

Valuation of a CDO and an n -th to Default CDS Without Monte Carlo Simulation

JOHN HULL AND ALAN WHITE

JOHN HULL

is a professor of finance at the Joseph L. Rotman School of Management at the University of Toronto.

ALAN WHITE

is a professor of finance at the Joseph L. Rotman School of Management at the University of Toronto.

Two fast procedures for valuing tranches of collateralized debt obligations and n -th to default swaps are developed here. The procedures are based on a factor copula model of times to default and are alternatives to using fast Fourier transforms. One involves calculating the probability distribution of the number of defaults by a certain time using a recurrence relationship; the other involves using a "probability bucketing" numerical procedure to build up the loss distribution.

Many different copula models can be generated by using different distributional assumptions in the factor model. The impact on valuations of default probabilities, default correlations, and a correlation of recovery rates with default probabilities is shown. An examination of the market pricing of index tranches indicates that a double t -distribution copula fits the prices reasonably well.

As the credit derivatives market has grown, products that depend on default correlations have become more popular. We focus on three of these products: collateralized debt obligations, index tranches, and n -th to default credit default swaps.

A collateralized debt obligation (CDO) offers a way to create securities with widely different risk characteristics from a portfolio of debt instruments. An example is shown in Exhibit 1, where four types of securities are created from a portfolio of bonds.

The first tranche of securities constitutes


5% of the total bond principal and absorbs all credit losses from the portfolio during the life of the CDO until they have reached 5% of the total bond principal. The second tranche has 10% of the principal and absorbs all losses during the life of the CDO in excess of 5% of the principal up to a maximum of 15% of the principal. The third tranche has 10% of the principal and absorbs all losses in excess of 15% of the principal up to a maximum of 25% of the principal. The fourth tranche with 75% of the principal absorbs all residual losses in excess of 25% of the principal.

The yields in Exhibit 1 are the rates of interest paid to tranche holders. These rates are paid on the balance of the principal remaining in the tranche after losses have been paid. In tranche 1, for example, initially the return of 35% is paid on the whole amount invested by the tranche 1 holders. But if losses equal to 1% of the total bond principal are experienced, tranche 1 holders have lost 20% of their investment, and the return is paid on only 80% of the original amount invested.

Tranche 1 is referred to as the *equity tranche*. A default loss of 2.5% on the bond portfolio translates into a loss of 50.0% of the tranche's principal. The creator of the CDO normally retains tranche 1 and sells the remaining tranches in the market. Tranche 4 is usually given an Aaa rating. Defaults on the bond portfolio must exceed 25% before the holders of this tranche are responsible for any credit losses.

EXHIBIT 1

Structure of a CDO



The CDO in Exhibit 1 is referred to as a *cash CDO*. An alternative structure is a synthetic CDO in which the creator sells a portfolio of credit default swaps to third parties. It then passes the default risk on to the synthetic CDO's tranche holders.

Analogously to Exhibit 1, the first tranche might be responsible for the payoffs on the credit default swaps until they have reached 5% of the total notional principal; the second tranche might be responsible for the payoffs between 5% and 15% of the total notional principal, and so on.

The income from the credit default swaps is distributed to the tranches in a way that reflects the risk they are bearing. For example, tranche 1 might get 3,000 basis points per year; tranche 2 might get 1,000 basis points per year, and so on. As in a cash CDO, this would be paid on a principal that declines as defaults for which the tranche is responsible occur.

Participants in credit derivatives markets have developed indexes to track credit default swap spreads. For example, the Dow Jones CDX NA IG five-year index gives the average five-year credit default swap spread for a portfolio of 125 investment-grade U.S. companies. Similarly, the Dow Jones iTraxx EUR five-year index is the average credit default swap spread for a portfolio of 125 investment-grade European companies.

The portfolios underlying indexes are used to define standardized index tranches similar to the tranches of a CDO. In the case of the CDX NA IG five-year index, successive tranches are responsible for 0% to 3%, over 3% to 7%, over 7% to 10%, over 10% to 15%, and over 15% to 30% of the losses. In the case of the iTraxx EUR five-

year index, successive tranches are responsible for 0% to 3%, over 3% to 6%, over 6% to 9%, over 9% to 12%, and over 12% to 22% of the losses.

Derivatives dealers have created a market to facilitate the buying and selling of index tranches that is proving very popular with investors. An index tranche is different from the tranche of a synthetic CDO in that an index tranche is not funded by the sale of a portfolio of credit default swaps. The rules for determining payoffs, however, ensure that an index tranche is economically equivalent to the corresponding synthetic CDO tranche.

A CDO is closely related to an *n*-th to default credit default swap (CDS). In an *n*-th to default CDS, the buyer of protection pays a specified rate (known as the *CDS spread*) on a specified notional principal either until the *n*-th default occurs among a specified set of reference entities or until the end of the contract's life. The payments are usually made quarterly. If the *n*-th default occurs before the contract maturity, the buyer of protection can present bonds issued by the defaulting entity to the seller of protection in exchange for the face value of the bonds. Alternatively, the contract may call for a cash payment equal to the difference between the post-default bond value and the face value.¹

We develop procedures for valuing both an *n*-th to default CDS and tranches of a CDO or index. Our model is a multifactor copula model similar to that used by researchers such as Li [2000], Laurent and Gregory [2003], and Andersen and Sidenius [2004].² Like other researchers, we calculate a distribution for the default loss by a certain time conditional on the factor values and then integrate over the factor values. The advantage of this procedure is that the conditional default losses for different companies are independent.

Laurent and Gregory use the fast Fourier transform method to calculate the conditional loss distribution on a portfolio as a convolution of the conditional loss distributions of each of the companies constituting the portfolio. We present two alternative approaches.

The first involves a recurrence relationship to determine the probability of exactly *k* defaults occurring by time *T* and works well for an *n*-th to default CDS and tranches of an index or a CDO when each company has the same weight in the portfolio and recovery rates are assumed to be constant. The second involves an iterative procedure that we refer to as *probability bucketing* for building up the portfolio loss distribution, which can be used in a wider range of situations. The second approach is the one we recommend for CDO pricing. It is a robust and flexible

approach that has some advantages over a similar approach developed by Andersen, Sidenius, and Basu [2003].

We evaluate the sensitivity of spreads for an n -th to default CDS and a CDO to a variety of different assumptions concerning default probabilities, recovery rates, and the correlation model chosen. We also explore the impact of dependencies between recovery rates and default probabilities, and examine the market pricing of index tranches and the interpretation of implied correlations.

I. THE DEFAULT CORRELATION MODEL

Default correlation measures the tendency of two companies to default at about the same time. Two types of default correlation models suggested by researchers are reduced-form models and structural models. Reduced-form models such as those in Duffie and Singleton [1999] assume that the default intensities for different companies follow correlated stochastic processes. Structural models are based on Merton's [1974] model, or one of its extensions, where a company defaults when the value of its assets falls below a certain level. Default correlation is introduced into a structural model by assuming that the assets of different companies follow correlated stochastic processes.

The reduced-form model and the structural model are computationally very time-consuming for valuing the types of instruments we are considering. This has led market participants to model correlation using a factor copula model, where the joint probability distribution for the times to default of many companies is constructed from the marginal distributions.

Consider a portfolio of N companies, and assume that the marginal risk-neutral probabilities of default are known for each company. Define:

- t_i : the time of default of the i -th company;
- $Q_i(t)$: the cumulative risk-neutral probability that company i will default before time t ; that is, the probability that $t_i \leq t$; and
- $S_i(t) = 1 - Q_i(t)$: the risk-neutral probability that company i will survive beyond time t ; that is, the probability that $t_i > t$.

