Journal of Derivatives, Fall 2006

Valuing Credit Derivatives Using an Implied Copula Approach

John Hull and Alan White*
Joseph L. Rotman School of Management

First Draft: June 2005

This Draft: November 2006

Abstract

We present an alternative to the Gaussian copula/base correlation model for valuing CDO tranches. Instead of implying copula correlations from market prices we imply the copula itself. Our model fits the market quotes for actively traded CDO tranches exactly. It is easy to understand and is a useful tool for pricing, trading, and risk management. It enables non-standard credit derivatives, such as bespoke CDOs and CDO squareds, to be priced consistently with market quotes for tranches of standard CDOs. Contrary to some of the criticisms that have been made of the approach, we find that it is more stable than the Gaussian copula/base correlation approach. Indeed our results suggest that the spreads given by the latter approach sometimes permit arbitrage.

* Early drafts of this paper were entitled "The Perfect Copula" and many analysts still refer to our approach as the perfect copula approach rather than as the implied copula approach. We are grateful to Bill Bobey for research assistance and to Leif Andersen, Steve Figlewski, Jon Gregory, David Lando, Roger Stein, and Harald Skarke for comments that have improved this paper. We are also grateful to participants at the following events for helpful feedback on earlier drafts of this paper: WBS Fixed Income Conference, Prague (Sept. 2005), GRETA Credit Risk Conference, Venice (Sept. 2005), Moody's Academic Advisory Committee Meeting (Nov. 2005), ICBI Global Derivatives, Geneva (Dec. 2005), WHU Campus-for-Finance Conference (Jan. 2006), Credit Derivatives Congress, New York (Apr. 2006), and an IAFE Presentation (May 2006). We are grateful for financial support from Moody's Investors Service.

Valuing Credit Derivatives Using an Implied Copula Approach

John Hull and Alan White

The Gaussian copula model has become the standard market model for valuing collateralized debt obligations and similar instruments. Many market participants like to imply what are known as base correlations for actively traded instruments using this model. Spreads for less actively traded instruments are often obtained by interpolating between these base correlations. In this respect the market uses implied base correlations and the Gaussian copula model in much the same way as it uses implied volatilities and the Black-Scholes model.[1]

If the Gaussian copula model fitted market prices well the implied base correlation would be approximately constant across tranches. However, this is not the case. As a result it is very difficult to determine the appropriate correlation when the Gaussian copula model is used to value non-standard credit derivatives such as bespoke CDOs and CDO-squareds.[2] This has led a number of researchers to look for copulas that fit market prices better than the Gaussian copula. Among the copulas that have been considered are the Student-*t*, double-*t*, Clayton, Archimedean, and Marshall Ohkin.

In this paper we take a different approach. We show how a copula model can be implied from market quotes. What we are doing in this paper is analogous to what Breeden and Litzenberger [1978] and Jackwerth and Rubinstein [1996] did when they implied a future stock price distribution from European option prices. Using a copula that provides a perfect fit to market quotes for all tranches is particularly useful when non-standard structures are considered.

The simplest version of the implied copula approach is the homogeneous case in which we assume that all companies being modeled have the same default probabilities and the same recovery rate. We choose a number of alternatives for the term structure of hazard

---

[1] For a description of the CDO market and the implementation of copula models see Hull and White [2004].

[2] Again we have an analogy with option pricing. Black-Scholes does not fit option prices well and so it is difficult to know the correct volatility to use when exotic options are valued using Black-Scholes assumptions.

rates and search for probabilities to apply to them so that market quotes are satisfied. We will refer to the result of this procedure as a set of hazard rate scenarios.

It is not immediately obvious that the implied copula approach, as we have just outlined it, has anything to do with copulas. However, we will show that the value of the factor in a one-factor copula model implies a hazard rate scenario for the underlying companies. The usual procedure for implementing a one-factor copula model involves integrating the value of the underlying instrument over the probability distribution of the factor using Gaussian quadrature. Equivalently it involves integrating the value of the instrument over the corresponding set of hazard rate scenarios. Specifying a one-factor copula model is therefore equivalent to specifying the corresponding set of hazard rate scenarios. Conversely, any set of hazard rate scenarios is equivalent to a copula.

Given this equivalence between a copula and a set of hazard rate scenarios, a natural suggestion is that, instead of specifying the copula in the usual way, we specify the set of hazard rate scenarios directly. This is the suggestion that underlies the ideas in this paper. The default correlation in the copula model corresponds to the dispersion of the hazard rate scenarios. As the dispersion increases the default correlation increases.

I. CDS AND CDO VALUATION

Credit default swaps (CDSs) and collateralized debt obligations (CDOs) provide protection against default in exchange for a fee. Although a CDO is a much more complicated contract than a CDS the approaches used to value the two contracts have some similarities.

A typical contract (CDS or CDO) has a life of 5 years during which the seller of protection receives periodic payments at some rate, s, times the outstanding notional. Usually these payments are made quarterly in arrears. When defaults occur[3] they reduce the outstanding notional principal and trigger two events. First there is an accrual payment to bring the periodic payments up to date. Second the seller of protection makes a payment equal to the loss to the buyer of protection. The loss is the reduction in the notional principal times one minus the recovery rate, R.

[3] For simplicity we assume that defaults occur in the middle of a period. More realistic assumptions can be used but they do not materially affect the value and serve only to obscure the notation.

The value of a contract is the present value of the expected cash flows. The calculation of this present value involves three terms. The first, *A*, is an annuity factor that is the present value of the regular payments at a rate of 1 per year. The second, *B*, is the accrual payment that occurs when defaults reduce the notional principal, and the third, *C*, is the payoff arising from defaults. Expressions for *A*, *B*, and *C* for a CDS are

$$
\begin{aligned}
A &= \sum_{i=1}^{n}\left(t_i - t_{i-1}\right)E\left[P\left(t_i\right)\right]e^{-rt_i}\\
B &= 0.5\sum_{i=1}^{n}\left(t_i - t_{i-1}\right)\left\{E\left[P\left(t_{i-1}\right)\right]-E\left[P\left(t_i\right)\right]\right\}e^{-r\left(t_i+t_{i-1}\right)/2} \qquad (1)\\
C &= \left(1-R\right)\sum_{i=1}^{n}\left\{E\left[P\left(t_{i-1}\right)\right]-E\left[P\left(t_i\right)\right]\right\}e^{-r\left(t_i+t_{i-1}\right)/2}
\end{aligned}
$$

where t_i for $1 \le i \le n$ is the end of period *i* when the i^{th} periodic payment is made, $P(t_i)$, is the notional principal outstanding at t_i, *r* is the risk-free rate of interest, and *E* denotes expectations taken over a risk-neutral density. The same equations apply to CDO tranches except that the expression for C should not be multiplied by 1–*R*. The total value of the contract to the seller of protection is $sA + sB - C$. The breakeven spread is $C/(A + B)$. The key element in pricing the contract is the determination of the expected notional principal at time t_i.

A CDS provides protection against a default by a particular firm. The variable $P(t)$ is initially equal to the notional principal, *K*. It reduces to zero when the firm defaults. The expected notional principal at time t_i is therefore

$$E\left[P\left(t_i\right)\right] = \left(1 - Q\left(t_i\right)\right)K$$

where $Q(t)$ is the risk-neutral probability that the entity has defaulted before time *t*.

A synthetic CDO provides protection against a subset of the total loss on a portfolio of CDSs. The portion of loss that is covered, known as a tranche, is defined by attachment point (AP), a_L, and detachment point (DP), a_H. The seller of protection agrees to cover all losses between $a_L K_{\text{Tot}}$ and $a_H K_{\text{Tot}}$ where K_{Tot} is the initial total notional principal of all the underlying CDSs. In exchange, the seller of protection receives payments at rate *s* on an initial notional $(a_H - a_L) K_{\text{Tot}}$. Each loss that is covered reduces the notional on which payments are based. Once the total portfolio losses exceed the detachment point no notional remains and all payments stop.

Two portfolios that attract a lot of trading are the CDX IG and the iTraxx portfolios. CDX IG is an equally weighted portfolio of 125 CDSs on investment grade North American companies. The standard {AP, DP} are {0, 3%), {3%, 7%}, {7%, 10%}, {10%, 15%}, {15%, 30%}, and {30%, 100%}. iTraxx is an equally weighted portfolio of 125 CDSs on investment grade European companies. The standard {AP, DP} are {0, 3%), {3%, 6%}, {6%, 9%}, {9%, 12%}, {12%, 22%}, and {22%, 100%}. The portfolios are revised periodically to reflect downgrades and defaults. Exhibit 1 shows the mid point of the bid and offer quotes from Reuters for CDX IG and iTraxx on August 30, 2005. The quotes for the 0 to 3% tranche show the upfront payment (as a percent of principal) that must be paid in addition to 500 basis points per year. The quotes for the other tranches are the annual payment rates in basis points per year that must be paid. The index quote indicates the cost of entering into a CDS on all 125 companies underlying the index. Consider for example the 50 basis point quote for the CDX IG index. This means that 125 five-year CDS contracts on the CDX IG companies, each with a notional principal of $800,000, could be purchased for a payment of 0.005×125×800,000 or $500,000 per year. This payment would reduce by 500,000/125 or $4,000 when a default occurs.

Suppose there are N CDSs in the CDO portfolio each with the same notional principal, K, and the same recovery rate, R. In this case the AP and DP can be mapped into the number of losses. A tranche with attachment point a_L and detachment point a_H is responsible for the n_L-th to n_H-th loss where $n_L = a_L N/(1–R)$ and $n_H = a_H N/(1–R)$. For example, if N = 125, $R = 0.4$, $a_L = 0.03$ and $a_H = 0.06$ then $n_L = 6.25$ and $n_H = 12.5$. In this case the initial notional for the CDO tranche is $(0.06–0.03)NK$ or $3.75K$. This is the remaining notional as long as there have been 6 or fewer defaults. The seller of protection is responsible for providing compensation for 75% of the seventh default so the seventh default reduces the notional by $(1–R)K×0.75$ or $0.45K$ to $3.30K$. Defaults 8 to 12 each reduce the remaining notional by $(1–R)K$ or $0.6K$. After the twelfth default the remaining notional is $0.30K$.

