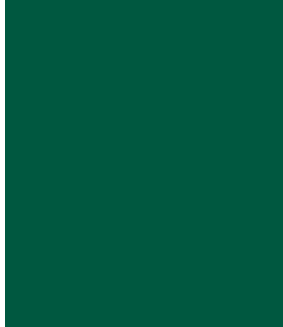
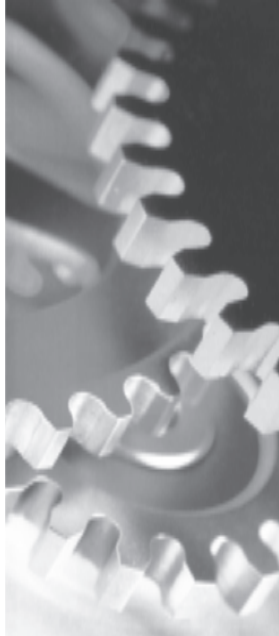


15 November 2004

Base Correlation Explained

Dominic O’Kane and Matthew Livesey

Since the advent of standardised single tranche CDOs on the liquid CDS indices of CDX and iTraxx, there has been a need for a commonly agreed method of quoting the implied correlation between the assets in the respective CDS index. Initially the market chose compound correlation as its quotation convention. More recently, base correlation has become more widely used. We define, discuss and compare both conventions. We conclude that base correlation possesses a number of desirable properties that make it a more powerful measure of tranche implied correlation. However, we argue that base correlation does not constitute a proper model for correlation skew.



Base Correlation Explained

Dominic O’Kane
+44 20 7102 2628
dokane@lehman.com

Matthew Livesey
+44 20 7102 5942
mlivesey@lehman.com

Since the advent of standardised single tranche CDOs on the liquid CDS indices of CDX and iTraxx, there has been a need for a commonly agreed method of quoting the implied correlation between the assets in the respective CDS index. Initially the market chose compound correlation as its quotation convention. More recently, base correlation has become more widely used in the market. In this paper we define, discuss and compare both conventions. We conclude that base correlation possesses a number of desirable properties that make it a more powerful measure of tranche implied correlation. However, we argue that base correlation does not constitute a proper model for correlation skew.¹

INTRODUCTION

The advent of standard CDO tranches with standard CDS indices as the reference portfolio has greatly enhanced liquidity and transparency in the synthetic CDO market. We are now able to observe daily pricing on a range of tranches linked to US, European and Japanese investment grade and high yield CDS indices. An example of tranche pricing on a selection of these indices is shown in Figure 1.

Figure 1. Indicative pricing for the five standard tranches linked to the CDX Investment Grade NA Series 3 and iTraxx Europe Series 2 indices, for 13 October 2004.²

Tranche	CDX Investment Grade North America Series 3		iTraxx Europe Series 2	
	Lower-Upper strike	Upfront / Running Spread (bp)	Lower-Upper Strike	Upfront / Running Spread (bp)
Equity	0-3%	37.125% + 500	0-3%	24.25% + 500
Junior Mezzanine	3-7%	259.5	3-6%	137.5
Senior Mezzanine	7-10%	101.0	6-9%	47.5
Senior	10-15%	38.5	9-12%	34.5
Super Senior	15-30%	11.5	12-22%	15.5

Note: On 13 October 2004, the CDX IG NA 3 index traded at 53.5bp and the iTraxx Europe 2 traded at 37bp. Both have a maturity date of 20 March 2010.

Source: Lehman Brothers

The price of a CDO tranche is a function of the default correlation between the assets in the reference portfolio. See O’Kane, Naldi *et al* [2003] for a discussion. An equity tranche investor can be shown to be long the default correlation between the credits in the underlying CDS index while a senior tranche investor is short this default correlation. Hence an equity tranche will increase in value and a senior tranche will fall in value if the default correlation of the underlying CDS index increases.

Before the advent of standard tranches, dealers looked to historical measures of default correlation. One widespread approach was to proxy the asset return correlation of latent variable models with the correlation of historical equity market returns. For a discussion of

¹ We would like to thank Wenjun Ruan, Saurav Sen, Minh Trinh and Lutz Schloegl for discussions and comments.

² Note that the convention for quoting prices is different for equity tranches. From Table 1 we can see that an investor who goes long the credit risk of the 0-3% equity tranche receives an upfront payment of 37.125 percent plus a running annual spread of 500bp. An investor who buys the 3-6% tranche receives an annualised spread of 259.5 bp (paid in quarterly instalments).

such models see O’Kane, Naldi *et al* [2003]. What has changed recently is that by observing the market prices of synthetic CDO tranches, we can begin to extract information about market-implied rather than historical default correlation.

