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 Management
University of Toronto

First Draft: April 2008
This Draft: June 2009

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" model. In this paper we develop a parametric version of the implied copula model and show how it can be used for valuing bespoke CDOs. A two-parameter version of the model is a simple and appealing alternative to the Gaussian copula model. One of the parameters in this model is used to match spreads. The other can be implied from tranche quotes and is much less variable across the capital structure than base correlation. Both homogeneous and heterogeneous versions of the model are presented and the differences between the results obtained from these two versions of the model are examined. Results are also presented for the situation where 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

An 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 option pricing model based on an alternative model of stock price behavior.

In the case of correlation-dependent 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 the tranches of nonstandard portfolios in the synthetic collateralized debt obligation (CDO) market and the prices of tranches of asset-backed securities (ABSs) and ABS CDOs in the cash CDO market.[1] This is a similar type of problem to the one just mentioned faced by options traders, but much more complicated. The usual procedure for extrapolating from the calibration to tranches of standard portfolios to the pricing of tranches of a non-standard portfolio involves what is called a base correlation mapping procedure. Many alternative mappings have been proposed. Unfortunately there is no theoretical or empirical basis to choose one over another or even to determine if the resulting extrapolated prices are reasonable.

[1] The turmoil in credit markets, starting in August 2007, has raised the profile ABSs and ABS CDOs. ABSs were used to tranche out the credit risk in portfolios of mortgages. ABS CDOs were used to tranche out the credit risk in the mezzanine tranches of ABSs.

This paper suggests an alternative way of proceeding. The model proposed is a parametric version of the implied copula model in Hull and White (2006). A two-parameter version of the model is found to fit market data well. One of the parameters is determined by the level of default risk of the portfolio underlying the CDO while the second parameter is related to the degree of default correlation among the names in the underlying portfolio. The value of the first parameter is unique to each portfolio but the value of the second parameter appears to be about the same for all portfolios at each point in time. This means that an implementation of the model that is calibrated to one portfolio can be used to estimate the value of tranches of a second portfolio by merely adjusting the first parameter to reflect the credit risk of the second portfolio. In this way a calibrated version of the model can be used to estimate the correct tranche quotes for different CDOs.

A second advantage of the parametric model is that extending the model from the homogeneous to the heterogeneous case, or from one factor to two factors, is straightforward. The paper takes advantage of this property to explore the difference between the heterogeneous and homogeneous version of the model, and the effect of assuming that hazard rates are driven by more than one factor.

The simplest two-parameter version of the model is an appealing alternative to the one-factor Gaussian copula model. Instead of using tranche quotes to imply base correlations, analysts can imply values for the second (correlation) parameter of the model. These values prove to be remarkably constant both across tranches and across portfolios.

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%. CDX NA HY is a portfolio of 100 high yield North American credit exposures. The tranches for this portfolio are 0 to 10%, 10 to 15%, 15 to 25%, 25 to 35%, and 35 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.[2]

As explained in Baheti and Morgan (2007), traders have tried various methods 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 methods is that they are rules-of-thumb with no theoretical or empirical support. As a result, it is not clear which if any of them will produce the correct results. Also, interpolating base correlations is fraught with difficulties. As shown by Hull and White (2006),

[2] 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.

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 MODEL

Hull and White (2006) suggest what has become known as the implied copula model. 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. 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.

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 start by considering a one-factor implied copula model where hazard rates are constant and the same for all companies. This model will be generalized and extended later in the paper.

The essence of the model is that a probability distribution is assumed for hazard rates. Many different distributions can be assumed. One which we have found useful is the "log Student-*t*" distribution. We 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

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

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), 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.

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

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 λ_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

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

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 $C - rA$. A 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]

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

Numerical Example

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

Results from Fitting iTraxx Europe and CDX NA IG Quotes

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

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 σ-parameters for the four sets of quotes are remarkably similar on any given day. Figure 3 shows that there is also substantial commonality in the best-fit ν.[6] 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 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

---

[5] Data were provided courtesy of Moody's Credit Quotes, www.bquotes.com.

[6] When the number of degrees of freedom is low small changes in ν cause large changes in the fatness of the tails of the distribution. As the number of degrees of freedom increases the distribution becomes more normal and the shape of the tails become much less sensitive to the number of degrees of freedom. As a result, as the best-fit number of degrees of freedom is increased we observe much more variability in the value.

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 the fit of the model to data, when measured using the average root mean square error (RMSE) of tranche spreads, was much worse after the start of the credit crisis in August 2007 than it was before the start of the crisis. Before the crisis the average RMSE was 0 to 3 basis points for five-year iTraxx and CDX IG and 3 to 9 basis points for ten-year iTraxx and CDX IG. After July 2007 they are higher. This is a reflection of a sharp increase in spreads in the period following July 2007. While pricing errors rose, they did not rise as much as the spreads themselves did. If pricing errors are measured as a proportion of the spread the average root mean square (proportional) error declines by about one-third after July 2007.

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. 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.