This is eliminated by the thirteenth default. Thereafter the notional is zero and no further payments occur.[4]

Let P_j be the remaining notional after the j^{th} default and define $m(x)$ as the smallest integer greater than x. Then the general expression for the remaining tranche notional after j defaults is

$$P_j(a_L, a_H) = \begin{cases} (a_H - a_L)NK & j < m(n_L) \\ a_H NK - j(1-R)K & m(n_L) \le j < m(n_H) \\ 0 & j \ge m(n_H) \end{cases} \tag{2}$$

The expected principal at time t_i is obtained by integrating this over the probability distribution of the number of defaults by time t_i.

II. ONE-FACTOR COPULA MODELS

A one-factor copula model is a way of modeling the joint defaults of n different obligors. The structure for this model was suggested by Vasicek [1987] and it was first applied to credit derivatives by Li [2000] and Gregory and Laurent [2005].

The first step is to define variables x_j ($1 \le j \le n$) by

$$x_j = a_j M + \sqrt{1 - a_j^2}\, Z_j \tag{3}$$

where M and the Z_j's have independent probability distributions with mean zero and standard deviation one. The variable x_j can be thought of as a default indicator variable for the j^{th} obligor: the lower the value of the variable, the earlier a default is likely to occur. Each x_j has two stochastic components. The first, M, is the same for all x_j while the second, Z_j, is an idiosyncratic component affecting only x_j. Equation (3) defines the correlation structure between the x_j's.

---

[4] A careful reader will notice that, if we have a set of CDO tranches that cover the full range of losses on the portfolio of CDSs, every time there is a default the total notional of the CDOs on which spread payments are based is reduced by $(1–R)K$. Meanwhile the total notional of the CDSs underlying the CDOs is reduced by K. Usually the notional of the most senior tranche, the tranche with the highest attachment point, a_L, is reduced by RK every time there is a loss that affects a more junior tranche and by K for each default affecting the most senior tranche. In this way the total notional for the CDOs remains the same as the total notional for the underlying CDSs.

Suppose that t_j is the time to default of the jth obligor and Q_j is the cumulative probability distribution of t_j. The copula model maps x_j to t_j on a "percentile to percentile" basis. The 5% point on the x_j distribution is mapped to the 5% point on the t_j distribution; the 10% point on the x_j distribution is mapped to the 10% point on the t_j distribution; and so on. In general, the point $t_j = t$ is mapped to $x_j = x$ where

$$x = F_j^{-1}\left[Q_j\left(t\right)\right] \tag{4}$$

or equivalently

$$t = Q_j^{-1}\left[F_j\left(x\right)\right]$$

and F_j is the cumulative probability distribution for x_j.

The copula model defines a correlation structure between the t_j's while maintaining their marginal distributions. The essence of the copula model is that we do not define the correlation structure between the variables of interest directly. We map the variables of interest into other more manageable variables (the x_j's) and define a correlation structure between those variables.

From Equation (3)

$$\text{Prob}\left(x_j < x \middle| M\right) = H_j\left[\frac{x - a_j M}{\sqrt{1-a_j^2}}\right]$$

where H_j is the cumulative probability distribution of Z_j. It follows from Equation (4) that

$$Q_j\left(t \middle| M\right) = \text{Prob}\left(t_j < t \middle| M\right) = H_j\left\{\frac{F_j^{-1}\left[Q_j\left(t\right)\right] - a_j M}{\sqrt{1-a_j^2}}\right\} \tag{5}$$

Conditional on M, defaults are independent. When using the model to value a CDO tranche we set up a procedure to calculate, as indicated in Equation (1), the components of the present value of expected cash flows on a tranche conditional on M and then integrate over M to obtain their unconditional values.

The Standard Market Model

In the standard market model M and the Z_j have standard normal distributions and all the a_j are equal. The time to default, $Q_j(t)$, is the same for all j and is usually determined by assuming a constant hazard rate that matches the CDS spread for the index. The recovery rate is assumed to be constant at 40%. The only free parameter in the model is therefore a, the common value of the a_j.

III. THE IMPLIED COPULA APPROACH

Equation (5) together with the probability distribution for M defines a probability distribution for the term structure of cumulative default rates. Define $\lambda_j(t|M)$ as the hazard rate at time t conditional on M for company j. The relationship between $\lambda_j(t|M)$ and $Q_j(t|M)$ is

$$Q_j(t|M) = 1 - \exp\left[-\int_0^t \lambda_j(\tau|M)\,d\tau\right]$$

or equivalently

$$\lambda_j(t|M) = \frac{dQ_j(t|M)/dt}{1 - Q_j(t|M)}$$

Equation (5) together with the probability distribution of M can therefore also be used to define a set of hazard rate scenarios. Furthermore, the set of hazard rate scenarios totally defines the model.

This suggests a new way of viewing factor-based copulas. Instead of formulating the model in terms of Equations (3) and (4) we can go directly to Equation (5) and formulate the model in terms of a set of hazard rate scenarios. It is clear that there is a set of hazard rate scenarios corresponding to any factor-based copula. From Sklar's theorem there must also be a copula corresponding to any set of hazard rate scenarios.

The implied copula approach involves implying a set of hazard rate scenarios from market quotes. We choose a set of hazard rate term structures. (The way this choice is made will be discussed later). We then search for probabilities to apply to the hazard rate term structures so that market quotes are matched.

The mechanics of implementing the implied copula model are quite similar to the mechanics of implementing a factor-based copula. In a factor-based copula we discretize the distribution of M and each point of the discrete distribution implies through Equation (5) a term structure for cumulative default rates. In the implied copula approach we choose a discrete set of hazard rate term structures. Each hazard rate term structure implies a term structure for cumulative default rates. In a factor-based copula the probabilities to apply to the term structure of cumulative default rates are based on the probability density for M. In the implied copula approach we imply the probabilities.

Correlation in the Implied Copula Model

Before describing how the hazard rate term structures are chosen and how the probabilities assigned to them are implied we make an important point about the relationship between default correlation and a set of hazard rate scenarios.

Default correlation arises in a factor-based copula model because of uncertainty about the default environment. When the model is expressed using Equation (3) it is the value of M that defines the default environment. Low values of M correspond to bad default environments while high values of M correspond to good default environments. When the model is expressed in terms of a set of hazard rate scenarios, high hazard rates correspond to bad default environments while low hazard rates correspond to good default environments.

Default correlation depends on the dispersion of the set of hazard rate scenarios. As the dispersion increases default correlation increases. This is because the hazard rate term structure sampled from the set of hazard rate scenarios applies to all companies in the portfolio. In the extreme situation where there is only one hazard rate term structure the default correlation is zero. In this case knowing the time to default for one company does not help in determining that for other companies. As the dispersion of the set of hazard rate scenarios increases this becomes progressively less true.

To express this point algebraically suppose that d_k is the cumulative probability of default over the life of the contract for an obligor for the kth hazard rate term structure. The probability that a company will default during the life of the contract is

$$q = \sum_k \pi_k d_k$$

where π_k is the probability of the kth hazard rate term structure. The probability that any two companies will both default during the life of the contract is

$$\sum_k \pi_k d_k^2$$

This increases as the variance of cumulative default rates increases. The latter increases as the dispersion of the set of hazard rate scenarios increases.
The probability that two companies will both default in the Gaussian copula model is

$$B(b, b, \rho)$$

where $b = N^{-1}(q)$, N is the cumulative normal distribution function, ρ is the copula correlation, and $B(x, y, \rho)$ is the cumulative probability that the first variable is less than x and the second variable is less than y in the bivariate normal distribution where the correlation is ρ. The Gaussian copula correlation for any two companies that is equivalent to a particular implied copula is given by finding the ρ that satisfies

$$\sum_k \pi_k d_k^2 = B(b, b, \rho) \tag{6}$$

This Equation shows that copula correlation can be regarded as a measure of the variance of the default rate distribution.

IV. IMPLEMENTATION OF MODEL

In the simplest implementation of the implied copula approach default probabilities are determined by assuming that defaults occur according to a Poisson process whose hazard rate is drawn from a discrete L-point distribution. Our conditional cumulative default probabilities are

$$Q(t|\lambda_k) = 1 - \exp(-\lambda_k t) \quad k = 1,\ldots,L \tag{7}$$

The probability that $\lambda = \lambda_k$ is π_k. The calibration problem is to choose a set of λ_k's and their associated π_k's.

The index quote and the quote on the first five iTraxx or CDX tranches are used for calibration. We first choose a set of λ_k's. For each λ_k, Equations (1) and (7) are used to calculate the value of the five tranches and the index. Let us denote these values as $V_m(\lambda_k)$ for $m = 1$ to 6, and $k = 1$ to L. The market is perfectly matched if we can choose the π_k's so that

$$\begin{aligned} &\sum_{k=1}^{L} \pi_k V_m(\lambda_k) = 0 \quad && m = 1,\ldots,6 \\ &\sum_{k=1}^{L} \pi_k = 1 && \\ &\pi_k \geq 0 && k = 1,\ldots,L \end{aligned} \tag{8}$$

If the recovery rate is constant we find that there is often no solution to Equation (8) when market spreads are used. That is, if recovery rates are assumed to be constant there is no distribution of default probabilities that is exactly consistent with observed prices. However, when a recovery rate model that reflects the negative correlation between default rates and recovery rates is used the prices can be fit. This negative correlation has been documented by Altman *et al* [2002], and Cantor *et al* [2002], and Hamilton *et al* [2005]. The best fit relationship reported by Hamilton *et al* [2005] is

$$R(Q) = \max\left[0.52 - 6.9 \times Q(1), 0\right]$$

Using this relationship does not make the analysis significantly more complicated. The only change is that a different recovery rate is used in the calculations for each λ_k. This applies to both CDOs and CDSs. In the interests of consistency we have used this recovery rate model in all the results we report for both the Gaussian copula model and the implied copula model.