To generate a one-factor model for the t_i , we define random variables x_i ($1 \leq i \leq N$):

$$x_i = a_i M + \sqrt{1 - a_i^2} Z_i \quad (1)$$

where M and the Z_i have independent zero-mean unit-

variance distributions and $-1 \leq a_i \leq 1$. Equation (1) defines a correlation structure between the x_i dependent on a single common factor M . The correlation between x_i and x_j is $a_i a_j$.

Let F_i be the cumulative distribution of x_i . Under the copula model, the x_i are mapped to the t_i using a percentile-to-percentile transformation. The five-percentile point in the probability distribution for x_i is transformed to the five-percentile point in the probability distribution of t_i ; the ten-percentile point in the probability distribution for x_i is transformed to the ten-percentile point in the probability distribution of t_i ; and so on. In general, the point $x_i = x$ is transformed to $t_i = t$ where $t = Q_i^{-1}[F_i(x)]$.

Let H be the cumulative distribution of the Z_i .³ It follows from Equation (1) that:

$$\text{Prob}(x_i < x | M) = H \left[\frac{x - a_i M}{\sqrt{1 - a_i^2}} \right]$$

When $x = F_i^{-1}[Q_i(t)]$, $\text{Prob}(t_i < t) = \text{Prob}(x_i < x)$. Hence

$$\text{Prob}(t_i < t | M) = H \left\{ \frac{F_i^{-1}[Q_i(t)] - a_i M}{\sqrt{1 - a_i^2}} \right\}$$

The conditional probability that the i -th bond will survive beyond time T is therefore:

$$S_i(T|M) = 1 - H \left\{ \frac{F_i^{-1}[Q_i(T)] - a_i M}{\sqrt{1 - a_i^2}} \right\} \quad (2)$$

Extension to Many Factors

The model we have presented can be extended to many factors. Equation (1) becomes

$$x_i = a_{i1} M_1 + a_{i2} M_2 + \dots + a_{im} M_m + Z_i \sqrt{1 - a_{i1}^2 - a_{i2}^2 - \dots - a_{im}^2}$$

where $a_{i1}^2 + a_{i2}^2 + \dots + a_{im}^2 < 1$, and the M_j have independent distributions with zero mean and unit variance. The correlation between x_i and x_j is then $a_{i1} a_{j1} + a_{i2} a_{j2} + \dots + a_{im} a_{jm}$.

Equation (2) becomes:

$$S_i(T|M_1, M_2, \dots, M_m) = 1 - H \left\{ \frac{F_i^{-1}[Q_i(T)] - a_{i1}M_1 - a_{i2}M_2 - \dots - a_{im}M_m}{\sqrt{1 - a_{i1}^2 - a_{i2}^2 - \dots - a_{im}^2}} \right\} \quad (3)$$

Distribution Assumptions

The advantage of the copula model is that it creates a tractable multivariate joint distribution for a set of variables that is consistent with known marginal probability distributions for the variables. One possibility is to let the M and the Z have standard normal distributions. A Gaussian copula then results. Any distributions can be used for the M and the Z , however, providing they are scaled so that they have zero mean and unit variance. Each choice of distributions results in a different copula model.

The choice of the copula governs the nature of the default-dependence. For example, as we will see, copulas in which the M have heavy tails generate models with a greater likelihood of a clustering of early defaults for several companies. Later we will explore the effect of using normal and t -distributions for the M and the Z .

Implementation

We present two new approaches for implementing the model so as to value an n -th to default CDS or the tranches of CDOs and indexes. The first, which involves calculating the probability distribution of the number of defaults by a time T , is ideal when companies have equal weight in the portfolio, and recovery rates are assumed to be constant. The second, which involves calculating the probability distribution of the total loss from defaults by time T , can be used for a wide range of assumptions.

II. FIRST IMPLEMENTATION APPROACH

Define $\pi_T(k)$ as the probability that exactly k defaults occur in the portfolio before time T . Conditional on the M , the default times, t_p , are independent. It follows that the conditional probability that all the N bonds will survive beyond time T is:

$$\pi_T(0|M_1, M_2, \dots, M_m) = \prod_{i=1}^N S_i(T|M_1, M_2, \dots, M_m) \quad (4)$$

where $S_i(T|M_1, M_2, \dots, M_m)$ is given by Equation (3). Similarly:

$$\pi_T(1|M_1, M_2, \dots, M_m) = \pi_T(0|M_1, M_2, \dots, M_m) \times \sum_{i=1}^N \frac{1 - S_i(T|M_1, M_2, \dots, M_m)}{S_i(T|M_1, M_2, \dots, M_m)}$$

Define:

$$w_i = \frac{1 - S_i(T|M_1, M_2, \dots, M_m)}{S_i(T|M_1, M_2, \dots, M_m)}$$

The conditional probability of exactly k defaults by time T is

$$\pi_T(k|M_1, M_2, \dots, M_m) = \pi_T(0|M_1, M_2, \dots, M_m) \sum w_{z(1)} w_{z(2)} \dots w_{z(k)} \quad (5)$$

where $\{z(1), z(2), \dots, z(k)\}$ is a set of k different numbers chosen from $\{1, 2, \dots, N\}$, and the summation is taken over the

$$\frac{N!}{k!(N-k)!}$$

different ways the numbers can be chosen. Appendix A provides a fast way to compute this.

The unconditional probability that there will be exactly k defaults by time T , $\pi_T(k)$, can be determined by numerically integrating $\pi_T(k|M_1, M_1, \dots, M_m)$ over the distributions of the M_j .⁴

The probability that there will be at least n defaults by time T is:

$$\sum_{k=n}^N \pi_T(k)$$

The probability that the n -th default occurs between time T_1 and time T_2 is the difference between the value of this expression for $T = T_2$ and its value for $T = T_1$.

This approach does give rise to occasional numerical stability problems.⁵ These can be handled using the approach given in Appendix A and other straightforward procedures.

III. SECOND IMPLEMENTATION APPROACH

The second implementation (our *probability bucketing* approach) is described in Appendix B. It calculates the prob-

ability distribution of the losses by time T . We divide potential losses into ranges as follows: $\{0, b_0\}$, $\{b_0, b_1\}$, ..., $\{b_{K-1}, \infty\}$. We refer to $\{0, b_0\}$ as the 0-th bucket, $\{b_{k-1}, b_k\}$ as the k -th bucket ($1 \leq k \leq K-1$), and $\{b_{K-1}, \infty\}$ as the K -th bucket. The loss distribution is built up one debt instrument at a time. The procedure keeps track of both the probability the cumulative loss is in a bucket and the mean cumulative loss conditional that the cumulative loss is in the bucket.

Andersen, Sidenius, and Basu [2003] develop a similar procedure, where discrete losses $0, u, 2u, 3u, \dots, n^*u$ are considered for some u (with n^*u the maximum possible loss), and the losses considered are rounded to the nearest discrete point as the loss distribution is built up. When u is a common divisor of all potential losses, our approach with a bucket width of u is the same as that approach. In other circumstances, we find it to be more accurate because it keeps track of the mean loss within each bucket.

Our approach can provide extra accuracy in some regions of the loss distribution by using smaller bucket sizes in those regions. Also, we truncate the loss distribution at b_{K-1} , so that we do not need to spend computation time on large losses that have virtually no chance of occurring. The method works well when recovery rates are stochastic.