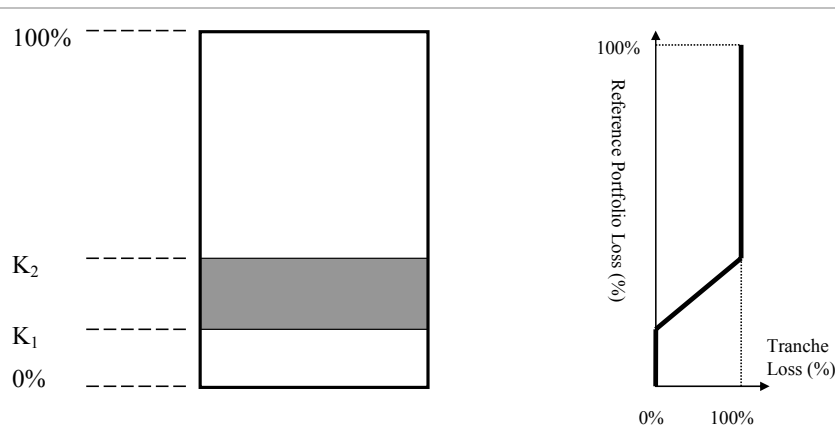
Initially, the market focused on compound correlation as the standard convention. More recently, dealers have begun to use base correlation instead. The main aim of this paper is to define, describe and compare these different measures of implied correlation. We begin with compound correlation.

COMPOUND CORRELATION

The first way to calculate an implied tranche correlation is to calculate the flat correlation that reprices each tranche to fit market prices. This method computes what is known as compound correlation.

We begin by defining a tranche simply in terms of its lower and upper “strikes”, which we denote by K_1 and K_2 . These are expressed as a percentage of the total notional of the reference portfolio. The lower strike is traditionally referred to as the tranche subordination or attachment point. The upper strike is referred to as the detachment point. The tranche loss is shown as a function of the percentage loss on the reference portfolio notional.

Figure 2. A mezzanine tranche with subordination (lower strike) K_1 and upper strike K_2 .



Source: Lehman Brothers

To calculate compound correlation we have to assume a mathematical framework for linking the defaults of all of the different assets in the reference collateral. The standard way of doing this is to use the Large Homogeneous Portfolio (LHP) model. For a derivation see Appendix A. The main modelling assumptions are:

1. The reference portfolio is homogeneous so that all assets share the same pairwise correlation, default probability and recovery rate.
2. The number of assets in the reference portfolio tends to infinity (see discussion below).
3. The default dependency structure is based on a Gaussian copula model.
4. Each tranche is priced off a single flat correlation (the compound correlation of the tranche).

Assumption (1) means that we model the actual reference portfolio as a portfolio of homogeneous assets each with the average spread and recovery rate of the actual reference portfolio. For standard tranches, it means that we assume that the spread and recovery rate of the individual names in the portfolio are the same as the index. This has the advantage that we do not need to exchange information about the individual CDS spreads and recovery rates of each name in the reference portfolio.

Assumption (2) means that the portfolio is infinitely granular so that all idiosyncratic risk has been diversified away. This has the advantage that it enables us to write a simple analytical expression for the tranche price and makes the calculation of the implied correlation very fast.

Assumptions (1), (3) and (4) taken together mean that the dependency structure for each tranche is characterised by a single correlation number. We can therefore solve for the compound correlation from one observed price.

Calculating Compound Correlation

Given a tranche denoted by its lower and upper strikes K_1 and K_2 , its present value today, time t , is given by:

$$PV_{tranche}(K_1, K_2, S_{K_1, K_2}, \rho_{K_1, K_2}) = U_{K_1, K_2} + S_{K_1, K_2} \sum_{n=1}^N Q_{K_1, K_2}(t_n) \Delta_n Z(t_n) - \sum_{m=1}^M (Q_{K_1, K_2}(t_{m-1}) - Q_{K_1, K_2}(t_m)) Z(t_m) \quad (1)$$

where

U_{K_1, K_2} is the tranche upfront payment,

S_{K_1, K_2} is the tranche contractual spread at issuance,

Δ_n is the accrual period between times t_{n-1} and t_n , usually paid quarterly, Actual 360,

$Z(t)$ is the LIBOR discount factor to time t .

The third term of equation (1) is the present value of the protection leg. Calculation of this involves an integration over time, which is usually discretised to quarterly time steps.

We define the tranche “survival probability” as follows:

$$Q_{K_1, K_2}(t) = 1 - \frac{\mathbf{E}_{\rho(K_1, K_2)}^{LHP} [\text{Min}(L(t), K_2)] - \mathbf{E}_{\rho(K_1, K_2)}^{LHP} [\text{Min}(L(t), K_1)]}{K_2 - K_1} \quad (2)$$

This survival probability is a measure of the expected percentage notional of the tranche remaining at some time t . The expectation is done using the Gaussian copula LHP model assuming a flat correlation as follows:

$$\mathbf{E}_{\rho}^{LHP} [\text{Min}(L(T), K)] = K\Phi(A) + (1 - R)N\Phi_{2, -\sqrt{\rho}}(C, -A) \quad (3)$$

where

N is the portfolio notional,

R is the average recovery rate of the reference portfolio,

$C = \Phi^{-1}(p(t))$ is the default threshold for the underlying reference portfolio,

$p(t)$ is the average cumulative default probability to time t of the issuers in the underlying reference portfolio ,

$\Phi(x)$ is the cumulative normal function,

$$A = \frac{1}{\sqrt{\rho}} \left(C - \sqrt{1-\rho} \cdot \Phi^{-1} \left(\frac{K}{N(1-R)} \right) \right)$$

ρ is the average pairwise asset correlation of the issuers in the reference portfolio,

$\Phi_{2,\rho}(x, y)$ is the cumulative bivariate normal with correlation coefficient ρ .