Choosing the λ's

To implement the implied copula approach we must choose the number of λ's that will be used, L. We then distribute these between zero (no chance of default) and some high number λ_{max} (probability of default over a period of 5 years of close to one).

Define the sum of the values of the instruments for a hazard rate of zero and a very high hazard rate as follows

$$V_{\min} = \sum_{m=1}^{6} V_m(0)$$
$$V_{\max} = \sum_{m=1}^{6} V_m(\lambda_{\max})$$

We choose values of λ_k ($1 \le k \le L$) so that

$$\sum_{m=1}^{6} V_m(\lambda_k) = V_{\min} + (k-1)(V_{\max} - V_{\min})/(L-1)$$

This ensures that the total values are evenly distributed between V_{min} and V_{max}.[5]

Choosing the π's

When $L > 7$ there may be many solutions to Equation (8). The easiest procedure for examining the set of possible solutions is to treat this as a linear programming problem in which we maximize

$$\max \sum_{k=1}^{L} w_k \pi_k$$

for some set of weights, w_k subject to the constraints given in Equation (8). This is solved using the Simplex method. By varying the weights, w_k, all possible solutions to the problem can be found. In practice we consider the L solutions defined by $\{w_k = 1; w_j = 0 \text{ for } j \neq k\}$ for k =1 to L. These solutions reveal the maximum value that

[5] A simple alternative approach is to distribute the λ's logarithmically

$$\ln(\lambda_j) = \ln(\varepsilon) + \left(\frac{\ln(0.5) - \ln(\varepsilon)}{L-1}\right)$$

for some suitably small value of ε.

each individual π_k can take on. We then choose the convex combination of these solutions that minimizes[6]

$$\sum_{k=2}^{L-1} \frac{\left(\pi_{k-1} + \pi_{k+1} - 2\pi_k\right)^2}{\left(\lambda_{k+1} - \lambda_{k-1}\right)}$$

This produces a maximally smooth distribution of probabilities. Exhibit 2 shows the results of fitting the implied copula model to the five-year iTraxx and CDX data shown in Exhibit 1. The convergence of the procedure as the value of L is increased is illustrated in Exhibit 3.

The probability distributions in Exhibits 2 and 3 are truncated. The implied copula approach assigns a small probability to extreme hazard rate term structures that lead to a cumulative default rate for the life of the model greater than 70%. We refer these extreme hazard rate term structures as "end-of-the-world" scenarios. Without assigning a small probability to these scenarios it is impossible to fit the spread for the most senior tranche.[7]

iTraxx and CDX Time Series Results

Exhibit 4 shows the results of fitting the implied copula model to daily five-year iTraxx and CDX data between July 26, 2004 and November 2, 2005. The model was calibrated every day for which we had an index quote plus three or more CDO tranche quotes (out of a possible 5).[8] To facilitate comparisons the set of λ's were chosen to be the same each day.

The main thing to note in Exhibit 4 is the structural change in the probability distribution of hazard rates that took place in May 2005 when Ford and GM were downgraded to non-investment grade. Prior to this there was a significant probability assigned to a hazard rate of zero. After this the probability that the hazard rate was zero (i.e., that there was no chance of default) was itself zero. This phenomenon is observed for both indices.

---

[6] This sum is the discrete analog of twice the integral of the square of the second derivative.

[7] This phenomenon will be familiar to anyone who has implemented the Gaussian copula. Even though they have a very low probability, it is necessary to consider extreme values of the probability distribution for M in order to calculate accurate spreads for senior tranches.

[8] It is not necessary to have a full data set to calibrate the model.

Obtaining Bounds for Tranches

As we will illustrate, the model we have outlined can provide an estimate of the price of non-standard instruments. It can also be used to obtain bounds for a price.[9] Suppose that $V(\lambda)$ is the value of a particular instrument when the hazard rate is λ. The maximum (minimum) value of the instrument is given by using linear programming to choose the π_k to maximize (minimize)

$$\sum_{k=1}^{L} \pi_k V\left(\lambda_k\right)$$

subject to the constraint that the selected probabilities satisfy the calibrating equations in Equation (8).

For a CDO tranche the tranche value is a function of the tranche spread, s. For any spread, s, and any hazard rate, λ, the value of the tranche is $V(s, \lambda)$. Given s, the upper and lower bounds on the tranche value are

$$V_{\max}\left(s\right) = \max \sum_{k=1}^{L} \pi_k V\left(s, \lambda_k\right) \quad \text{and} \quad V_{\min}\left(s\right) = \min \sum_{k=1}^{L} \pi_k V\left(s, \lambda_k\right)$$

subject to the constraint that the selected probabilities satisfy the calibrating equations. The maximum and minimum permissible (no arbitrage) spreads are the highest and lowest spreads for which $V_{\min}(s) \leq 0 \leq V_{\max}(s)$.

The normal market practice is to imply a spread for the super senior tranche (30% to 100% in the case of CDX and 22% to 100% in the case of iTraxx) from the spreads for other tranches and the index. The rationale for this is that buying a portfolio of the super-senior tranche plus the other five tranche is similar to buying the index.[10] To test this we used the approach just presented to calculate bounds for the super-senior spread. Our results confirm that the market practice is reasonable. The bounds appear to be quite tight. For example, using the five-year iTraxx data from August 30, 2005 we find that the spread for the 22% to 100% tranche must lie between 4.92 and 4.97 basis points.

[9] The issue of providing bounds on prices is also considered by Walker [2006]

[10] The payoffs in the event of default are the same for the two portfolios. However, the pattern of payments is quite different.

V. COMPARISON OF THE BASE CORRELATION AND IMPLIED COPULA APPROACHES

As already mentioned the standard market model is a one-factor Gaussian copula model. There is one free parameter. This is a, the common value of a_j. The copula correlation in the model is a^2.

Copula correlations can be implied from the spreads quoted in the market for particular tranches. These correlations are known as tranche correlations or compound correlations. Exhibit 5 shows implied tranche correlations for the five-year quotes in Exhibit 1. It can be seen that the tranche correlations were not constant on April 30, 2005. They exhibit a pronounced 'smile'. This is the usual situation.

An alternative and more popular correlation measure is known as the base correlation. Using Equation (1) define $C(a_L, a_H, \rho)$ as the expected loss for a tranche with attachment point a_L, detachment point a_H, and correlation ρ. This can be calculated for each tranche using the tranche implied correlation. The tranches are numbered so that the attachment point for tranche m is the detachment point for tranche m–1. The ν^{th} base correlation is chosen so that the total expected loss is correct for the first ν tranches. This means that

$$C\left(0, a_{H\nu}, \rho_{B\nu}\right) = \sum_{m=1}^{\nu} C\left(a_{Lm}, a_{Hm}, \rho_m\right)$$

where $\rho_{B\nu}$ is the base correlation for the ν^{th} tranche and ρ_m is the m^{th} tranche correlation. Base correlations are usually monotonically increasing with the number of tranches included, as illustrated in Exhibit 5.

We now compare the use of base correlations with our implied copula approach.

Non-Standard Attachment Points

The usual procedure for calculating a spread for nonstandard a_L and a_H is as follows. The base correlations for $\{0, a_L\}$ and $\{0, a_H\}$ are calculated by interpolating between the base correlations for the neighboring points for which data is available. The expected loss for the tranche is estimated as the expected loss for the $\{0, a_H\}$ tranche minus the expected loss for the $\{0, a_L\}$ tranche. The $\{a_L, a_H\}$ tranche correlation is implied from this expected loss and the breakeven tranche spread is calculated using this tranche correlation.

The implied copula approach can be used for a nonstandard tranche in two different ways. First, the tranche spread can be calculated directly from the λ_k's and π_k's. Second, the bounding approach described above can be used to find the highest and lowest possible tranche spreads.

The results from using the implied copula approach to calculate the spread directly for iTraxx and CDX data on August 30, 2005 are shown in Exhibit 6. They are compared with two different implementations of the base correlation approach, one using linear interpolation, the other using a cubic-spline interpolation scheme. In all cases the tranche spread is determined and then converted into an implied tranche correlation using the Gaussian copula model. The tranches considered have {AP, DP} equal to {4%, 5%}, {5%, 6%}, {6%, 7%}, … The exhibit shows that the interpolation scheme does have a noticeable effect on the base correlation results. Also, the spreads from using the implied copula approach are smoother and more stable across tranches than those from the base correlation approach.

The results from using the bounding approach for iTraxx and CDX data on August 30, 2005 are shown in Exhibit 7. The tranches considered are the same as in Exhibit 6. The results from the two implementations of the base correlation approach are compared with the bounds given by the implied copula approach. The results show that when linear interpolation is used in the base correlation approach the tranche correlations are regularly outside the bounds. When the spline interpolation scheme is used in the base correlation approach the tranche correlations are better behaved but occasionally fall marginally outside the bounds for iTraxx and are far outside the bounds for CDX.

The market's spreads for the non-standard tranches we consider are not known but it is clear that the results from the implied copula approach are superior to those from the base correlation approach. The implied copula results are smoother and more stable across tranches than the base correlation results. Many of the base correlation results lie outside the permissible range for the set of copula models that are implicitly considered by the implied copula approach.