Define:

$p_T(k)$: the probability that the loss by time T lies in the k -th bucket, and

$P(k, T)$: the probability that the loss by time T is greater than b_{k-1} (i.e., that it lies in the k -th bucket or higher).

The approach in Appendix B calculates the conditional probabilities $p_T(k | M_1, M_2, \dots, M_m)$. As in the case of $\pi_T(k)$, we can calculate the unconditional probability $p_T(k)$ by integrating over the distributions of the M_j .

We have compared our approach for calculating the $p_T(k)$ with the fast Fourier transform (FFT) approach of Laurent and Gregory [2003]. Our approach has the advantage of being more intuitive. We find that the two approaches, for a given bucket size, are very similar in terms of computation speed. We also compare both approaches with Monte Carlo simulation, using a very large number of trials (so that the Monte Carlo results could be assumed to be correct). In our comparisons we use the same bucketing scheme for both approaches. The approach in Appendix B always gives reasonably accurate answers. We find the performance of the FFT method is quite sensitive to the bucket size. When the bucket size

is such that FFT works well, the two approaches are about equally accurate; in other circumstances, Appendix B works much better.⁶

Both methods can be used to compute Greek letters quickly. Both methods can be used to calculate the probability distribution of the number of defaults by time T by setting the principal for each reference entity equal to one and the recovery rate equal to zero.

After the procedure in Appendix B has been carried out, we assume that losses are concentrated at the midpoints of the buckets for the purposes of integrating over factors. The probabilities $P(n, T)$ are given by

$$P(n, T) = \sum_{k=n}^K p_T(k)$$

We estimate the probability that a loss equal to $0.5(b_{k-1} + b_k)$, the midpoint of bucket k , first happens between T_1 and T_2 as

$$0.5[P(k, T_2) + P(k+1, T_2)] - 0.5[P(k, T_1) + P(k+1, T_1)]$$

IV. RESULTS FOR AN N-TH TO DEFAULT CDS

In an n -th to default CDS, we assume that the principal amounts and the expected recovery rates are the same for all underlying reference assets. The valuation procedure is similar to that for a regular CDS where there is only one reference entity.⁷ In a regular CDS, the valuation is based on the probability that a default will occur between times T_1 and T_2 . Here the valuation is based on the probability that the n -th default will occur between times T_1 and T_2 .

We assume the buyer of protection makes quarterly payments in arrears at a specified rate until the n -th default occurs, or until the end of the life of the contract. In the event of an n -th default, the seller pays the notional principal times $1 - R$. Also, there is a final accrual payment by the buyer of protection to cover the time elapsed since the previous payment. The contract can be valued by calculating the expected present value of payments and the expected present value of payoffs in a risk-neutral world. The break-even CDS spread is the one at which the expected present value of the payments equals the expected present value of the payoffs.⁸

Consider first a five-year n -th to default CDS on a basket of ten reference entities when the expected recovery rate, R , is 40%. The term structure of interest rates is assumed to be flat at 5%. The default probabilities for the ten entities are generated by Poisson processes with constant default intensities, λ_i , ($1 \leq i \leq 10$) so that:⁹

$$S_i(t) = e^{-\lambda_i t}; \quad Q_i(t) = 1 - e^{-\lambda_i t}$$

In the base case, we use a one-factor model where $\lambda_i = 0.01$ for all i (so that all entities have a probability of about 1% of defaulting each year). The correlation between all pairs of reference entities is 0.3. This means that $a_i = \sqrt{0.3}$ for all i in Equation (1). M and the Z_i are assumed to have standard normal distributions.

Exhibit 2 shows that in the base case the buyer of protection should be willing to pay 440 basis points per year for a first-to-default swap, 139 basis points per year for a second-to-default swap, 53 basis points per year for a third-to-default swap, and so on.

Impact of Default Probabilities and Correlations

Exhibits 2 and 3 show the impact of changing the default intensities and correlations. Increasing the default intensity for all firms raises the cost of buying protection in all n -th to default CDS. The cost of protection rises at a declining rate for low n and at an increasing rate for high n .

Increasing the pairwise correlations between all firms while holding the default intensity constant reduces the cost of protection in an n -th to default CDS if n is small and increases it if n is large. To understand the reason for this, consider what happens as we increase the pairwise correlations from zero to one. When the correlation is zero, the cost of default protection is a sharply declining function of n . In the limit, when the default times are perfectly correlated, all entities default at the same time, and the cost of n -th to default protection is the same for all n . As correlations increase, we are progressing from the first case to the second case, so that the cost of protection declines for low n and increases for high n .

Impact of Distributional Assumptions

Exhibit 4 shows the effect of using a t -distribution instead of a normal distribution for M and the Z_i in Equation (1). The variable n_M is the number of degrees of freedom of the distribution for M , and n_Z is the number of degrees of freedom for the distribution of the Z_i . As the degrees of freedom become high, the t -distribution converges to a standard normal.

We consider three alternatives to the base case. In the first, M has a t -distribution with 5 degrees of freedom, and the Z_i are normal. In the second, M is normal, and the Z_i have t -distributions with 5 degrees of freedom. In

EXHIBIT 2

Spread to Buy 5-Year Protection for n -th Default

n	Default Intensity for all Firms		
	0.01	0.02	0.03
1	440	814	1165
2	139	321	513
3	53	149	263
4	21	71	139
5	8	34	72
6	3	15	36
7	1	6	16
8	0	2	6
9	0	1	2
10	0	0	0

All ten firms in basket have the same probability of default. The correlation between each pair of names is 0.3. The spread is in basis points per year.

EXHIBIT 3

Impact of Correlations

n	Pairwise Correlations		
	0.00	0.30	0.60
1	603	440	293
2	98	139	137
3	12	53	79
4	1	21	49
5	0	8	31
6	0	3	19
7	0	1	12
8	0	0	7
9	0	0	3
10	0	0	1

All pairs of firms have the same correlation. The default intensity for each firm is 0.01. The spread is in basis points per year.

the third, both M and the Z_i have t -distributions with 5 degrees of freedom.

Because a standard t -distribution with f degrees of freedom has a mean of zero and a variance of $f/(f-2)$, the random variable used for M in Equation (1) is scaled by

$$\sqrt{(n_M - 2)/n_M}$$

so that it has unit variance, and the random variable used for the Z_i is scaled by

$$\sqrt{(n_Z - 2)/n_Z}$$

for the same reason.

EXHIBIT 4

Effect of Different Distributional Assumptions on Spread to Buy Protection for n -th Default

N	Degrees of Freedom of t -Distributions (n_M / n_Z)			
	∞ / ∞	5 / ∞	$\infty / 5$	5 / 5
1	440	419	474	455
2	139	127	127	116
3	53	51	44	44
4	21	24	18	22
5	8	13	7	13
6	3	8	3	8
7	1	5	1	5
8	0	3	0	4
9	0	2	0	2
10	0	1	0	1

All pairs of firms have a 0.3 copula correlation. The default intensity for each firm is 0.01. The spread is in basis points per year.