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 to the level and slope of the term structure[7], the S&P500, the VIX index, and the 5-year CDX index[8] by regressing the changes in σ on the changes and lagged changes in the explanatory variables. The implied σ was the best-fit value for the case in which all four series are calibrated using a common value of σ. The regression was done using the full daily observation sample and a weekly observation sample. In other experiments the sample was divided into the period up to July 31, 2007 and the period from August 1, 2007.

The results are poor. Using daily observations, only changes in the level of the CDX index are always statistically significant in explaining changes in implied σ. In all cases the economic significance and the R^2 are very small. Using weekly changes, only changes in the level of the CDX index are statistically significant in explaining changes in implied σ in the full sample and nothing is statistically significant in the sub-samples, perhaps because of the reduced sample size. Again, the economic significance and the R^2 are very small.

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, we find that equity volatility as reflected in the VIX index had no power to explain changes in the implied σ or by analogy, default correlation.

Relation to Base Correlations

In the two-parameter log-*t* implied copula σ plays a role similar to the base correlation, ρ, in the one-factor Gaussian copula model. Indeed, if a set of σ's are implied from tranches quotes, they contain the same information as base correlations. However, they are much less variable. 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 $\nu = 2.5$. 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.
[8] The 5-year CDX index was chosen to represent the level of credit spreads. All four indices are highly correlated and the 5-year CDX index had the fewest missing observations.

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.[9]

Table 2 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 0.62 percentage points. Increasing the index spread by 1 basis point reduces the average base correlation by about 0.25 percentage points. The relationship between the log-*t* model parameters and the slope of the base correlation curve is more complex but is small in general. 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. A PARAMETRIC IMPLIED COPULA: HETEROGENEOUS VERSION

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

[9] 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, $t_{\nu,k}$ $(1 \leq k \leq n)$, 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

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

$q_k = 0.5(t_{\nu,k}+t_{\nu,k+1})$ and F is the cumulative probability distribution function for the t-distribution with ν degrees of freedom.

The hazard rate for the kth company is $\lambda_{k,i} = \exp(t_{\nu,k}\sigma_i+\mu_i)$. 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 if a three-parameter model is used, ν)
2. Find, for each company i, the value of μ_i that matches its CDS spread.
3. Determine hazard rates $\lambda_{k,i}$ $(1 \leq k \leq n)$ for the ith firm using $\lambda_{k,i}=\exp(t_{\nu,k}\alpha\sigma_i^*+\mu_i)$,
4. Use procedures in Andersen et al (2003) and Hull and White (2004) to value CDO tranches.
5. Search for values of α (and, if the three-parameter model is used, 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 either a different μ and or a different σ, or both. 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 and that the equity tranche spread remains constant as heterogeneity is introduced.

We continue to assume $\nu = 2.5$. 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 to match the predetermined index spread. 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 results are shown in Tables 3 and 4. 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 model 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 (s_1 = 0, all the firms have the same CDS spread) but the σ's differ from company to company. As before $\nu = 2.5$ and we consider index spreads of 50 and 200 basis points. The σ_i are drawn from a lognormal distribution 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 left-hand column of Tables 3 and 4. 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 2, this roughly corresponds to a range of base correlations of about 25%.[10] The results are shown in Tables 5 and 6 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 34.8%.[11] 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 5).

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. In the second case, $s_1 = 0.50$ 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 7 and 8. The results in Table 7 should be compared with the results in the second column of Table 3 ($s_1 = 0.25$) and the results in the third column of Table 5 ($s_2 = 0.20$). The results in Table 8 should be compared with the results in the third column of Table 3 ($s_1 = 0.50$) and the results in the fifth column of Table 5 ($s_2 = 0.40$). Table 7 shows that when $s_1 = 0.25$ and $s_2 = 0.20$ the correlation between the credit spread and σ,

[10] 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.

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

the default correlation, does not play an important role. The tranche spreads are close to the average of the corresponding spreads in Tables 3 and 5. Table 8 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. A TWO-FACTOR MODEL

The model we have proposed can be extended so that hazard rates are driven by more than one factor. Suppose that there are two portfolios. The hazard rates of the first portfolio have log-*t* distributions and are perfectly correlated; the hazard rates of the second portfolio have log-*t* distributions and are perfectly correlated. However, the hazard rates for the two portfolios are less than perfectly correlated. We are interested in determining that value of tranches on a portfolio which is the sum of the two individual portfolios. Correlation between the two portfolios is achieved by using a copula to mix the two distributions of hazard rates.[12]

The procedure is as follows:

1. Select *n* representative hazard rates for the first portfolio and *n* representative hazard rates for the second portfolio with their associated marginal probabilities, as described in Section III.
2. A Gaussian (or some other) copula is used to define the joint probability distribution of the hazard rates for the two portfolio. This results in an *n* by *n* table. Each cell in the table represents a hazard rate for the first portfolio, a hazard rate for the second portfolio, and the probability that these two hazard rates will occur together.
3. 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

[12] An alternative approach would be two specify two independent random factors t_1 and t_2 which are *t*-distributed. The hazard rate for any firm, *i*, is then $\ln(\lambda_i) = \mu_i + \sigma_{i1}t_1 + \sigma_{i2}t_2$. Correlation between sub-portfolios is determined by the relative magnitudes of σ_1 and σ_2.

determine the global present values of expected payoffs and payments. These are used to determine a spread for the bespoke.