One approach that is sometimes used to avoid the problems with the base correlation approach that are illustrated in Exhibits 6 and 7 is to calculate expected losses for the 0 to *X*% tranche. Market quotes enable this to be estimated for *X* equal to 3, 6, 9, 12, and 22 in the case of iTraxx and *X* equal to 3, 7, 10, 15, and 30 in the case of CDX. The expected losses for other values of *X* are obtained by interpolation. Interpolating expected losses is more direct than interpolating base correlation and gives much better results. It can be shown that for no arbitrage the expected loss for the 0 to *X*% tranche must be an increasing function of *X* with a negative second derivative. It is not difficult to ensure that the interpolation scheme satisfies this requirement.

The implied copula approach does not have any advantages over the interpolate-expected-loss approach for valuing non-standard tranches. But it is much easier to use than that approach for many of the non-standard structures that are encountered.

Non-Standard Number of Names

We now test the effect of changing the number of companies in the iTraxx portfolio. We consider a {3%, 6%} (mezzanine) five-year tranche. We are assuming homogeneity. The names in our portfolio therefore are assumed to have exactly the same characteristics as the average iTraxx name. Three models are considered:

1. A model where base correlations for the 3% and 6% detachment points are assumed to be the same as those for the regular 125-company iTraxx portfolio
2. A model where the {3%, 6%}tranche correlation is assumed to be the same as the {3%, 6%} tranche correlation for the regular 125-company iTraxx portfolio
3. The implied copula model, fitted to the data in Exhibit 1

The results are shown in Exhibit 8. The three models agree that the tranche spread decreases quite sharply as the number of companies increase. As far as we know, this result is not widely known. It arises from the properties of the binomial distribution. Consider the case where the default probability is 0.03. The probability that the number of defaults is greater than 5% and less than or equal to 10% is 11.1%, 9.2%, and 7.0% for portfolios of size 40, 80, and 120 companies respectively.

Exhibit 8 shows that the results obtained from the Gaussian copula model are critically dependent on how the model is used. Two plausible ways of using the model give very different results. Again we conclude that the implied copula model is more stable across tranches than the Gaussian copula model.

VI. EXTENSIONS OF THE IMPLIED COPULA APPROACH

The implied copula approach as we have described it so far assumes that all companies have the same constant default rates. The life of the CDO being considered is assumed to be the same as that of the reference (iTraxx or CDX) CDOs. We now consider some ways in which the implied copula approach can be extended.

Different Maturities

Exhibit 9 shows the cumulative probability distribution for the 5-year and 10-year hazard rate for iTraxx on August 30, 2005. It is natural to suggest that the implied copula methodology be extended to find a 10-year hazard rate path distribution that simultaneously matches the market data for both five and ten year maturities. It is possible to find such a hazard rate path distribution, but it is important to avoid over fitting the model. The model is a description of the average hazard rate environment between time zero and some time T as seen at time zero. The model does not say anything about the dynamics of hazard rates. A different type of model is needed to answer a question such as "if the hazard rate between now and year 5 is λ, what is the probability distribution for hazard rates between years 5 and 10?"

When valuing CDO tranches and other similar instruments that have maturities between five and ten years, it makes sense to interpolate between the terminal default rates at each quantile of the cumulative distributions. For example, in the case of iTraxx 90% of the five- and ten-year default rates are below 0.0446 and 0.0975 respectively. From this we estimate that 90% of the seven-year default rates are less than or equal 0.0658. Similar calculations can be carried out for every quantile to generate a complete probability distribution for seven-year default rate.

Different CDS Spreads

We now discuss how the model we have presented can be extended so that the homogeneity assumption is relaxed. In our experience a non-homogeneous model gives similar results to the corresponding homogeneous model in most circumstances. However, a non-homogeneous model can be important for risk management.

At this stage we continue to assume flat hazard rate term structures. Continuing with our earlier notation λ_k ($1 \le k \le L$) is the value of the kth hazard rate in the homogeneous model. We define:

λ_{kj} ($1 \le k \le L$): the kth hazard rate for the jth company (to be determined)

s_j ($1 \le j \le N$): the CDS spread for the jth company

s_{index}: the CDS spread for the index (This is assumed to be the CDS spread for all companies in the homogeneous case)

$U(\lambda,s)$: the value of a CDS to buy protection on a company on when the principal is \$1, hazard rate is λ and the CDS spread is s. (From Equation (1) this is $C(\lambda) - [A(\lambda)+B(\lambda)]s$.)

We choose the λ_{kj} so that

$$\frac{U\left(\lambda_{kj}, s_j\right)}{U\left(\lambda_k, s_{index}\right)} = c_j \tag{9}$$

for all k where c_j is a constant independent of k. The c_j cannot be chosen arbitrarily. In general there is a small range of values of c_j for which Equation (9) can be satisfied. We then choose the π_k so that a) the five CDO tranche quotes are matched in the non-homogeneous model and b) s_{index} is matched in the homogeneous model. Matching s_{index} in the homogeneous model ensures that

$$\sum_{k=1}^{L} \pi_k U\left(\lambda_k, s_{index}\right) = 0$$

From Equation (9) this implies that

$$\sum_{k=1}^{L} \pi_k U\left(\lambda_{kj}, s_j\right) = 0$$

so that all N CDS spreads are matched.

Implementing the non-homogeneous model is therefore very similar to implementing the homogeneous model. The only differences are a) it is necessary to develop the procedure for calculating the c_j's and λ_{kj}'s and b) it is necessary to use a algorithm such as that in Andersen et al [2003] or Hull and White [2004] to calculate loss distributions on a non-homogeneous portfolio.

Matching the CDS Term Structure

The implied copula model can be extended so that the term structure of CDS spreads is matched. This is achieved by allowing the alternative hazard rate term structures to be non-flat. We find that for most applications of the model matching the CDS term structure makes very little difference. Roughly speaking, it changes the times when defaults occur without changing the cumulative number of defaults.

Suppose that the CDS spread is known for company j for maturities t_1, t_2, t_3, …, t_n. We first calculate a set of constant hazard rates for each company as just described in the previous section to match the t_n maturity spread. Define d_{kj} as the cumulative default rate for the kth hazard rate (= $1-\exp(-\lambda_{kj}t_n)$ where λ_{kj} is the kth hazard rate for company j).

To match all CDS spreads for company j we adjust the model so that the kth hazard rate for the company is constant between time t_{i-1} and time t_i (t_0=0). We choose the hazard rates so that the ratio of the value of the t_i year CDS to the value of the t_n year CDS is independent of k while ensuring that the kth hazard rate term structure is consistent with d_{kj}. This means that when we match the t_n maturity spread we automatically match the t_i maturity spread for $1 \leq i \leq n-1$. The ratio of the value of the t_i year CDS to the value of the t_n year CDS can be chosen as $(w_i s_i)/(w_n s_n)$ where w_i is the present value of payments at the rate of \$1 per year on a t_i year CDS when the hazard rate is zero and s_i is the t_i year CDS spread.

There is perfect correlation between obligors in the generalized model we have outlined in that when we know the hazard rate term structure for one obligor we know the hazard rate term structure for all other obligors. Searching for the optimal set of probabilities in the model is no more difficult than in the homogeneous model because the model is set up so that when we match the CDS spread for the index we automatically match the term structure of CDS spreads for all companies underlying the index. As already mentioned an algorithm such as that in Andersen et al [2003] or Hull and White [2004] must be used to calculate loss distributions on the portfolio for each i, but such an algorithm must similarly be used for each value of the factor that is considered in the nonhomogeneous Gaussian copula model.

It is possible to extend the model still further so that the dispersion of hazard rate paths is different for different companies. However, just as it is difficult to know how to choose unequal correlations in the Gaussian copula model, so it is difficult to know how to vary the dispersion of hazard rates from company to company in the implied copula model.

VII. BESPOKE PORTFOLIOS

CDOs involving bespoke portfolios can be handled using the implied copula approach. Consider first the situation where the bespoke portfolio is considered to be as well diversified as the portfolio underlying the reference index (iTraxx or CDX). The simplest approach is to first calibrate the model to the index assuming homogeneity and then adjust the hazard rates. Suppose that the index gives a probability π_k of a hazard rate λ_k. For the bespoke portfolio we assume that there is a probability π_k of a hazard rate λ_k^* where

$$\lambda_k^* = \beta\lambda_k$$

The constant β is chosen so that the average CDS spread for the companies in the bespoke portfolio is matched.

A more sophisticated approach is to use the extension of the implied copula model where each company has a different set of hazard rates. This works exactly as described Section VI. The model is set up so that when we match the index CDS quote in the homogeneous model we automatically match all CDS spreads for the companies in the bespoke portfolio.

Dealing with portfolios that are less (or more) well diversified than the index requires some judgment and, whether the Gaussian copula/base correlation or implied copula approach is used, is inevitably somewhat *ad hoc*. The theory underlying structural models of credit risk suggests that the Gaussian copula correlation should be similar to the correlation between equity returns. Based on this observation one way of proceeding is as follows. Calculate, y, the average pairwise correlation between equity returns for companies in the portfolio and y^* the average pairwise correlation between equity returns for companies in the index. Assume that the average pairwise Gaussian copula correlation for the bespoke portfolio is $\rho y/y^*$ where ρ is its value for the index. Increase the dispersion of hazard rate paths for each obligor while maintaining the correct CDS spread so that the results are consistent with Equation (6) when ρ is replaced by $\rho y/y^*$.

VIII. CDO SQUAREDS

A CDO squared is an example of a credit derivative with a non-standard structure. It is similar to a CDO or a CDS in that, in exchange for a set of periodic payments, it provides protection against losses due to default. The distinction is that in a CDO-squared the default losses arise from a portfolio of CDO tranches and the loss covered is a subset of the total loss. The CDO-squared is referred to as the parent CDO while the CDOs from which the CDO tranches underlying the CDO-squared are drawn are known as child CDOs.