EXHIBIT 5

Effect of Varying Probability of Default Across Firms

n	Correlation = 0		Correlation = 0.30	
	$\lambda = 0.01$	Disperse λ	$\lambda = 0.01$	Disperse λ
1	602.6	602.6	439.9	443.0
2	97.8	97.0	138.7	138.0
3	12.0	11.7	52.8	51.8
4	1.0	1.0	21.1	20.4
5	0.1	0.1	8.4	8.0
6	0.0	0.0	3.2	3.0
7	0.0	0.0	1.1	1.0
8	0.0	0.0	0.3	0.3
9	0.0	0.0	0.1	0.1
10	0.0	0.0	0.0	0.0

The average default probability is kept constant. In the base case the default intensity is 1% for each firm. In the comparison case the default intensity varies linearly from 0.0055 to 0.0145. The spread is in basis points per year.

Using heavier tails for M (small n_M) lowers the cost of protection in an n -th to default CDS if n is small and increases it if n is large. Using heavier tails for the Z_i distributions (small n_Z) raises the cost of protection in an n -th to default CDS if n is small and lowers it if n is large.

These results can be explained as follows. The value of x_i in Equation (1) can be thought of as determined by a sample from the distribution for M and a sample from the distribution for Z_i . When M has heavy tails and the Z are normal, an extreme value for a particular x_i is more likely to arise from an extreme value of M than an extreme

value of Z_i . It is therefore more likely to be associated with extreme values for the other x_i . Similarly, when the Z have heavy tails and M is normal, an extreme value for a particular x_i is more likely to arise from an extreme value of Z_i than an extreme value of M . It is therefore less likely to be associated with extreme values for the other x_i .

We deduce that an extreme case where the default times of several companies occur sooner becomes more likely as the tails of the distribution of M become heavier and less likely as the tails of the distribution of the Z become heavier. This explains the results in Exhibit 4. The overall effect of making the tails of M heavier is much the same as increasing the correlations among all entities, and the overall effect of making the tails of the Z heavier is much the same as reducing the correlations among all entities.

When both M and Z_i have t -distributions, we refer to the model as a *double t -distribution copula*. The cost of protection increases (over the base case) for small and large n and declines for intermediate values of n . As we will see later, the double t -distribution copula fits market prices reasonably well.

Impact of Dispersion in Default Intensities

Exhibit 5 shows the effect of setting the default intensities equal to

$$\lambda_i = 0.0055 + 0.001(i - 1)$$

The default intensities average 0.01 as in the base case, but they vary from 0.0055 to 0.0145.

When the default correlations are zero, the probability of no defaults by time T is:

$$\exp\left(-\sum_{i=1}^N \lambda_i T\right)$$

and the probability that the first-to-default will occur before time T is:

$$1 - \exp\left(-\sum_{i=1}^N \lambda_i T\right)$$

This is dependent only on the average default intensity. We should therefore expect the value of the first-to-default CDS to be independent of the distribution of default intensities. Exhibit 5 shows that this is what we find.

From the equations in Section II, the probability of one default by time T is:

$$\exp\left(-\sum_{i=1}^N \lambda_i T\right) \sum_{i=1}^N (e^{\lambda_i T} - 1)$$

Because of the convexity of the exponential function:

$$\sum_{i=1}^N (e^{\lambda_i T} - 1) > N(e^{\bar{\lambda} T} - 1)$$

dispersion in the default intensities increases the probability of exactly one default occurring by time T . The probability that the second default occurs before time T is therefore reduced. The value of the second-to-default should therefore decline. Again this is what we find. Similarly, the value of the n -th to default where $n > 2$ also declines.

Exhibit 5 also considers the circumstance that all pairs of firms have a correlation of 0.3. In this case, allowing each firm to have a different default probability while keeping the average default probability constant increases the cost of default protection over the base case.

To understand this, consider a pairwise correlation of 1. In this case, there is only a single value of x for all firms. This value of x is mapped into ten possible default times for the ten firms. The first of these default times is always the time for the firm with the highest default intensity. So, as we spread the default intensities while keeping the average intensity the same, the first default (for any value of x) becomes earlier than when the intensities are all the same. As a result, the first-to-default protection becomes more valuable.

When the correlation is less than perfect, this effect is still present but is more muted. The effect of dispersion in the default intensities on n -th to default swaps where $n > 1$ is a combination of two effects. The correlation makes it more likely that the n -th default will occur by time T when there is dispersion. The convexity of the exponential function makes it less likely that this will happen. In our example, the second effect is greater.

Impact of Dispersion in Pairwise Correlations

We now consider the valuation of a CDS when each firm has a different coefficient, a_i , in the single-factor model. The coefficients vary linearly across firms, but are chosen so that the average pairwise correlation is 0.30.

EXHIBIT 6

Effect of a Relationship Between Probability of Default and Pairwise Correlation

n	Base Case	Case 1	Case 2	Case 3
1	440	436	418	460
2	139	135	140	129
3	53	54	59	48
4	21	23	26	20
5	8	10	11	8
6	3	4	4	3
7	1	3	3	3
8	0	3	3	0
9	0	3	0	0
10	0	3	0	0

The average default probability and the average correlation are kept constant. In the base case the default intensity is 1% for each firm and all correlations are 0.30. In cases 1 to 3 the factor weights generating correlations vary linearly from 0.30 to 0.7995. In case 1 the default intensity is 1% for each firm. In cases 2 and 3 the default intensity varies linearly from 0.0055 to 0.0145. In case 2 the relation between default intensity and factor weight is positive while in case 3 it is negative.

Three cases other than the base case are considered:

Case 1: $\lambda_i = 0.01$ $a_i = 0.30 + 0.0555(i-1)$ $i = 1, \dots, 10$
Case 2: $\lambda_i = 0.0055 + 0.001(i-1)$ $a_i = 0.30 + 0.0555(i-1)$ $i = 1, \dots, 10$
Case 3: $\lambda_i = 0.0145 - 0.001(i-1)$ $a_i = 0.30 + 0.0555(i-1)$ $i = 1, \dots, 10$

In cases 2 and 3, both default intensities and correlations vary across firms. In case 2, the default intensities and correlations are positively related, while in case 3 the relation is negative. The results are shown in Exhibit 6.

Building dispersion into the pairwise correlations while holding default intensities constant has a modest effect on the cost of protection for first- and second-to-default swaps but greatly increases the cost of protection for eighth- to tenth-to-default swaps. When the correlations are correlated with the default probabilities, we observe very large changes in the cost of protection. The changes are similar to those observed when we move from a normal distribution to t -distributions with few degrees of freedom for M and the Z .

When high-default probability firms have high correlations (case 2), the cost of n -th to default protection is sharply reduced relative to the base case for $n = 1$ while higher for $n > 1$. When high-default probability firms have low correlations (case 3), the cost of n -th to default protection is increased for low and high n , while it is lower for intermediate values of n .

EXHIBIT 7

Effect of a Two-Factor Model on Spreads

<i>n</i>	Base Case	Case 1	Case 2	Case 3
1	440	392	386	401
2	139	151	151	150
3	53	68	69	65
4	21	30	30	27
5	8	11	11	9
6	3	2	2	2
7	1	1	0	1
8	0	1	0	1
9	0	0	0	0
10	0	0	0	0

The average default probability and the average correlation are kept constant. In the base case the default intensity is 1% for each firm and all pairwise correlations are 0.30. In cases 1 to 3 there are two sectors. The factor weights are chosen so that pairwise correlations are 0.6 for companies in the same sector and zero for companies in different sectors. In case 1 the default intensity is 1% for each firm. In case 2 companies in one sector have a default intensity of 0.5% while those in the other sector have a default intensity of 1.5%. In case 3 the default intensity in each sector varies linearly from 0.5% to 1.5%.