When the portfolios are not homogeneous 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.

To provide a concrete example, we consider two homogeneous portfolios. The first portfolio contains 63 firms and the second 62 firms. Every name in each portfolio has a CDS spread of 50 basis points and a σ of 0.5. The combined portfolio contains 125 names all with the same credit spread and σ. We can consider this to be a portfolio with firms from two sectors or industries.

We use the two factor model to determine the spreads on tranches with the same attachment and detachment points as iTraxx. If the copula correlation used to mix the two distributions of hazard rates is 1 then the portfolio is the same as the case consider in the first column of Tables 3 and 5. The resulting tranche spreads for correlations between 0 and 1 are shown in Table 9. The equity (0–3%) and super-senior (22–100%) tranche spreads are most affected by the copula correlation. As the correlation rises from zero to one the equity tranche spread falls monotonically while the super-senior spread rises monotonically. For all other tranches the spread first rises and then falls.

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.

VI. OUT-OF-SAMPLE MODEL PERFORMANCE

One of the major applications of derivatives models is to use the prices of actively traded instruments to estimate prices of other less actively traded instruments. We have noted that when the two parameter version of the parametric implied copula model is calibrated to different instruments the parameter that determines default correlation, σ, is remarkably similar across instruments. This suggests that if the model is calibrated to one set of instruments it may produce good estimates for the prices of other instruments. To explore this we tried two tests.

The first test involved calibrating the two-parameter model ($\nu = 2.5$) to 5-year iTraxx quotes. On each day the resulting best-fit value of σ was then combined with the μ that fit the CDX 5-year index spread and the resulting parameters were used to calculate the model spreads for the CDX tranches. The calculated model spreads were compared to the actual CDX tranche spreads and the daily root mean square pricing error was calculated. We will refer to this daily root mean square pricing error as the out-of-sample root mean square pricing error, $RMSE_O$. These daily out-of-sample values of the RMSE are then compared with the corresponding in-sample root mean square pricing error, $RMSE_I$, which arises when the two-parameter model is calibrated to the CDX tranche data.

The in-sample RMSE is a measure of how well the 2-parameter model can fit the data. Over the entire sample period from September 28, 2005 to February 29, 2008 the average value of $RMSE_I$ is 4.1 basis points. For the two periods up to July 31, 2007 and after August 1, 2007 the corresponding average values of $RMSE_I$ are 2.8 basis points and 8.2 basis points respectively. As discussed earlier, the increase in $RMSE_I$ during the financial crisis is related to markedly higher spreads. In proportion to the average spreads the average $RMSE_I$ declines during the crisis.

The out-of-sample RMSE is always at least as large as $RMSE_I$. However, if the model is to be useful for extrapolating from one market to another, the out-of sample fit should not be materially worse than the in-sample fit. The average differences between $RMSE_O$ and $RMSE_I$ for the entire sample period and the two sub-periods, up to July 2007 and after July 2007, are 1.5, 0.3 and 5.3 basis points respectively. This indicates that the two-parameter model could have been used quite successfully to predict 5-year CDX tranche spreads from 5-year iTraxx spreads. This experiment was repeated calibrating the model to the 10-year iTraxx quotes and fitting the 10-year CDX quotes. Again the out-of-sample results are not materially worse than the in-sample results. The results are shown in Table 10.

It is not surprising that calibrating σ to iTraxx and then using the model to price CDX tranches works well. The results in Figure 6 suggest that this is likely to be the case. This led to our second test of out-of-sample pricing. In this test the model σ was determined by calibrating the model to all 5- and 10-year CDX and iTraxx investment grade tranche data. The calibrated σ was then used to determine the model quotes for the 5-year CDX HY tranches. The portfolio underlying the CDX HY tranches is a portfolio of 100 high yield names. The HY tranches are 0 to 10%, 10 to 15%, 15 to 25%, 25 to 35% and 35 to 100%. The quoted spreads for the first two tranches are an up front payment in percentage terms. The quotes for the three higher tranches are in basis points per year. [13]

Calibrating to investment grade tranches and pricing the HY tranches is a fairly strong test of the ability to price out-of-sample. The HY tranches differ substantially from the investment grade tranches both in the risk of the underlying portfolio and the tranche attachment points. The HY quotes are also somewhat more volatile than the IG quotes. In spite of this, calibrating to investment grade tranches (with $\nu = 2.5$) and pricing HY tranches works quite well. The results are in Table 10. For the period before the credit crisis the quality of out-of-sample fit to the HY data was comparable to that for the 5- and 10-year CDX data. (The tranche that is least well priced is the 35% to 100% tranche.)

VII. 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 parameter homogeneous model described previously to the index and the tranches of iTraxx Europe and then assume that the σ 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

[13] A typical set of quotes for the five tranches in June 2007 was 70%, 35%, 400bp, 125bp, and 25bp.

estimated for iTraxx Europe applies to the bespoke. In this case a different μ_i is calculated for each company underlying the bespoke.