In a typical structure there are 10 child CDOs each containing about 80 CDSs. Some of the names appear in more than one of the child CDOs. We define ‘overlap’, a measure of commonality between two CDOs, as the number of names in common divided by the average number of firms in the two CDOs. For example, if one CDO has 80 names in it and one has 100 names, and there are 25 names in common, the overlap is 25/90 or about 28%. In a typical CDO-squared the average overlap between the child CDOs is about 25%.

In this section we will construct a typical CDO-squared and calculate the breakeven tranche spread using the implied copula model. We will then determine what copula correlation would produce the same spread if the value were determined using the standard Gaussian copula.

The Example

Our example is based on a pool of 500 homogeneous names. We calibrated the implied copula model to the five-year August 30, 2005 iTraxx market data and set the spread for each of the 500 names equal to that for the index. Ten portfolios of 80 names were created using a quasi-random selection process. The average of the 45 pairwise overlap measures, as we have just defined them, was 23%; the highest overlap between two portfolios was 30% and the lowest overlap was 14%.

Each of the 10 portfolios has a CDO tranche based on it with attachment points a_{Lm} and a_{Hm}, m=1 to 10. If K is the initial notional on each of the CDSs, the maximum possible loss on each of the 10 tranches is $80(a_{Hm} - a_{Lm})K$. The total possible loss across all 10 tranches is

$$K_{\max} = \sum_{m=1}^{10} 80\left(a_{Hm} - a_{Lm}\right)K$$

$K_{\max}$ can be thought of as the notional underlying the CDO-squared. The CDO-squared also has attachment point a_L and detachment point a_H, and the seller of protection is required to cover all losses on the 10 CDO tranches between $a_L K_{\max}$ and $a_H K_{\max}$. The seller of protection receives spread income at rate s on an initial notional of $(a_H - a_L)K_{\max}$. As losses that affect the CDO-squared arise, the notional on which the spread income is earned is reduced as in a regular CDO. Once total losses on the ten CDO tranches reach $a_H K_{\max}$ no further spread income is earned.

Valuation

The value of the CDO-squared using the implied copula approach is calculated by Monte Carlo simulation. The procedure is as follows. For each of the 500 names we draw a random variable, x_j, j=1 to 500. The x_j are uniformly distributed between zero and one. For each λ_k, k=1 to L, we calculate the default time, τ_j, for each name based on the random sample,

$$\tau_j = Q^{-1}\left(x_j \middle| \lambda_k\right)$$

We also calculate the recovery rate based on the value of λ_k. For each portfolio *m* we accumulate all the τ_j's and the associated loss given default for all the names in that portfolio for which $0 < \tau_j < T$ where *T* is the life of the CDO, five years. These are sorted in increasing time of default. Losses within each portfolio are aggregated over time and the losses that lie between 80 a_{Lm} *K* and 80 a_{Hm} *K* along with their default times are retained. These losses are pooled across all portfolios, resorted in order of time of default and the losses between $a_L K_{\max}$ and $a_H K_{\max}$ together with their default times are retained. These are the losses that affect the CDO-squared. The size and timing of these losses are used to calculate the present value of spread income (when paid at a rate of $1 per year) and the present value of losses for the CDO squared.

Analogous to Equation (1) we will refer to the present value of spread income, the present value of accrual payments, and the present value of losses as $A(\lambda_k)$, $B(\lambda_k)$, and $C(\lambda_k)$ respectively. Using the same set of random numbers this procedure is repeated for every λ_k. The first simulation result is then

$$A_i = \sum_{k=1}^{L} \pi_k A(\lambda_k)$$

$$B_i = \sum_{k=1}^{L} \pi_k B(\lambda_k)$$

$$C_i = \sum_{k=1}^{L} \pi_k C(\lambda_k)$$

for *i*=1. At this point a new set of 500 random numbers is drawn and the procedure is repeated to calculate A_2, B_2 and C_2. This is repeated *n* times and the final sample estimates are

$$\hat{A} = \frac{1}{n}\sum_{i=1}^{n} A_i \quad \hat{B} = \frac{1}{n}\sum_{i=1}^{n} B_i \quad \hat{C} = \frac{1}{n}\sum_{i=1}^{n} C_i$$

$$\hat{s} = \frac{\hat{C}}{\hat{A} + \hat{B}}$$

where $\hat{s}$ is the estimated breakeven spread.

Comparison of Models

Very little market data for CDO squareds is available. The implied copula model has the key advantage that it fits all available market data on CDOs and the index. The prices produced by that model are therefore consistent with the observed prices; that is, these prices do not allow arbitrage opportunities. To test the viability of using the Gaussian copula model to value CDO squareds we imply a Gaussian copula correlation from the prices given by the implied copula model. We then examine whether this correlation might reasonably have been used by market participants.

Usually the child CDO tranches underlying a CDO-squared are mezzanine tranches. So as our base case we choose a case in which all the child CDO tranches have attachment and detachment points a_{Lm}=3% and a_{Hm}=6%, the same as the iTraxx mezzanine tranche. The child CDO contains 80 names. Ten different parent CDO cases are considered with attachment and detachment points of 0 to 10%, 10 to 20% and so on up to 90 to 100%.

Four other cases were considered:

Senior Case 1: The basic structure is the same as the base case except that the child CDO tranches are all 6 to 9% senior tranches. The same ten parent CDO cases are considered.

Senior Case 2: The basic structure is the same as the base case except that the child CDO tranches are all 9 to 12% senior tranches. The same ten parent CDO cases are considered.

Mixed Case 3: In this case several child CDO tranches are mixed together. Three of the child CDO tranches are 3 to 6% mezzanine tranches, four are 6 to 9% senior tranches, and three are 9 to 12% senior tranches. The same ten parent CDO cases are considered. (This type of structure is rarely seen in the market.)

Mixed Case 4: This is similar to case 3 except three of the child CDO tranches are 0 to 3% equity tranches, four are 3 to 6% mezzanine tranches, and three are 6 to 9% senior tranches. The same ten parent CDO cases are considered.

For each of the cases the breakeven spread for each of the ten parent tranches was calculated. The results are shown in Exhibit 10. All spreads are quoted in basis points per year (running rates). For the first three cases in which all the child CDOs are the same, the breakeven spread for the relevant CDO tranche is shown in the rightmost column for reference purposes.

The implied Gaussian copula correlations are shown in Exhibit 11. It is difficult to discern any pattern in the implied copula correlations relative to the correlations for the underlying CDO tranches. For parent tranche 4 {30%, 40%} the implied Gaussian copula correlation is approximately the same as that for the child tranche, but this is not true of other tranches.[11]

For 'Senior Case 2' in which all the underlying tranches are 9 to 12% tranches the implied correlation for tranches 2 to 7 are close to the underlying tranche correlation. For the other two cases the implied correlation is usually significantly higher or lower than the underlying tranche correlation. The base case exhibits the sort of correlation smile that is usually observed in the CDO market.

Our overall conclusion is that it would be difficult to devise a scheme that would allow us to determine the copula correlation that would be appropriate for any particular CDO-squared tranche.

Impact of Overlap

One of the factors that affect the pricing of a CDO-squared is the degree of overlap between the names in the child CDOs. Our examples so far have used 10 child CDOs of 80 names with an average overlap of about 23%. This means that on average when two child CDOs are compared there will be about 19 or 20 names in common between the two portfolios. Exhibit 12 compares this situation with situations where the overlap is 46% and 70% for the base case. Increasing the degree of overlap for a CDO-squared has an effect similar to increasing the correlation in valuing the tranches of a CDO. Increasing the degree of overlap increases the cost of CDO-squared tranches with high attachment points and decreases the cost of tranches with low attachment points.

[11] As explained earlier the Gaussian copula implied correlation for a particular tranche of an 80-company portfolio is different from that for a 125-company portfolio.

IX. CONCLUSIONS

We have presented a new approach to modeling default dependence. The approach has a number of advantages over the Gaussian copula model and its extensions. First, the model can be exactly fitted to the market quotes for the actively traded CDO tranches of standard portfolios. Second, the model is more intuitive than the base correlation model. Third, by perturbing market quotes such as those in Exhibit 1 Greek letters showing sensitivity to the quotes can be calculated. Fourth, the model can be used to value CDOs on bespoke portfolios and more exotic structures such as CDO-squareds. Finally, the model is ideally suited for trading. A trader can first calculate the implied distributions such as those in Exhibit 4. She can then investigate the effect on market prices of modifying the distributions to reflect her beliefs. A natural extension is to a two-factor model that fits iTraxx and CDX tranches simultaneously. This extension can be achieved by assuming a copula correlation structure for the hazard rate distributions for North America and Europe.

We have shown that the implied copula approach is more stable than the Gaussian copula/base correlation approach. The latter is liable to produce results that are inconsistent with market quotes. The implied copula approach is sometimes criticized because it depends on the number of hazard rates chosen and the smoothness condition used. However, the Gaussian copula/base correlation model is no less arbitrary. It depends on the number of values chosen for the factor and on the interpolation scheme used for base correlation.

There are limitations of both the implied copula and the Gaussian copula/base correlation approach. They do not involve the dynamic evolution of hazard rates or credit spreads. They are therefore inappropriate for some instruments. For example, the models are not appropriate for valuing a one-year option on a five-year CDO because this depends on the hazard rate distribution between years one and five conditional on what we observe happening during the first year.

References

Altman, E.I., B. Brady, A. Resti, and A. Sironi. "The Link Between Default Rates and Recovery Rates: Implications for Credit Risk Models and Procyclicality." Working Paper, Stern School of Business, New York University, April 2002.

Andersen, L. and J. Sidenius. "Extensions of the Gaussian Copula Model." *Journal of Credit Risk*, 1, 1 (2004/2005), 29-70.

