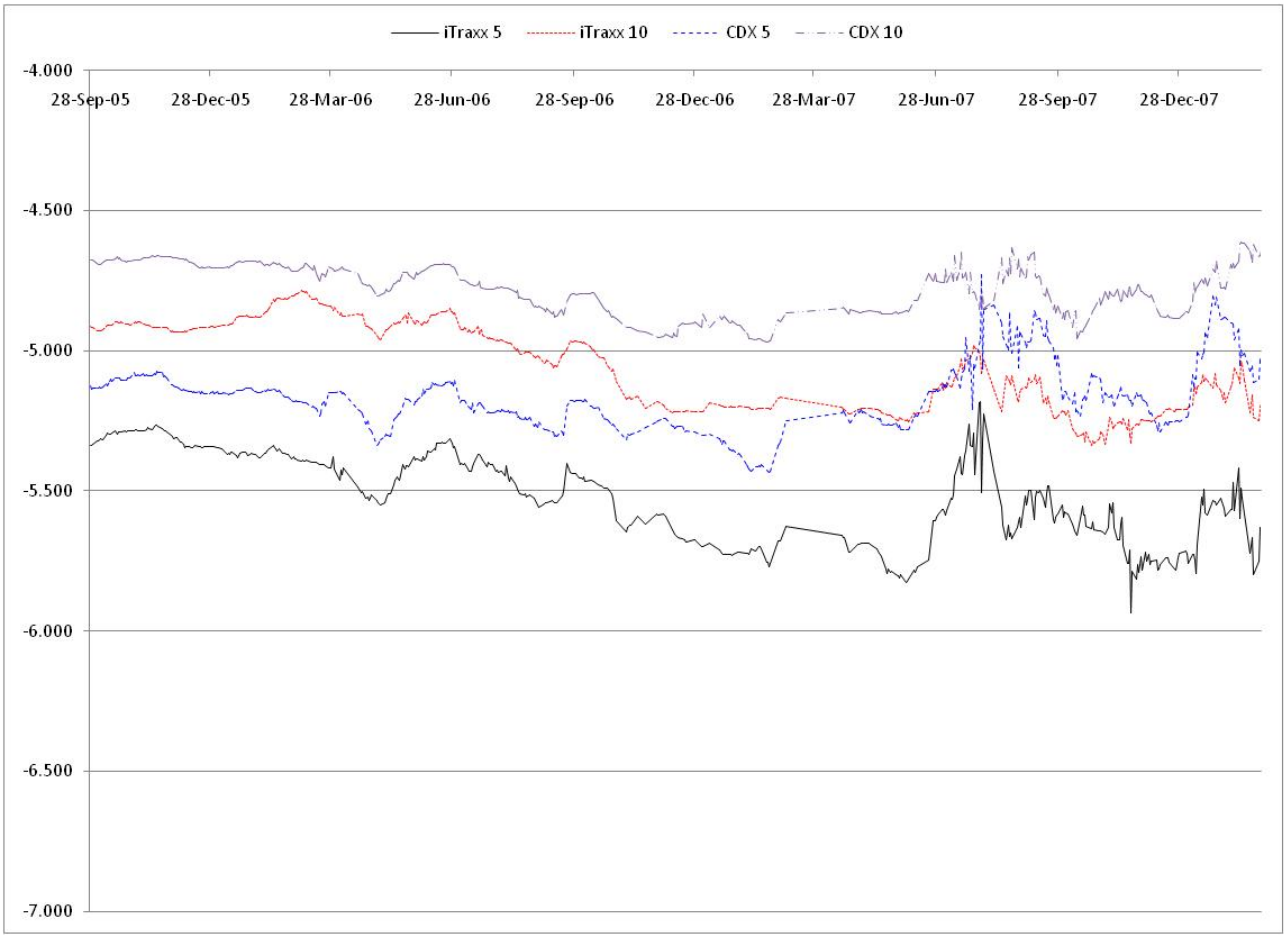

Figure 6

Value of σ when the Two-Parameter Model is Fitted to

5-year and 10-year iTraxx and CDX IG; ν=2.5

The "All" result is from fitting all four data sets with a single value of σ