Two Factors

In Exhibit 7 we investigate the effect of using a two-factor model. We maintain the average correlation at 0.3 and the average default intensity at 1%. In case 1, there are two sectors each with five entities, each with a default intensity of 1%. The pairwise correlations within a sector are 0.6, and the pairwise correlations between sectors are zero. Because there are five entities in each sector, the impact of moving from the base case to case 1 when fewer than five defaults are considered is similar to the effect of increasing the correlation in the base case. For more than five defaults, we need entities from both sectors to default, so the impact of the two sectors is more complex.

In case 2, one sector has a default intensity of 1.5%, and the other has a default intensity of 0.5%. This produces results very similar to case 1 results. In case 3, the default intensity for each sector varies linearly from 0.5% to 1.5%. Here the results are similar to those for case 1, but the difference from the base case is less pronounced. (See Exhibit 5 for the 0.3 correlation.)

V. RESULTS FOR A CDO

The approach in Appendix A can be used to value a CDO when the principal amounts associated with all the underlying reference entities are the same. The recovery

rates must be non-stochastic and the same.

Consider, for example, the tranche responsible for between 5% and 15% of losses in a 100-name CDO. Suppose that the recovery rate is 40%. This tranche bears 66.67% of the cost of the 9th default, and all of the costs of the 10th, 11th, 12th, ..., and 25th defaults. The cost of defaults is therefore 66.67% of the cost of a 9th to default CDS plus the sum of the costs of the *n*-th to default for all values of *n* between 10 and 25, inclusive. Assume that the principal of each entity is *L* and there is a promised percentage payment of *r* at time *τ*. The expected payment in this case is:

$$10Lr \sum_{k=0}^8 \pi_{\tau}(k) + (10 - 0.6667 \times 0.6)Lr\pi_{\tau}(9) + \\ (10 - 1.6667 \times 0.6)Lr\pi_{\tau}(10) + \dots + \\ (10 - 15.6667 \times 0.6)Lr\pi_{\tau}(24)$$

The approach in Appendix B can be used in a more general set of circumstances. The principal amounts for the underlying names can be different. Also, there can be a different probability distribution for the recovery rate for each name. Furthermore, the recovery rate and factor loadings can be dependent on the factor values.

Cash versus Synthetic CDOs

The valuation approaches for a cash CDO and a synthetic CDO are very similar. In a cash CDO, the tranche holder has made an initial investment, and the valuation procedure calculates the current value of the investment (which can never be negative). In a synthetic CDO, there is no initial investment, and the value of a tranche can be positive or negative.

If we assume that interest rates are constant, the value of a cash CDO tranche is the value of the corresponding synthetic CDO tranche plus the remaining principal of the tranche. The break-even rate for a tranche in a new cash CDO is the risk-free zero rate plus the break-even rate for a tranche in the corresponding CDO.

To see that these results are true, we note that under constant interest rates a cash CDO tranche is the same as a synthetic CDO tranche plus a cash amount equal to the remaining principal of the tranche. As defaults occur, the synthetic tranche holder pays for them out of the cash.

The cash balance at any given time is invested at the risk-free rate. Losses reduce both the principal to which the synthetic CDO spread is applied and the cash balance. The

total income from the synthetic CDO plus the cash is therefore the same as that on the corresponding cash CDO.

Numerical Results

The break-even rate for a synthetic CDO is the payment that makes the present value of the expected cost of defaults equal to the present value of the expected income. The break-even promised payments (per year) for alternative tranches for a 100-name synthetic CDO for a range of model assumptions are shown in Exhibit 8. Payments are assumed to be made quarterly in arrears. The recovery rate is assumed to be 40%, and the default probabilities for the 100 entities are generated by Poisson processes with constant default intensities set to 1% per year. The term structure of interest rates is flat at 5%. The parameters n_M and n_Z are the degrees of freedom used in the t -distribution for M and Z_i in Equation (1).

The results in Exhibit 8 are consistent with the CDS results reported in Exhibits 2 through 6. Increasing the correlations reduces the value and the break-even rate for the junior tranches that bear the initial losses and increases the break-even rate for the senior tranches that bear the later losses. Making the tails of the M distribution heavier has the same effect as increasing the correlation, while making the tails of the Z distribution heavier generally has the opposite effect.

The double t -distribution copula has the same sort of effect as for the n -th to default. The break-even spreads for the most junior and senior tranches increase, while those for intermediate tranches decline.

Note that for low-risk (senior) tranches, increasing the size of the tranche narrows the spread paid on the tranche. This is because increasing the tranche size does not materially increase the number of defaults that are likely to be incurred, but it does increase the notional amount on which the payments are based. As a result, the spread paid on the tranche is approximately proportional to the inverse of the size of the tranche.

For example, in the 0.1 correlation case in Exhibit 8, tranches 10% to 15%, 10% to 20%, and 10% to 30% have sizes of 5%, 10%, and 20%, respectively. The spreads for the three tranches are 11, 6, and 3, almost exactly inversely proportional to the size of the tranche. (The probability of losses totaling more than 15% in this case is close to zero.)

In Exhibit 9 we consider the effect of moving to two sectors. The default intensities for all entities are 1%, and the average correlation is 0.30 for companies in the

EXHIBIT 8

Break-Even Spread on Various Tranches in a 100-Name Synthetic CDO

Correlation n_M / n_Z	0.1 ∞ / ∞	0.3 ∞ / ∞	0.3 $\infty / 5$	0.3 $5 / \infty$	0.3 $5 / 5$
Tranche (%)	Break-Even Tranche Spread (basis points per year)				
0 to 3	2279	1487	1766	1444	1713
3 to 6	450	472	420	408	359
6 to 10	89	203	161	171	136
10 to 100	1	7	6	10	9

same sector and zero otherwise. The 100 names are divided into two sectors, not necessarily equal in size. The results are consistent with those for case 1 in Exhibit 7. The impact of the two-factor model is to reduce the break-even spread for the very junior tranches and increase it for the more senior ones.

VI. CORRELATION BETWEEN DEFAULT RATE AND RECOVERY RATE

Altman et al. [2002] show that recovery rates tend to be negatively correlated with default rates. Cantor, Hamilton, and Ou [2002, p. 19] estimate the negative correlation to be about -0.67 for speculative-grade issuers. The phenomenon is quite marked. For example, in 2001 the annual default rate was about 10% and the recovery rate was about 20%; in 1997 the annual default rate was about 2% and the recovery rate was about 55%.

In the one-factor version of our model, the level of defaults by time T is measured by the factor M . The lower the value of M , the earlier defaults occur. We model the dependence between the recovery rate, R , and the level of defaults by letting R be positively dependent on M . We use a copula model to define the nature of the dependence. The math is similar to that in Section I.

Define a random variable, x_R :

$$x_R = a_R M + \sqrt{1 - a_R^2} Z_R$$

where $-1 < a_R < 1$, and Z_R has a zero-mean, unit-variance, distribution that is independent of M . The copula model maps x_R to the probability distribution of the recovery rate on a percentile-to-percentile basis. If H_R is the probability distribution for Z_R , F_R is the unconditional probability distribution for x_R , and Q_R is the unconditional probability for R , then:


EXHIBIT 9

Effect of Two-Sector Model on Break-Even Spread on Various Tranches

Sector Size	100 / 0	50 / 50	75 / 25
n_M / n_Z	∞ / ∞	∞ / ∞	∞ / ∞
Tranche (%)	Break-even Tranche Spread (basis points per year)		
0 to 3	1487	1161	1352
3 to 6	472	475	455
6 to 10	203	238	213
10 to 100	7	11	9

EXHIBIT 10

Relationship Between Default Probability and Expected Recovery Rate



$$\text{Prob}(R < R^* | M) = H_R \left\{ \frac{F_R^{-1}[Q_R(R)] - a_R M}{\sqrt{1 - a_R^2}} \right\} \quad (6)$$

Exhibit 10 shows the relationship between the expected recovery rate and the expected default rate when we use parameters similar to those observed by Cantor, Hamilton, and Ou for speculative-grade issuers in the copula model. The nature of the relationship is quite similar to that reported by Cantor, Hamilton, and Ou (see their Exhibit 21).