In some cases it is appropriate to use two reference portfolios. 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. When both CDX NA IG and CDX NA HY are used for calibration, the values of σ that are appropriate for the bespoke can be 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 σ_{BE} are the σ-parameters of the log Student-*t* distribution that are fitted to CDX NA IG while CDX NA HY. The σ–parameter for the bespoke would be chosen as

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

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

When the bespoke portfolio includes both European and North American names the two-factor model outlined in Section V can be used.

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 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

VIII. 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 model is much easier. At first blush, it appears that the model provides a worse fit during credit crisis period that started in the summer of 2007 than it did before that. However, a closer examination shows that the average percentage fit of the model to tranche quotes is just as good, if not better, during the crisis as before the crisis.

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 is an attractive alternative to the one-factor Gaussian copula model. It is as easy to implement as that model and leads to a correlation parameter that is relatively constant across tranches and across portfolios on any given day. The model provides a natural approach to valuing bespokes. 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 companies in two sectors and a Gaussian copula model defines the correlation between the hazard rates of the two types of companies. The impact of the correlation between the sector 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.

Finally, one application of the model may be to the valuation of tranches of asset backed securities (ABSs) and ABS CDOs from ABX-HE indices.

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 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 = 2.5. *t*-statistics appear in parentheses.

Intercept	Sigma	Index (bp)	R^2	*n*
0.1865	0.6196	–0.0025	97.03%	63
(16.91***)	(39.96***)	(–20.70***)		

*** significant at 1%

Table 3
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 = 2.5.

	$s_1 = 0$	$s_1 = 0.25$	$s_1 = 0.50$	$s_1 = 0.75$	$s_1 = 1.00$
0-3%	34.8	34.8	34.8	34.8	34.8
3-6%	210	211	213	217	222
6-9%	73.5	73.8	74.8	76.5	78.7
6-12%	43.2	43.4	43.9	44.7	45.6
12-22%	23.4	23.4	23.5	23.6	23.7
22-100%	4.4	4.4	4.3	4.0	3.8

Table 4
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 = 2.5.

	$s_1 = 0$	$s_1 = 0.25$	$s_1 = 0.50$	$s_1 = 0.75$	$s_1 = 1.00$
0-3%	84.2	84.2	84.2	84.2	84.2
3-6%	2407	2413	2431	2463	2504
6-9%	1143	1147	1159	1175	1189
6-12%	574	577	583	593	599
12-22%	206	207	208	210	211
22-100%	18.5	18.3	17.5	16.7	15.9

Table 5
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 34.8%, and ν = 2.5.

	$s_2 = 0$	$s_2 = 0.10$	$s_2 = 0.20$	$s_2 = 0.30$	$s_2 = 0.40$
0-3%	34.8	34.8	34.8	34.8	34.8
3-6%	210	210	208	205	203
6-9%	73.5	73.5	73.8	74.6	75.8
6-12%	43.2	43.4	44.2	45.7	47.6
12-22%	23.4	23.5	23.9	24.7	25.9
22-100%	4.4	4.4	4.4	4.3	4.1

Table 6
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 84.2%, and ν = 2.5.

	$s_2 = 0$	$s_2 = 0.10$	$s_2 = 0.20$	$s_2 = 0.30$	$s_2 = 0.40$
0-3%	84.2	84.2	84.2	84.2	84.2
3-6%	2407	2402	2387	2364	2335
6-9%	1143	1140	1131	1115	1095
6-12%	574	573	571	567	561
12-22%	206	207	210	214	220
22-100%	18.5	18.6	18.8	19.3	19.8

Table 7
Model Quotes for Situation Where Both Spreads and Correlations Exhibit a 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 $\nu = 2.5$. β is the correlation between CDS spreads and firm σ's.

	β = –0.50	β = –0.25	β = 0.0	β = 0.25	β = 0.50
0-3%	34.8	34.8	34.8	34.8	34.8
3-6%	206.7	207.8	208.8	209.8	210.9
6-9%	73.1	73.6	74.2	74.7	75.2
6-12%	43.8	44.1	44.3	44.6	44.9
12-22%	23.8	23.9	23.9	24.0	24.1
22-100%	4.5	4.4	4.4	4.3	4.2

Table 8
Model Quotes for Situation Where Both Spreads and Correlations Exhibit a 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 $\nu = 2.5$. β is the correlation between CDS spreads and firm σ's.

	β = –0.50	β = –0.25	β = 0.0	β = 0.25	β = 0.50
0-3%	34.8	34.8	34.8	34.8	34.8
3-6%	198.1	202.4	206.7	210.8	215.0
6-9%	72.7	75.0	77.3	79.5	81.4
6-12%	45.5	46.7	47.8	48.8	49.7
12-22%	25.0	25.3	25.6	25.7	25.8
22-100%	4.6	4.3	4.0	3.7	3.5

Table 9
Model Quotes for a Portfolio Containing Two less than Perfectly Correlated Sub-Portfolios

The portfolio contains 125 names. Every name has a credit spread of 50 basis points and a σ of 0.5. The two sub-portfolios contain 63 and 62 names respectively. Results are shown for varying values of the copula correlation, ρ, between the two sub-portfolios. In all cases ν = 2.5.