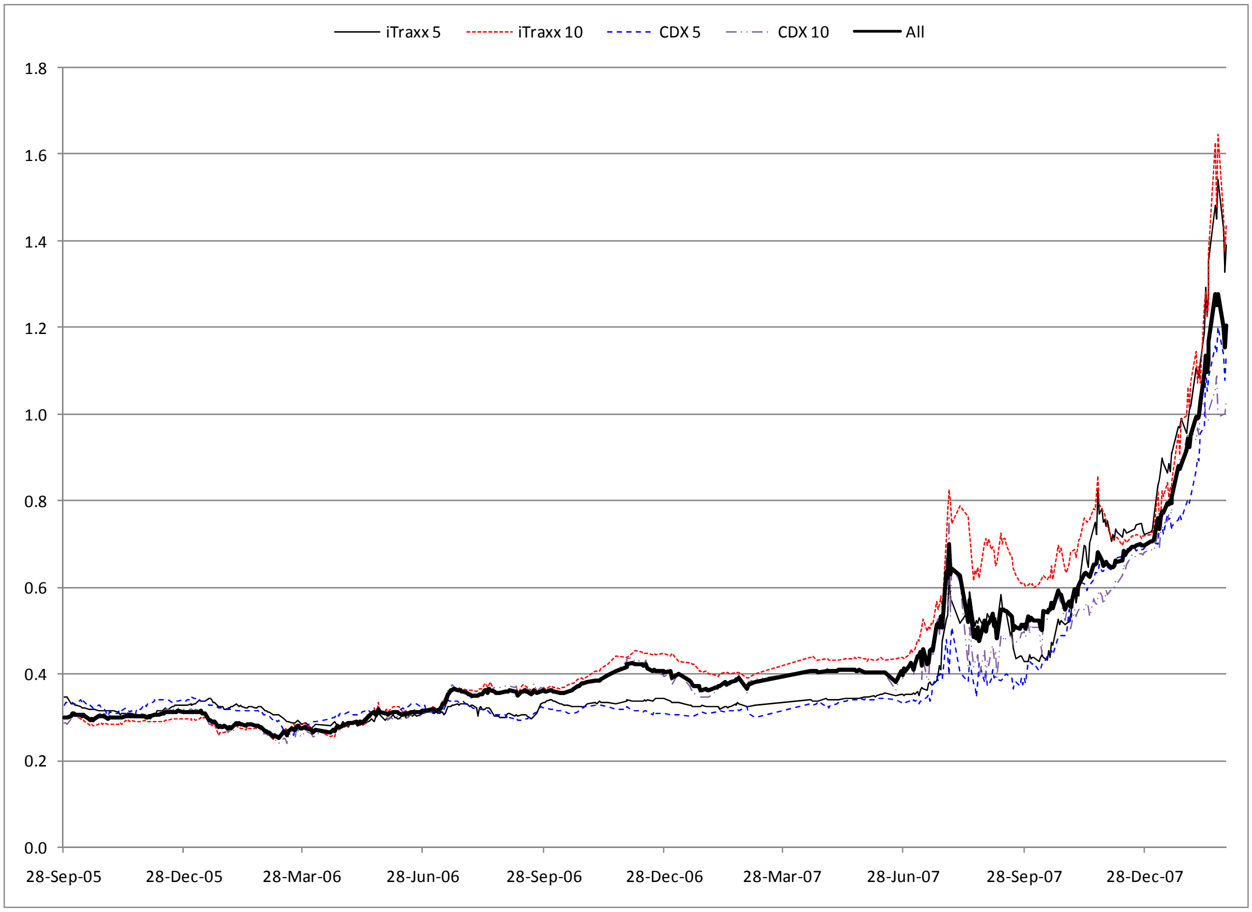

Figure 7
Value of Root Mean Square Error when the Two-Parameter Model is Fitted to 5-year and 10-year iTraxx and CDX IG; ν=2.5
The "All" result is from fitting all four data sets with a single value of σ