The procedure in Appendix B can be extended to accommodate a model like ours when the recovery rate (assumed to be the same for all companies) and the value

of M are correlated. When a value of M is chosen, we first use Equation (6) to determine the conditional probability distribution for R . We then proceed as described in Appendix B.

Exhibit 11 shows the impact of a stochastic recovery rate on the break-even spread for tranches in a CDO when the recovery rate is assumed to have a trinomial distribution with probabilities assigned to recovery rates of 0.25, 0.50, and 0.75 and an average recovery rate of 0.50. When there is no correlation between the value of the factor levels and the recovery rate, a stochastic recovery rate has little impact. When the two are correlated, the impact is significant, particularly for senior tranches. Without the correlation, these tranches are relatively safe. With the correlation, they are vulnerable to a bad (low M) year when probabilities of default are high and recovery rates are low.

When recovery rates are correlated with the probability of loss, the expected loss is increased if the default intensity is the same as in the uncorrelated case. As a result, the break-even rate for every tranche is increased, and senior tranches are more seriously affected. This phenomenon is shown in the second to last column of Exhibit 11.

To adjust for the change in expected loss when the recovery rate is correlated with default probabilities, we reduce the default intensity to a level at which the break-even spread on a single-name CDS is the same as in the uncorrelated case. We then recalculate the break-even spread for every tranche of the CDO. The results are in the final column of Exhibit 11. The break-even spread for the lowest-quality tranches is reduced from the zero-correlation case, while that for the highest-quality tranches is increased.¹⁰

VII. DETERMINING PARAMETERS AND MARKET PRACTICE

A model for valuing a CDO or an n -th to default CDS requires estimation of many parameters. Recovery rates can be estimated from data published by rating agencies. The required risk-neutral default probabilities can be estimated from credit default swap spreads or bond prices using the

EXHIBIT 11

Effect of a Stochastic Recovery Rate

Tranche %	PD/RR corr. = -0.5		PD/RR corr. = -0.5 Same Default Intensity	PD/RR corr. = -0.5 Same CDS Spread
	Const RR	SD of RR = 0.2		
0 to 3	1401	1368	1403	1208
3 to 6	395	403	480	391
6 to 10	139	144	211	164
10 to 100	3	3	8	6

recovery rates. The copula default correlation between two companies is often assumed to be the same as the correlation between their equity returns. This means that the factor copula model is related to an equivalent market model. For the one-factor model in Equation (1), a_i is set equal to the correlation between the equity returns of company i and the returns from a well-diversified market index.¹¹ For the multifactor model in Equation (3), a multifactor market model with orthogonal factors would be used to generate the a_{ij} .

The standard market model has become a one-factor Gaussian copula model with constant pairwise correlations, constant CDS spreads, and constant default intensities for all companies in the reference portfolio. A single recovery rate of 40% is assumed. This simplifies the calculations because the probability of k or more defaults by time T conditional on the value of the factor M can be calculated from the properties of the binomial distribution. In Equation (1) the a_i are all the same and equal to $\sqrt{\rho}$ where ρ is the pairwise correlation.

It is becoming common practice for market participants to calculate implied correlations from the spreads at which tranches trade using the standard market model. (This is similar to the practice of calculating implied volatilities from option prices using the Black-Scholes model.) The implied correlation for a tranche is the correlation that causes the value of the tranche to be zero.

Sometimes base correlations are quoted instead of tranche correlations. Base correlations are the correlations that cause the total value of all tranches up to a certain point to be zero. For example, in the case of the DJ CDX IG NA five-year index, the 0% to 10% base correlation is the correlation that causes the sum of the values of the 0% to 3%, the over 3% to 7%, and the over 7% to 10% tranches to be zero.

VIII. MARKET DATA

Market data for the pricing of index tranches are beginning to be available. We will look at the Dow Jones CDX NA IG five-year and the Dow Jones iTraxx EUR five-year tranches on August 4, 2004. Exhibit 12 shows midmarket quotes collected by GFI, a credit derivatives broker, and tranche spreads calculated using the standard market model for different correlations. The CDX index level on August 4, 2004, was 63.25 basis points, and the iTraxx index level was 42 basis points. We assume a recovery rate of 40%, and estimate the swap zero curves on that day in the usual way.

Note that the first (equity) tranche is by convention

quoted in a different way from other tranches. The market quote of 41.8% for the CDX means that the protection provider receives 500 basis points per year on the outstanding principal plus an initial payment of 41.8% of the tranche principal. Similarly, the market quote of 27.6% for iTraxx means that the protection provider receives 500 basis points per year plus an initial payment of 27.6% of the principal.

Exhibit 12 shows that the spreads in the market are not consistent with the standard market model. Consider the correlation of 0.25. The standard market model comes close to giving the correct break-even spread for the 0%-3% and the over 15% to 30% tranche but produces a break-even spread that is too high for the other tranches. The break-even spread is particularly high for the over 3%-7% tranche.

We have considered a number of reasons the pricing given by the standard market model may be wrong. Most model changes that we have considered have the effect of either 1) reducing spreads for all tranches up to a certain level of seniority and increasing spreads for all tranches beyond that level of seniority, or 2) increasing spreads for all tranches up to a certain level of seniority and reducing spreads for all tranches beyond that level of seniority.

Examples of model changes that have this effect include:

1. Changing the pairwise correlation.
2. Making the tails of the distribution of M heavier or less heavy.
3. Making the tails of the distribution of the Z_i heavier or less heavy.
4. Allowing the recovery rate to be stochastic and correlated with M .
5. Adding a second factor.

To match market data, we require a model that increases break-even spreads for the equity and very senior tranches and reduces them for intermediate tranches. Of those we have looked at, the only model that does this is the double t -distribution copula where both M and the Z_i have heavier tails than the normal distribution. This model fits the data in Exhibit 12 quite well. This is illustrated in Exhibit 13, which shows model prices for the iTraxx data on August 4, 2004, when both M and Z_i had four degrees of freedom. (The fit with four degrees of freedom is slightly better than the fit with five degrees of freedom.)