Andersen, L., J. Sidenius, and S. Basu. "All Your Hedges in One Basket." *RISK*, November, 2003.

Breeden, D.T., and R. H. Litzenberger. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business*, 51 (1978), 621-51.

Cantor, R., D.T. Hamilton, and S. Ou. "Default rates and recovery rates of corporate bond issuers." Moody's Investor's Services, February, 2002.

Gregory, J., and J-P Laurent. "Basket Default Swaps, CDO's and Factor Copulas." Journal of Risk, 7, 4 (2005), 103-22.

Hamilton, D.T., P. Varma, S. Ou, and R. Cantor. "Default and Recovery Rates of Corporate Bond Issuers." Moody's Investor's Services, January 2005.

Hull, J., and A. White. "Valuation of a CDO and nth to Default CDS without Monte Carlo Simulation." *Journal of Derivatives*, 12, 2 (Winter 2004), 8-23.

Jackwerth, J. C., and M. Rubinstein. "Recovering Probability Distributions from Option Prices." *Journal of Finance*, 51 (December 1996), 1611-31.

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

Vasicek, O. ``Probability of Loss on a Loan Portfolio." Working Paper, KMV, 1987. (Published in Risk in December 2002 under the title ``Loan Portfolio Value.")

Walker, M. "CDO Models-Towards the Next Generation: Incomplete Markets and Term Structure." Working Paper, University of Toronto, 2006.

EXHIBIT 1: Quotes for CDX IG and iTraxx tranches on August 30, 2005. Quotes for the 0 to 3% tranche are the percent of the principal that must be paid up front in addition to 500 basis points per year. Quotes for other tranches and the index are in basis points. Source: Reuters

	CDX IG Tranches					
	0% to 3%	3% to 7%	7% to 10%	10% to 15%	15% to 30%	Index
5-year Quotes	40%	127	35.5	20.5	9.5	50

	ITraxx Tranches					
	0% to 3%	3% to 6%	6% to 9%	9% to 12%	12% to 22%	Index
5-year Quotes	24%	81	26.5	15	9	36.375
10-year Quotes	53%	395	90	52	29	57.625

EXHIBIT 2: Results of fitting the implied copula model to 5-year iTraxx and CDX data

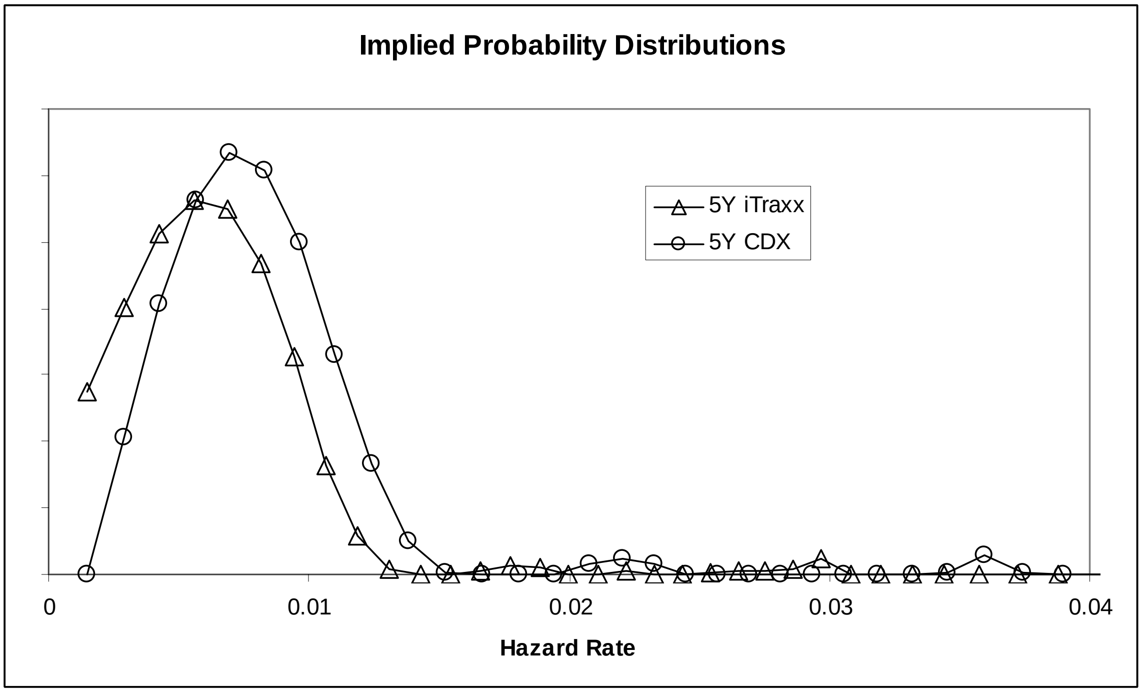

EXHIBIT 3: Impact on the hazard rate probability distribution of increasing the number of points, *L*, that are used.

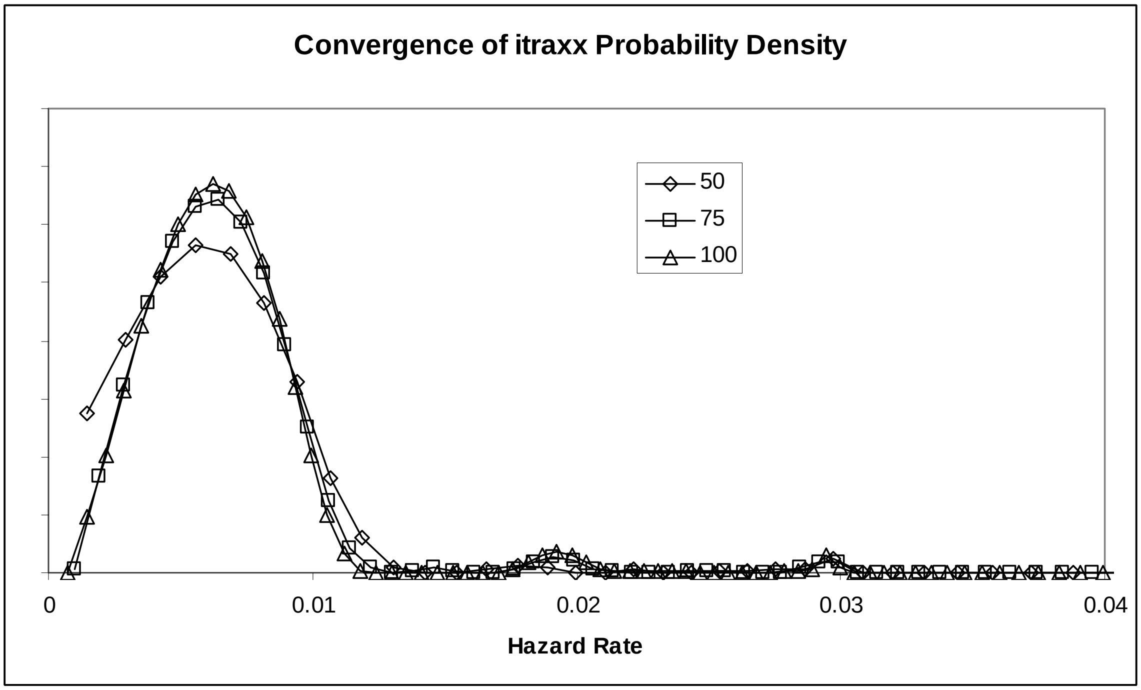

EXHIBIT 4: Change in the iTraxx (upper chart) and CDX (lower chart) hazard rate probability distribution between July 26, 2004 and November 2, 2005