	ρ = 0.0	ρ = 0.25	ρ = 0.50	ρ = 0.90	ρ = 1.0
0-3%	38.9	38.0	37.0	35.3	34.8
3-6%	210	216	218	214	210
6-9%	68.5	72.3	75.0	75.2	73.5
6-12%	43.7	44.8	45.3	43.8	43.2
12-22%	23.3	24.5	25.6	26.1	23.4
22-100%	1.5	1.7	2.1	3.5	4.4
Index	50.0	50.0	50.0	50.0	50.0

Table 10
In- and Out-of-Sample Model Fit

The table shows the average root mean square pricing error for three portfolios. The in-sample results are the average root mean square pricing errors when μ and σ in the two-parameter model (ν = 2.5) are chosen to minimize the RMSE. The out-of-sample results are the average root mean square pricing errors when σ is chosen to minimize the RMSE for some other set of tranche spreads and is then used to estimate the CDX portfolio spreads. The reference portfolio for 5-year CDX is 5-year iTraxx, for 10-year CDX it is 10-year iTraxx, and for CDX HY it is 5- and 10-year CDX and iTraxx. Note that the CDX HY data set starts in March 2006.

		5 Yr CDX	10 Yr CDX	5 Yr CDX HY
Sep05 to Feb08	In Sample	4.1	10.2	12.3
	Out of Sample	5.5	14.0	16.3
Sep05 to Jul07	In Sample	2.8	8.9	6.1
	Out of Sample	3.1	10.1	8.0
Aug07 to Feb08	In Sample	8.2	14.1	25.7
	Out of Sample	13.5	26.8	34.0

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

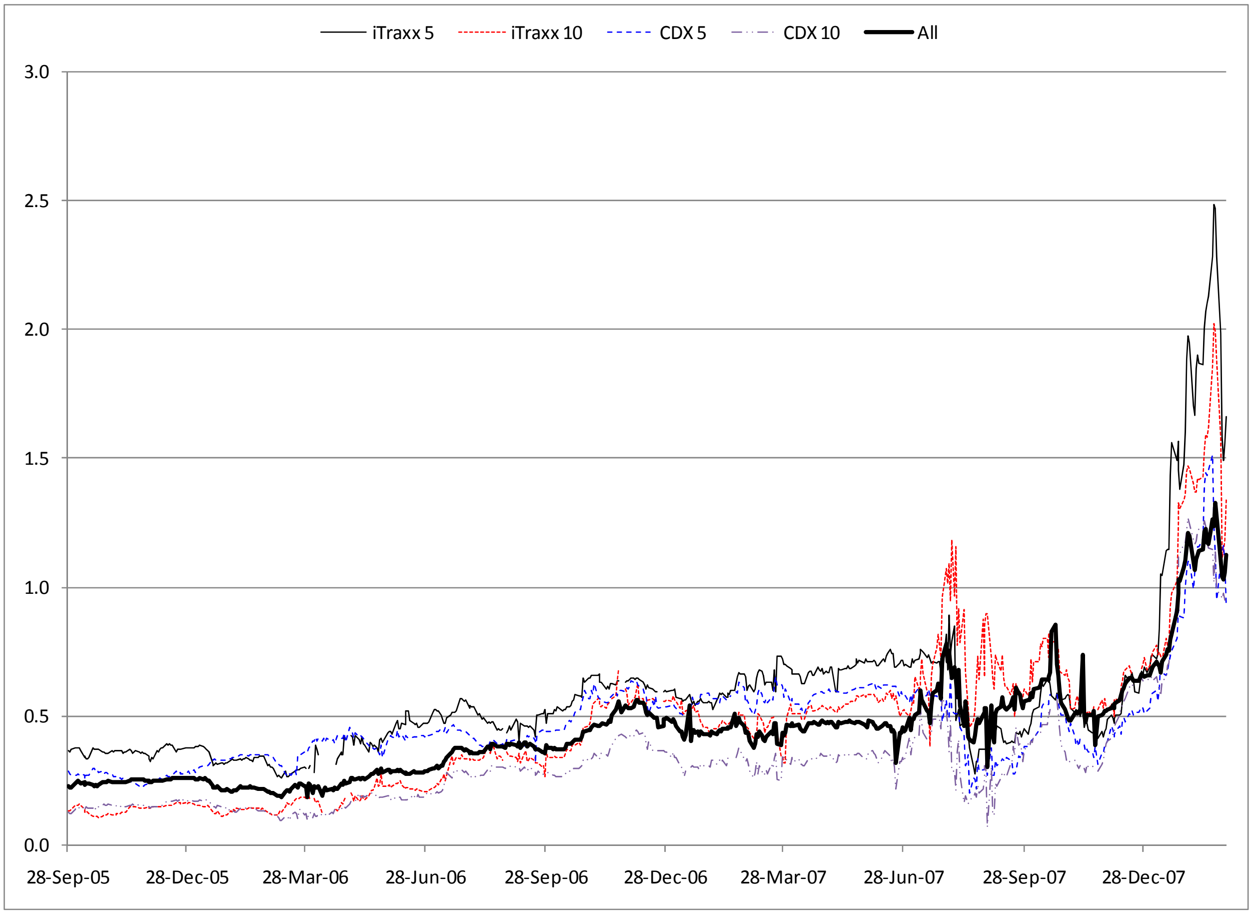

Figure 3
Value of ν when the Three-Parameter Model is Fitted to
5-year and 10-year iTraxx and CDX IG.