A final point is that tranche implied correlations must be interpreted with care. For the equity tranche (the most risky tranche in a CDO, typically 0% to 3% of the notional),

EXHIBIT 12

Gaussian Copula Model Quotes for CDX and iTraxx Tranches—August 4, 2004

DJ CDX IG NA					
Tranche	0 - 3%	3 - 7%	7 - 10%	10 - 15%	15 - 30%
Market Quote	41.8%	347	135.5	47.5	14.5
Correlation	Gaussian Copula Model Quotes				
0.00	67.9%	251	1	0	0
0.05	59.5%	365	29	2	0
0.10	53.0%	418	76	13	0
0.15	47.5%	444	118	31	2
0.20	42.6%	455	151	51	6
0.25	38.2%	457	177	72	11
0.30	34.0%	453	198	89	18
0.40	26.3%	434	227	116	35
Tranche Implied Corr.	0.210	0.042	0.177	0.190	0.274
Base Implied Corr.	0.210	0.279	0.312	0.374	0.519

DJ iTraxx EUR					
Tranche	0 - 3%	3 - 6%	6 - 9%	9 - 12%	12 - 22%
Market Quote	27.6%	168	70	43	20
Correlation	Gaussian Copula Model Quotes				
0.00	44.3%	69	0	0	0
0.05	39.7%	161	10	1	0
0.10	35.4%	222	36	6	0
0.15	31.5%	258	64	18	2
0.20	27.9%	281	90	33	6
0.25	24.5%	294	110	49	11
0.30	21.2%	300	127	64	18
0.40	15.2%	299	151	86	34
Tranche Implied Corr.	0.204	0.055	0.161	0.233	0.312
Base Implied Corr.	0.204	0.288	0.337	0.369	0.448

higher implied correlation means lower value to someone buying protection. For the mezzanine tranche (the second-most risky tranche in the CDO, typically 3% to 6% or 3% to 7%), the value of the tranche is not particularly sensitive to correlation, and the relationship between correlation and break-even spread, as illustrated in Exhibit 12, may not be monotonic. For other tranches, higher implied correlation means higher value to someone buying protection.

Base implied correlations are even more difficult to interpret. For the equity tranche, the base correlation is the same as the implied correlation; higher implied correlation means lower value to someone buying protection.

Consider the calculation of the base correlation for the mezzanine 3%-7% tranche in the CDX case. This is the correlation that causes the sum of values of the 0% to 3% and the over 3% to 7% tranche to be zero. When the correlation equals 0.210, the 0% to 3% tranche has a zero

value and the over 3% to 7% tranche has a positive value to a buyer of protection. Increasing the correlation reduces the value of the 0% to 3% tranche to a buyer of protection and increases the value of the over 3% to 7% tranche. The 0% to 3% tranche is much more sensitive to correlation than the over 3% to 7% tranche. As the correlation increases from 0.210, the total value of the two tranches therefore declines. When the correlation reaches 0.279, the total value of the two tranches is reduced to zero. Similar arguments explain why base correlations continue to increase as we move to more senior tranches.

It is evident from this that implied correlations, particularly base correlations, are not at all intuitive. On August 4, 2004, a correlation smile for tranche implied correlations translates into a steeply upward-sloping skew for base implied correlations.

EXHIBIT 13

Double t-Distribution Copula Model and Model Prices for iTraxx EUR Index Tranches—August 4, 2004

DJ iTraxx EUR					
Tranche	0 - 3%	3 - 6%	6 - 9%	9 - 12%	12 - 22%
Market Quote	27.6%	168	70	43	20
Correlation	Double-t Distribution Copula Model Quotes				
0.00	43.7%	66	0	0	0
0.05	41.0%	107	9	3	1
0.10	37.9%	133	23	10	4
0.15	34.8%	150	37	18	8
0.20	31.7%	161	49	26	13
0.25	28.6%	167	60	35	18
0.30	25.5%	171	69	42	23
0.40	19.5%	173	84	56	34
Tranche Implied Corr.	0.266	0.258	0.303	0.304	0.270
Base Implied Corr.	0.266	0.266	0.260	0.253	0.241

IX. CONCLUSIONS

We have presented two fast procedures for valuing an n -th to default CDS and a CDO tranche. The procedures (particularly the probability bucketing approach in Appendix B) are attractive alternatives to Monte Carlo simulation and have advantages over the fast Fourier transform approach.

In a general procedure for generating a wide range of different copulas, we find that the double t -distribution copula where both the market factor and the idiosyncratic factor have heavy tails provides a good fit to CDX and iTraxx market data.

The implied correlations now being reported by dealers and brokers for index tranches are often higher than typical equity correlations, and can be very difficult to interpret. Implied correlations are typically not the same for all tranches. This leads to a *correlation smile* phenomenon.

APPENDIX A

Calculation of the Probability Distribution of the Time of the n -th Default

For any given set of numbers $c_1, c_2, \dots, c_N$, we define:

$$U_k(c_1, c_2, \dots, c_N) = \sum c_{z(1)} c_{z(2)} \dots c_{z(k)}$$

where $k < N$ and $\{z(1), z(2), \dots, z(k)\}$ is a set of k different integers chosen from $\{1, 2, \dots, N\}$ and the summation is taken over the:

$$\frac{N!}{k!(N-k)!}$$

different ways the integers can be chosen.

We also define

$$V_n(c_1, c_2, \dots, c_N) = \sum_{j=1}^N c_j^n$$

There is an easy-to-compute recurrence relationship for determining the U_k from the V_k . Dropping arguments, the recurrence relationship is

$$\begin{aligned} U_1 &= V_1 \\ 2U_2 &= V_1 U_1 - V_2 \\ 3U_3 &= V_1 U_2 - V_2 U_1 + V_3 \\ 4U_4 &= V_1 U_3 - V_2 U_2 + V_3 U_1 - V_4 \\ &\vdots \\ kU_k &= V_1 U_{k-1} - V_2 U_{k-2} + V_3 U_{k-3} - \dots + (-1)^k V_{k-1} U_1 + (-1)^{k+1} V_k \end{aligned}$$

With $c_i = w_i$, this recurrence relation allows calculation of the probabilities in Equation (5).

To prove the recurrence relationship we define $Y_{k,i}$ as the value of $U_k(c_1, c_2, \dots, c_N)$ when $c_i = 0$. We define:

$$X_{k,n} = \sum_{i=1}^N c_i^n Y_{k-1,i}$$

It follows that

$$V_n U_k = X_{k+1,n} + X_{k,n+1}$$

when $k > 1$ and

$$V_n U_1 = X_{2,n} + V_{n+1}$$

These results lead to

$$\begin{aligned} V_1 U_{k-1} - V_2 U_{k-2} + V_3 U_{k-3} - \dots + (-1)^k V_{k-1} U_1 + (-1)^{k+1} V_k \\ = (X_{k,1} + X_{k-1,2}) - (X_{k-1,2} + X_{k-2,3}) + \\ (X_{k-2,3} + X_{k-3,4}) - \dots + (-1)^k (X_{2,k-1} + V_k) + (-1)^{k+1} V_k \\ = X_{k,1} \end{aligned}$$

Because $X_{k,1} = kU_k$ it follows that

$$\begin{aligned} kU_k = V_1 U_{k-1} - V_2 U_{k-2} + \\ V_3 U_{k-3} - \dots + (-1)^k V_{k-1} U_1 + (-1)^{k+1} V_k \end{aligned}$$

This is the required relationship.

Note that:

$$U_k = c_1 c_2 \dots c_N U_{N-k} \left(\frac{1}{c_1}, \frac{1}{c_2}, \dots, \frac{1}{c_N} \right)$$

This is a useful result for calculating U_k for large k .

APPENDIX B

Probability Bucketing

In a second approach we build up the probability distribution of the loss by time T , conditional on the values of the factors $M_1, M_2, \dots, M_m$, one debt instrument at a time. It is not necessary for the principal amounts to be equal, and the recovery rates can be stochastic.