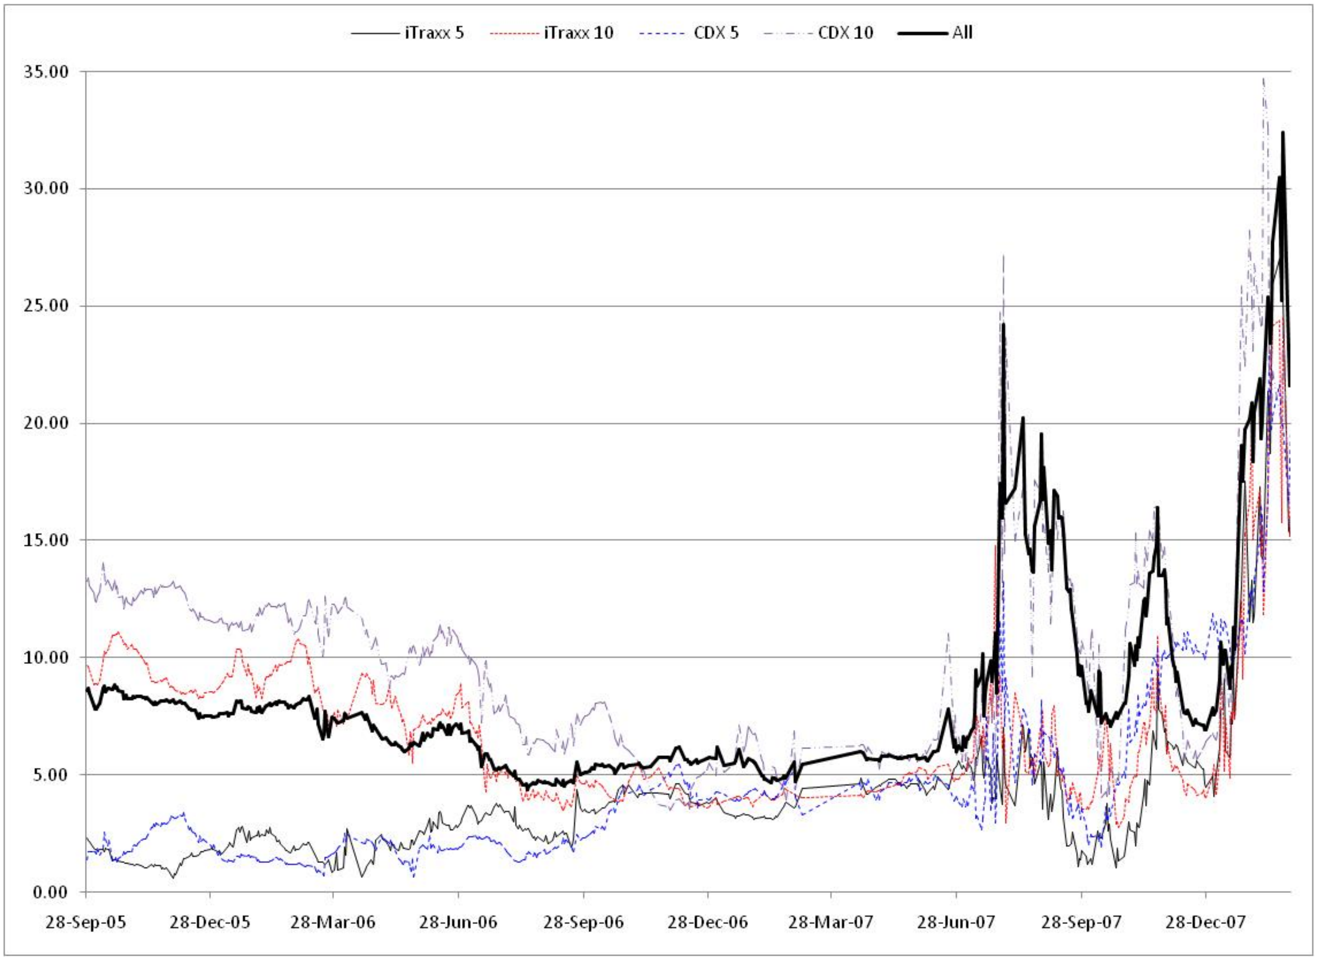

Figure 8

Breakeven tranche spread for different copula correlations as a proportion of the spread when the correlation is 0.99. The scale for the 30 to 100% tranche is shown on the right hand axis.