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

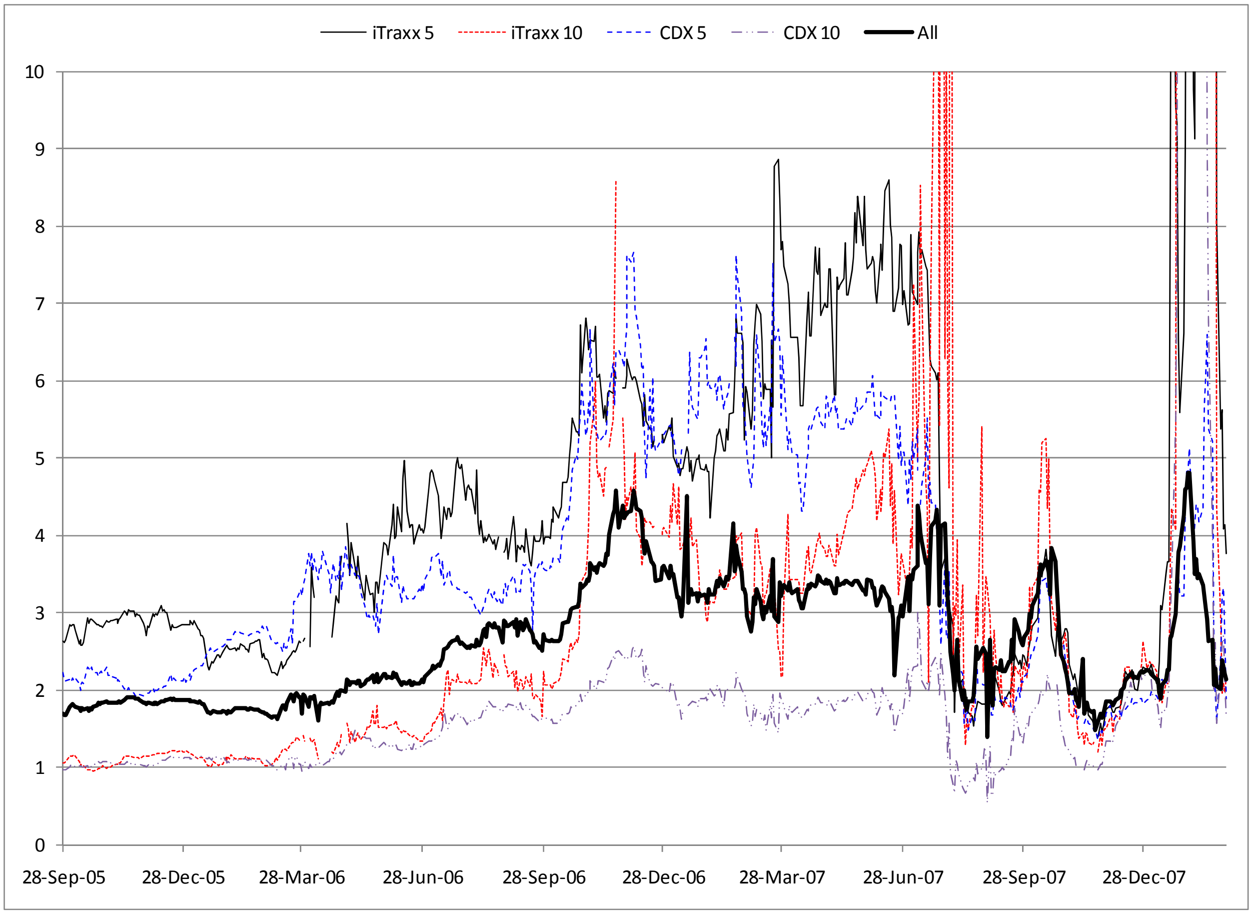

Figure 4
Value of Root Mean Square Error when the Three-Parameter Model is Fitted to 5-year and 10-year iTraxx and CDX IG.