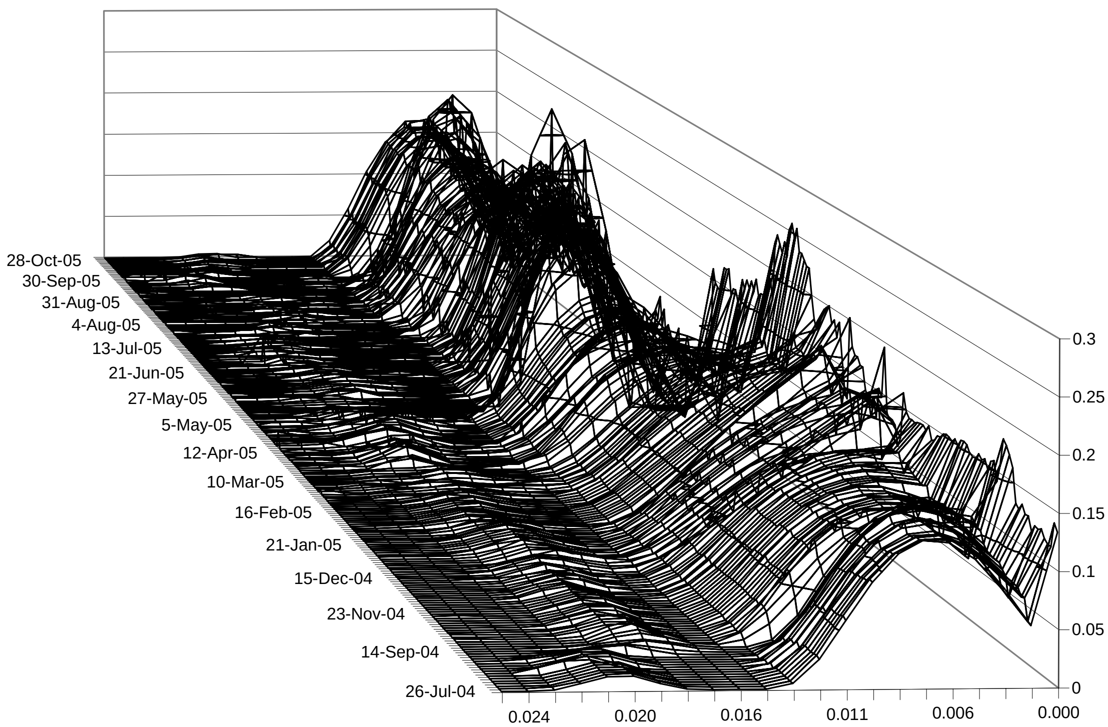

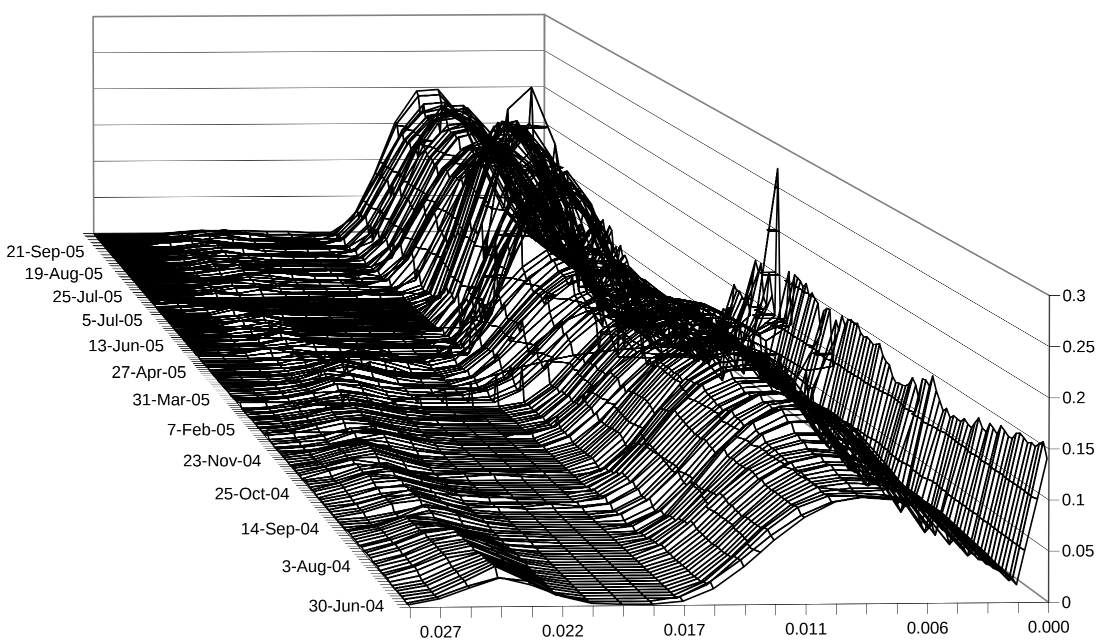

EXHIBIT 5: Implied tranche and base correlations for 5-year CDX IG and 5-year iTraxx tranches on August 30, 2005.

	CDX IG					
	0% to 3%	3% to 7%	7% to 10%	10% to 15%	15% to 30%	Index
Tranche Correlation	0.091	0.012	0.068	0.106	0.156	n/a
Base Correlation	0.091	0.177	0.223	0.280	0.448	n/a

	iTraxx					
	0% to 3%	3% to 6%	6% to 9%	9% to 12%	12% to 22%	Index
Tranche Correlation	0.134	0.030	0.079	0.111	0.154	n/a
Base Correlation	0.134	0.210	0.266	0.308	0.428	n/a

EXHIBIT 6: Tranche correlations when the tranche width is 1% calculated using the base correlation and implied copula approaches. Linear and spline interpolation schemes are used in the implementation of the base correlation approach. The upper panel shows the iTraxx results, the lower CDX results.

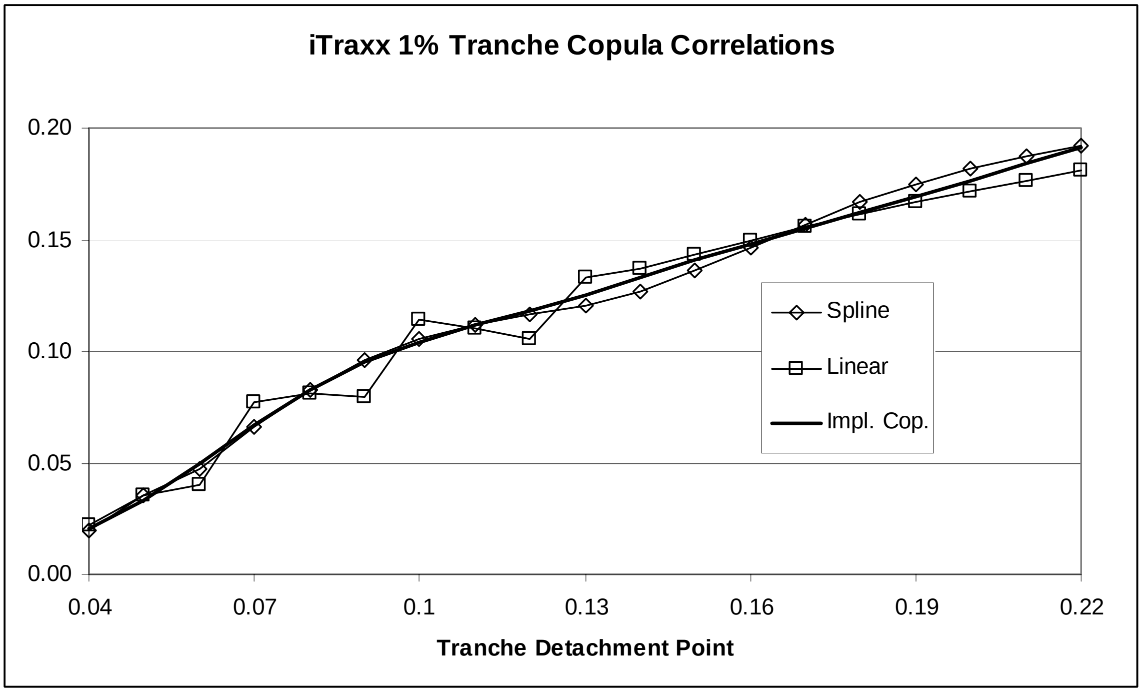

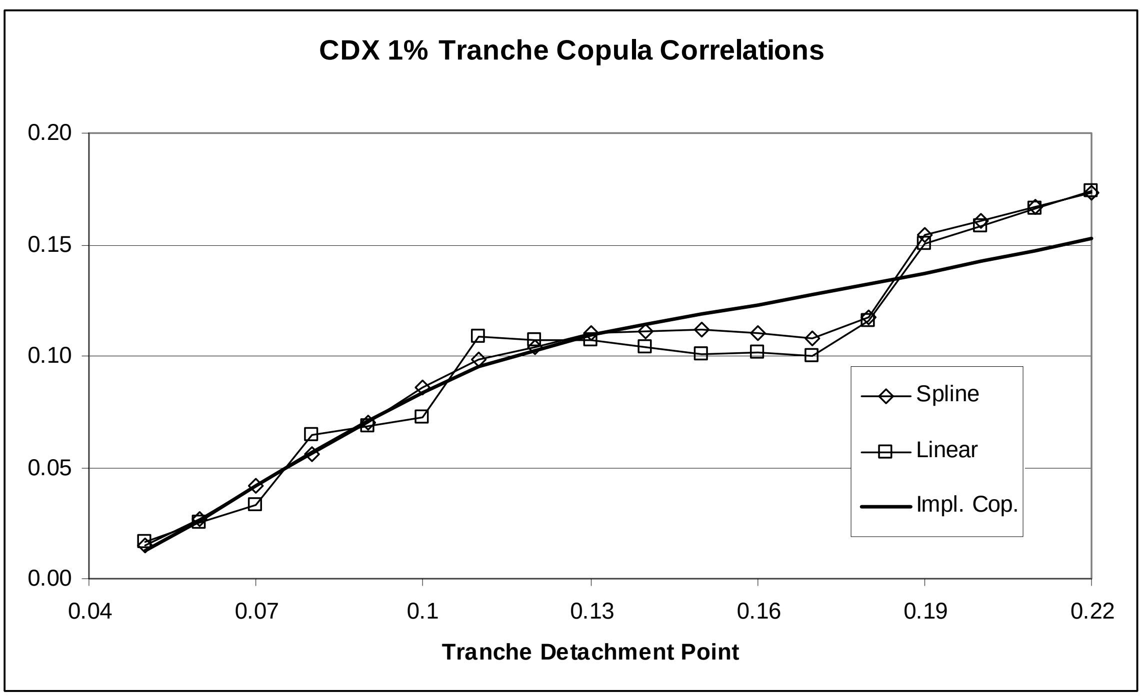

EXHIBIT 7: Tranche correlations when the tranche width is 1% implied using the base correlation method. These are compared with the maximum and minimum possible correlations implied from the implied copula approach. Linear and spline interpolation schemes are used in the implementation of the base correlation approach. The upper panel shows the iTraxx results, the lower CDX results.