Consider first that the recovery rate is known, and suppose there are N debt instruments. We choose intervals or buckets $\{0, b_0\}, \{b_0, b_1\}, \dots, \{b_{K-1}, \infty\}$ for the loss distribution. $\{0, b_0\}$ is the 0-th bucket; $\{b_{k-1}, b_k\}$ is the k -th bucket ($1 \leq k \leq K-1$); and $\{b_{K-1}, \infty\}$ is the K -th bucket. Our objective is to estimate the probability that the total loss lies in the k -th bucket for all k .

In some circumstances, it is best to set $b_0 = 0$ and $b_k - b_{k-1} = u$ ($1 \leq k \leq K-1$) for some constant u . The first bucket then corresponds to a loss of zero, and the other buckets except for the final one have equal widths. In other circumstances, when we are interested in valuing only one tranche, it makes sense to use narrow buckets for losses corresponding to the tranche and wide buckets elsewhere.

For the purposes of this appendix we abbreviate $p_T(k|M_1, M_2, \dots, M_m)$, the conditional probability that the loss by time T will be in the k -th bucket, as p_k . Let A_k be the mean loss conditional that the loss is in the k -th bucket ($0 \leq k \leq K$). We calculate p_k and A_k iteratively by first assuming that there are no debt instruments, then assuming that there is only one debt instrument, then assuming that there are only two debt instruments, and so on. Our only assumption in the iterative procedure is that

all the probability associated with bucket k is concentrated at the current value of A_k . We find that in practice this assumption leads to accurate loss probability distributions.

When there are no debt instruments, we are certain there will be no loss. Hence $p_0 = 1$ and $p_k = 0$ for $k > 0$. Also $A_0 = 0$. The initial values A_k for $k > 0$ are not important, but for the sake of definiteness we can set $A_k = 0.5(b_{k-1} + b_k)$ for $1 \leq k \leq K-1$ and $A_K = b_{K-1}$.

Suppose we have calculated the p_k and A_k when the first $j-1$ debt instruments are considered. Suppose that the loss given default from the j -th debt instrument is L_j , and the probability of a default is α_j . Define $u(k)$ as the bucket including $A_k + L_j$ for $0 \leq k \leq K$. The impact of the j -th debt instrument is to move an amount of probability $p_k \alpha_j$ from bucket k to bucket $u(k)$ ($0 \leq k \leq K$). When $u(k) > k$, the updating formulas are:

$$\begin{aligned} p_k &= p_k^* - p_k^* \alpha_j \\ p_{u(k)} &= p_{u(k)}^* + p_k^* \alpha_j \\ A_k &= A_k^* \\ A_{u(k)} &= \frac{p_{u(k)}^* A_{u(k)}^* + p_k^* \alpha_j (A_k^* + L_k)}{p_{u(k)}^* + p_k^* \alpha_j} \end{aligned}$$

where $p_k^*, p_{u(k)}^*, A_k^*$ and $A_{u(k)}^*$ are the values of $p_k, p_{u(k)}, A_k$, and $A_{u(k)}$ before the probability shift is considered. When $u(k) = k$, the updating formulas are:

$$\begin{aligned} p_{u(k)} &= p_k^* \\ p_k &= p_k^* \\ A_k &= A_k^* + \alpha_j L_k \\ A_{u(k)} &= A_k^* + \alpha_j L_k \end{aligned}$$

When all N debt instruments have been considered, we obtain the total loss distribution. If the recovery rate for each debt instrument is stochastic, we must first discretize the recovery rate distribution. Then the j -th debt instrument has loss L_{ji} with probability α_{ji} ($i = 1, 2, 3, \dots$). The total loss probability is $\sum_i \alpha_{ji}$. The iterative procedure can easily be adapted to accommodate this. As each entity is considered, we shift probability mass to multiple future buckets rather than to a single future bucket.

The procedure calculates the impact on the distribution of losses of adding a company to the portfolio. We can analogously calculate the impact of removing a company from the portfolio. This is useful in the calculation of Greek letters. For example, to calculate the impact of increasing the probability of default for a company, we can remove the company from the portfolio, and then restore it at a higher default probability. This type of approach has been independently developed by Andersen, Sidenius, and Basu [2003].

It is worth noting that when debt instruments have dif-

ferent principal amounts or different recovery rates, the valuations and the Greek letters in our approach may depend on the sequence in which names are added to the loss distribution. (An exception occurs when the bucket width is constant and a divisor of potential losses.) Our tests show that the sequence in which names are added makes very little difference to the results, however.

ENDNOTES

The authors thank Michael Gibson, Dominic O'Kane, and Michael Walker for comments on earlier versions of the article.

¹An n -th to default swap is sometimes defined so that there is a payoff for the first n defaults rather than just for the n -th default. Also, sometimes the rate of payment is reduced as defaults occur.

²It has also been used outside the credit risk area by researchers such as Hull [1977] and Hull and White [1998].

³For notational convenience we assume that the Z_i are identically distributed.

⁴The integration can be accomplished in a fast and efficient way using an appropriate Gaussian quadrature.

⁵These numerical stability problems occur because very large and very small numbers are sometimes involved in the recurrence relationship calculations. A computer stores only a finite number of digits for each number. For example, when 16 digits are stored, if $X - Y$ is calculated where X and Y are both between 10^{20} and 10^{21} and have the same first 17 digits, the result will be unreliable.

⁶In FFT the number of buckets must be $2^N - 1$ for some integer N , but not all choices for N work well.

⁷For a discussion of the valuation of a regular CDS, see Hull and White [2000, 2003].

⁸The valuation methodology can be adjusted to accommodate variations on the basic n -th to default structure.

⁹Constant default intensities provide a convenient way to generate marginal default probabilities. Our valuation procedures can be used for any set of marginal default time distributions.

¹⁰Results similar to those in the final column of Exhibit 11 are produced if the recovery rate is constant at 0.5, the default intensity is 1% per year, and the factor weightings, a_i , are negatively related to M . This is similar to case 2 in Exhibit 6.

¹¹A variation on this procedure is to assume that a_i is proportional to the correlation between the return from company i 's stock price and the return from a market index, and then choose the (time-varying) constant of proportionality so that available market prices are matched as closely as possible.

REFERENCES

Altman, E.I., B. Brady, A. Resti, and A. Sironi. "The Link Between Default 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 to the Gaussian Copula: Random Recovery and Random Factor Loadings." Working paper, Bank of America, June 2004.

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

Cantor, R., D.T. Hamilton, and S. Ou. "Default and Recovery Rates of Corporate Bond Issuers." Moody's Investors Services, February 2002.

Duffie, D., and K. Singleton. "Modeling Term Structures of Defaultable Bonds." *The Review of Financial Studies*, 12 (1999), pp. 687-720.

Hull, J. "Dealing with Dependence in Risk Simulations." *Operational Research Quarterly*, 28, 1 (1977), pp. 201-218.

Hull, J., and A. White. "Valuing Credit Default Swap Options." *The Journal of Derivatives*, 10, 3 (Spring 2003) pp. 40-50.

———. "Valuing Credit Default Swaps: No Counterparty Default Risk." *The Journal of Derivatives*, Vol. 8, No. 1 (Fall 2000), pp. 29-40.

———. "Value at Risk when Daily Changes Are Not Normally Distributed." *The Journal of Derivatives*, Vol. 6, No. 1 (Fall 1998), pp. 5-19.

Laurent, J.-P., and J. Gregory. "Basket Default Swaps, CDOs and Factor Copulas." Working paper, ISFA Actuarial School, University of Lyon, 2003.

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

Merton, R.C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 29 (1974), pp. 449-470.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC.

The most recent two editions of this title are only ever available at <http://www.euromoneyplc.com>