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

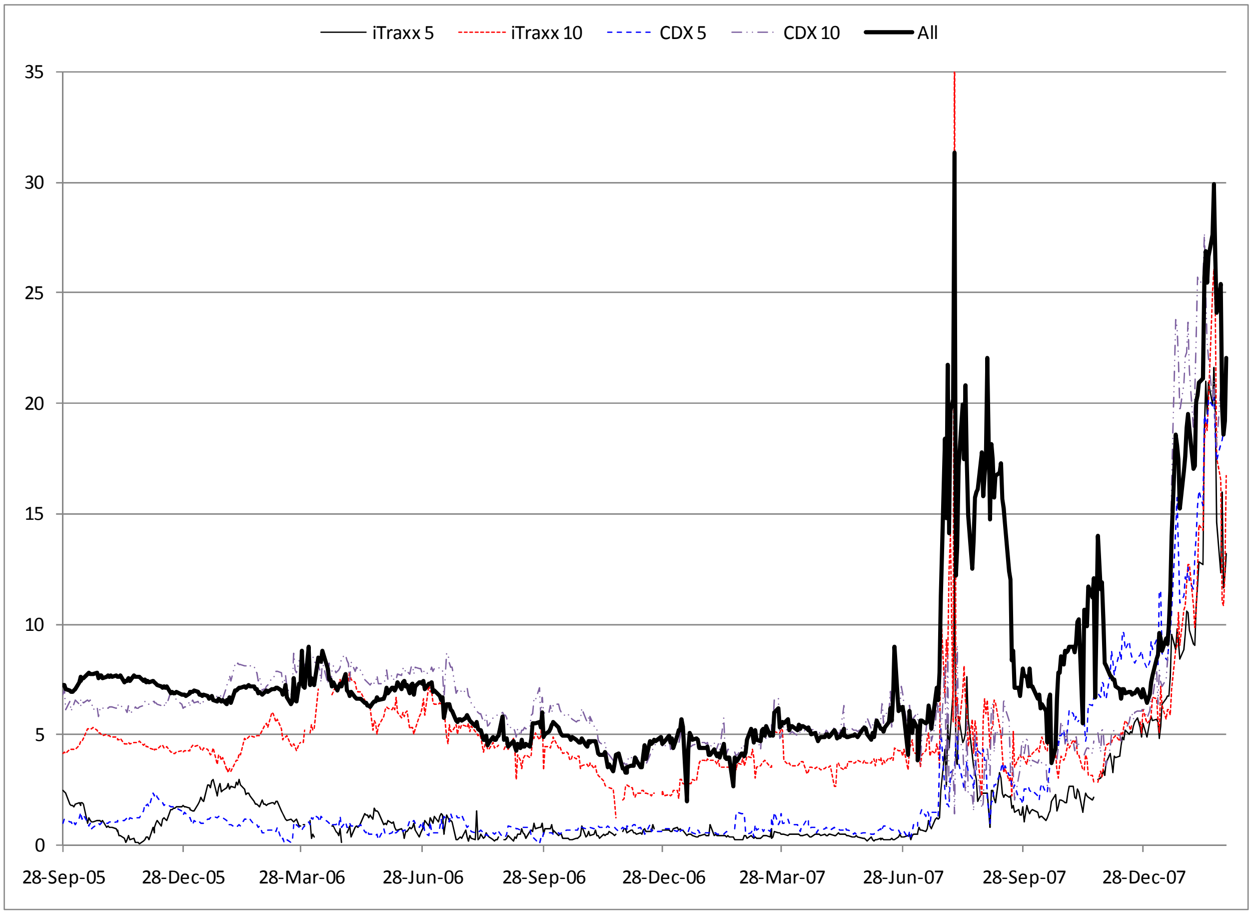

Figure 5
Value of μ when the Two-Parameter Model is Fitted to
5-year and 10-year iTraxx and CDX IG; ν=2.5