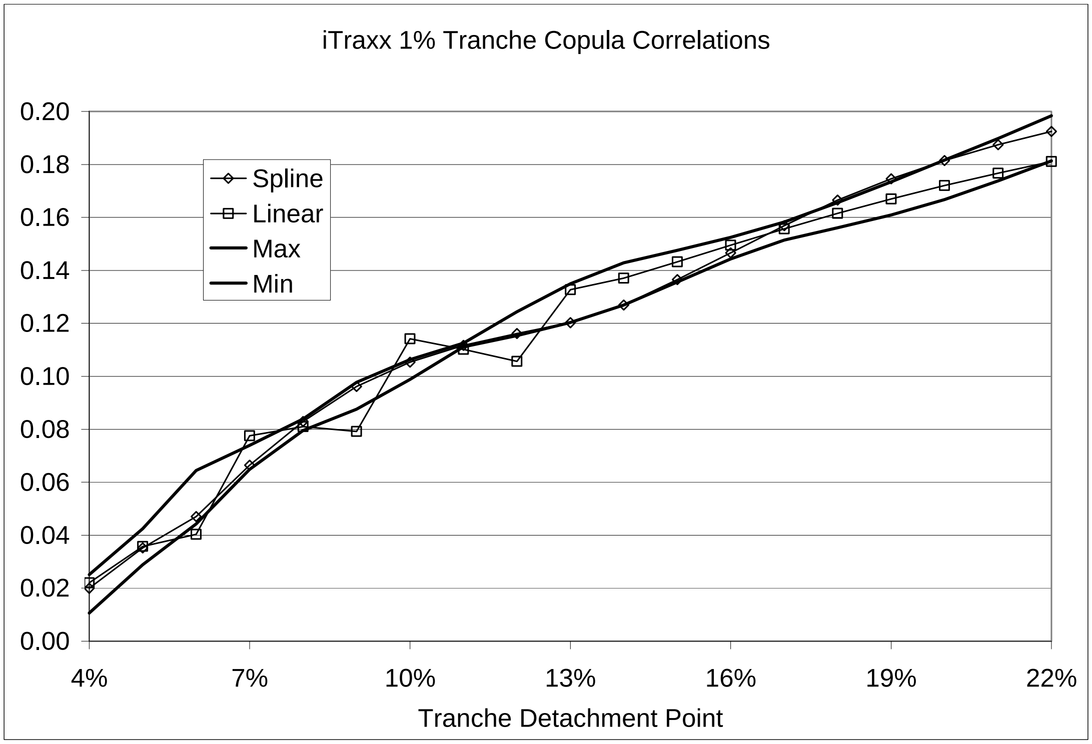

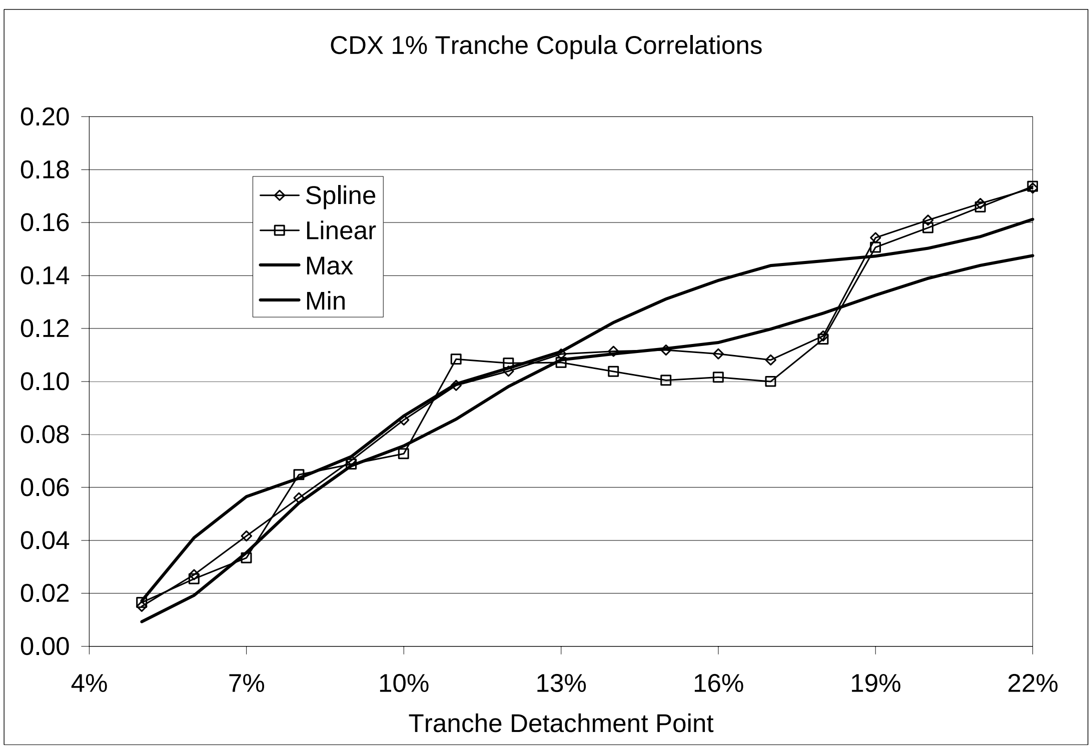

EXHIBIT 8: Breakeven Spread for 5-Year iTraxx Mezzanine Tranche with Different Portfolio Sizes on August 30, 2005.

	Number of Firms in Portfolio						
	40	60	80	100	125	150	200
Implied Copula	154.4	122.4	103.7	91.6	81.0	75.2	67.0
Gaussian Copula: Tranche Correlation	174.3	133.3	109.4	94.0	81.0	73.1	62.8
Gaussian Copula: Base Correlation	124.8	103.9	93.4	86.8	81.0	77.6	73.0

EXHIBIT 9: Cumulative five- and ten-year default rates for iTraxx on August 30, 2005

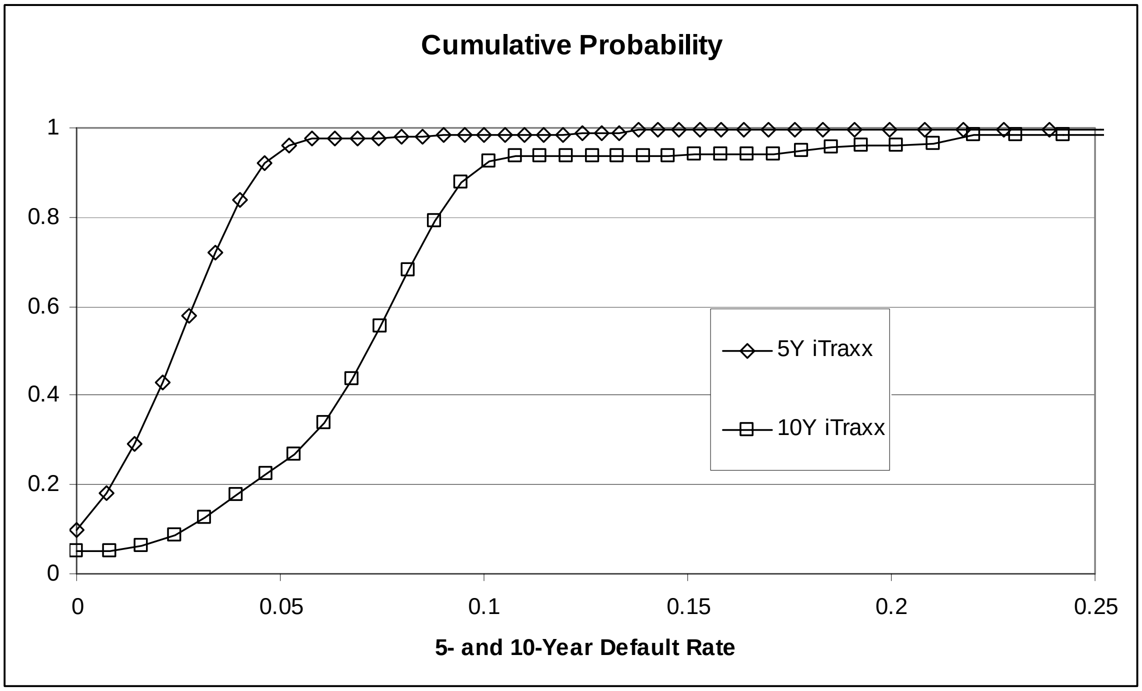

EXHIBIT 10: Breakeven spreads in basis points for ten different CDO-squareds with five different underlying CDO structures. When all child CDOs have the same attachment and detachment points the CDO tranche spread is also reported.

	CDO-Squared Tranche										
	1	2	3	4	5	6	7	8	9	10	CDO
Base Case	499	190	104	68	52	44	39	35	31	24	104
Senior Case 1	53	37	33	30	28	25	22	19	15	12	27
Senior Case 2	29	24	20	17	14	13	11	10	10	9	16
Mixed Case 3	234	63	40	33	29	25	20	16	12	10	
Mixed Case 4	3063	1056	395	140	65	43	35	30	24	16	

EXHIBIT 11: Implied copula correlations for ten different CDO-squareds with five different underlying CDO structures. When all child CDOs have the same attachment and detachment points the CDO tranche correlation is also reported.

	CDO-Squared Tranche										
	1	2	3	4	5	6	7	8	9	10	CDO
Base Case	0.240	0.019	0.024	0.028	0.035	0.043	0.052	0.062	0.076	0.094	0.026
Senior Case 1	0.036	0.056	0.068	0.079	0.088	0.095	0.101	0.106	0.112	0.122	0.071
Senior Case 2	0.079	0.098	0.105	0.108	0.110	0.112	0.117	0.124	0.134	0.151	0.106
Mixed Case 3	0.012	0.033	0.050	0.068	0.083	0.096	0.104	0.110	0.116	0.132	
Mixed Case 4	0.112	0.117	0.000	0.021	0.053	0.060	0.067	0.078	0.094	0.109	

EXHIBIT 12: Breakeven spreads for CDO-squared with different degrees of overlap in basis points

			Overlap		
Tranche	AP	DP	23%	46%	70%
1	0%	10%	498.7	428.3	359.0
2	10%	20%	190.2	187.6	183.7
3	20%	30%	103.6	116.1	127.1
4	30%	40%	68.3	81.4	94.9
5	40%	50%	52.3	62.5	74.3
6	50%	60%	43.9	51.0	60.3
7	60%	70%	38.7	43.6	50.1
8	70%	80%	34.6	37.8	43.0
9	80%	90%	30.5	32.9	37.1
10	90%	100%	23.7	26.5	30.3

Acknowledgements