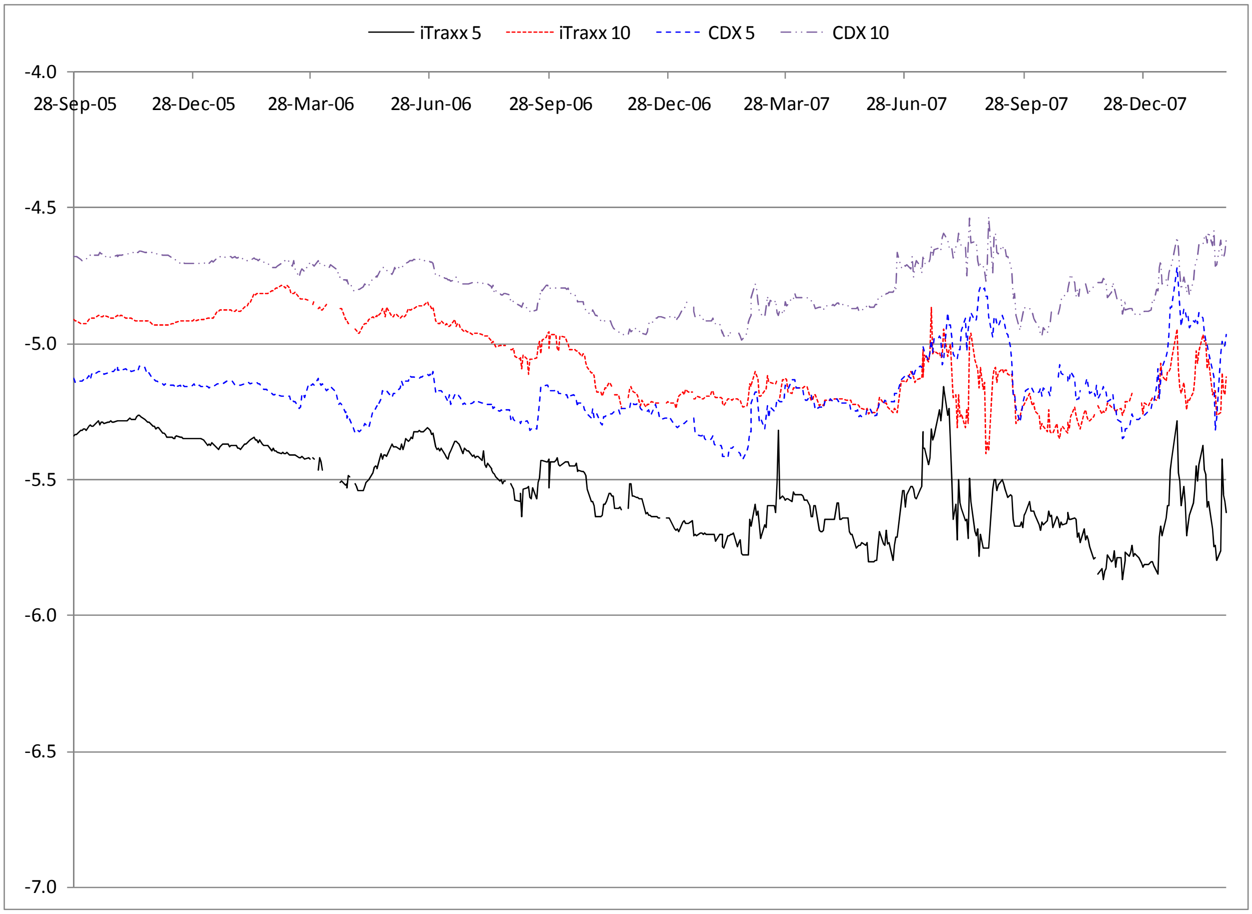

Figure 6

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

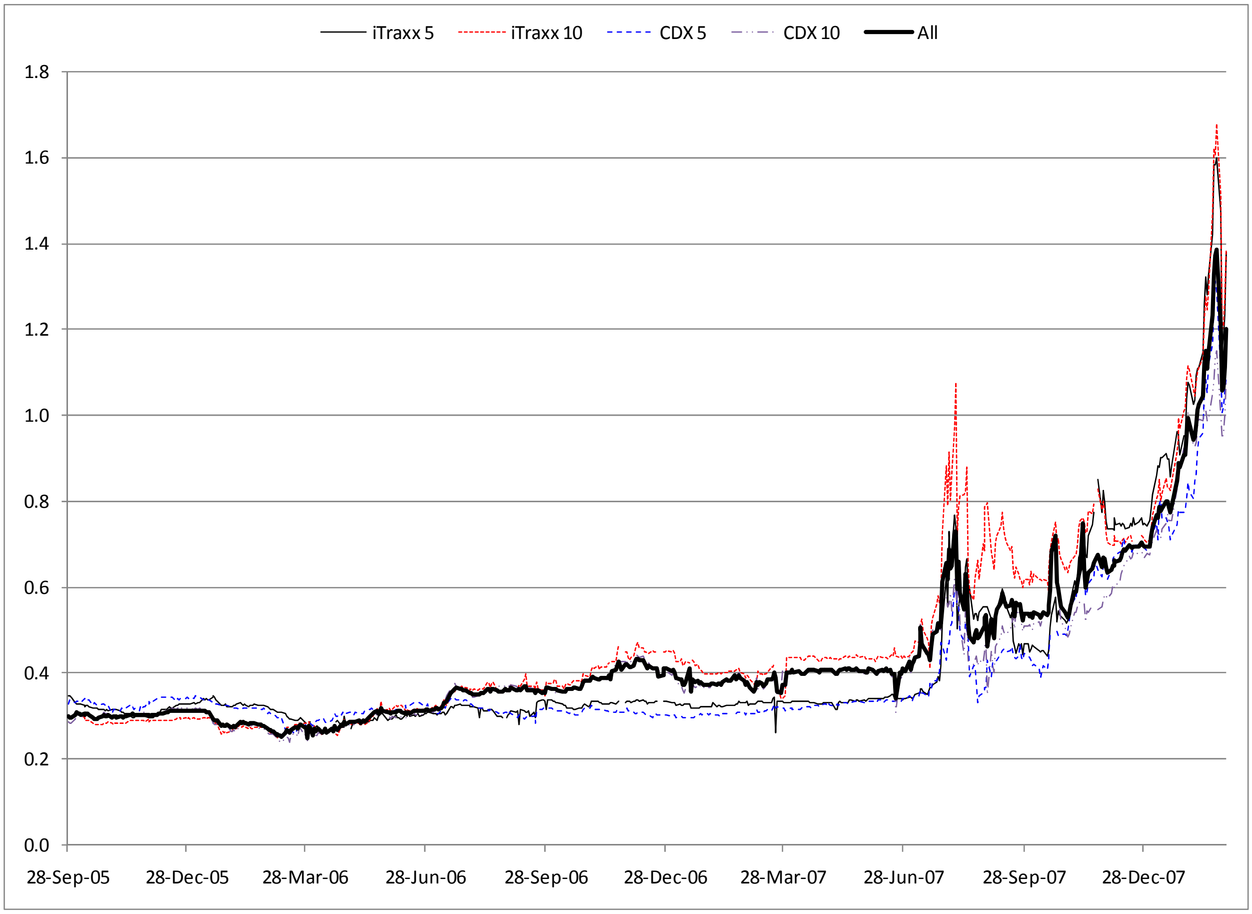

Figure 7

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