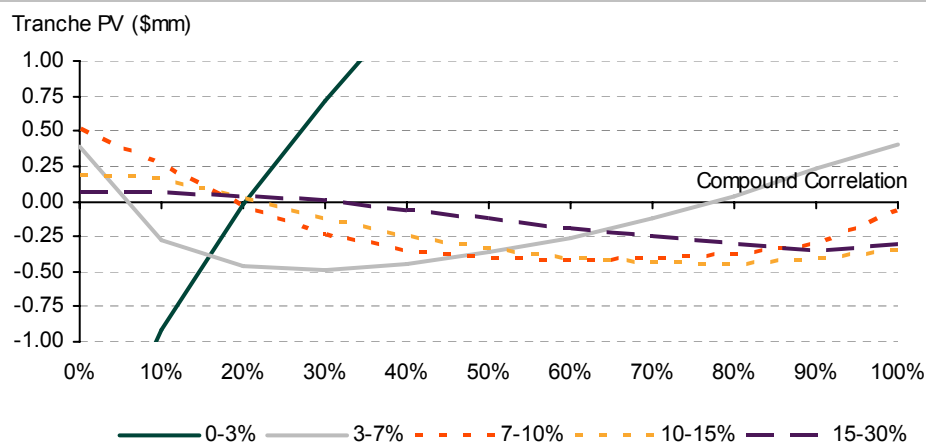
To calculate the compound correlation of a market tranche, we set the contractual spread equal to the observed market quote and, by definition, the present value of the tranche should be equal to zero:

$$PV_{tranche}(K_1, K_2, S_{K_1, K_2}, \rho_{K_1, K_2}) = 0$$

and we solve for $\rho = \rho_{K_1, K_2}$.

Solving this equation is straightforward, requiring a simple one-dimensional root searching algorithm. This works fine in almost all cases. However, there is sometimes a problem in that either we cannot find a root or that we get two roots. Why this is the case is shown in Figure 3 where we have plotted the present value of the five CDX tranches as a function of the compound correlation.

Figure 3. The present value of the five standard CDX tranches with different compound correlations – from the perspective of a protection seller (investor)



Source: Lehman Brothers

As expected, we see in Figure 3 that the equity tranche investor is long compound correlation while the senior tranche investor is short compound correlation.

We see that for all tranches, there is a single solution at which the PV is zero, except the 3-7% mezzanine which has two solutions at 5% and 78% compound correlation.

For mezzanine investors, the relationship between changes in the tranche PV and changes in correlation itself changes with correlation. At low correlations, mezzanine tranche investors are short correlation while at high correlations, mezzanine tranche investors are long correlation. Clearly, the two compound correlation solutions of 5% and 78%, while producing the same tranche PV, imply radically different risk profiles. Typically, we choose the lower correlation as it is closer to the other compound correlation solutions for adjacent tranches and because it better fits historical observations of equity return correlation which are widely used as a proxy for asset return correlation. See O’Kane, Naldi *et al* [2003] for a discussion.

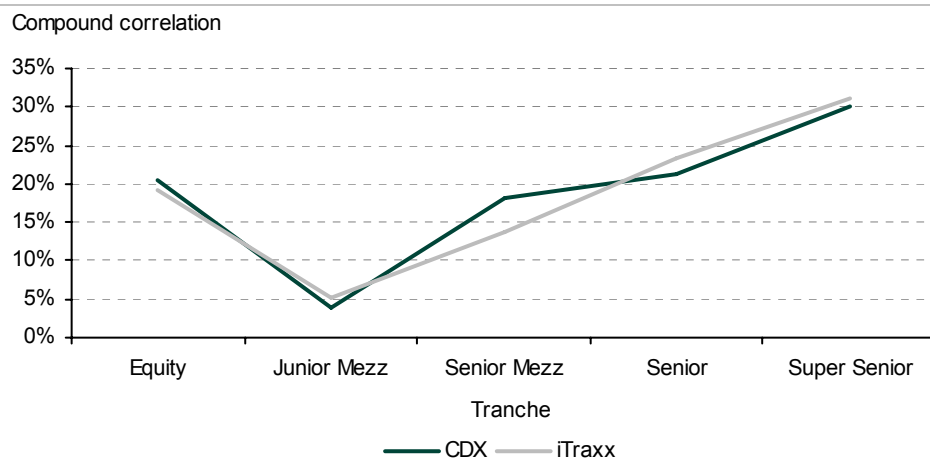
If we keep the reference portfolio spreads and recovery rates fixed and increase observed contractual tranche spreads, all of the curves in Figure 3 are shifted upwards. This causes the compound correlation of the equity tranche to decrease and the compound correlation of the other tranches to increase. If the tranche spreads are sufficiently large, there may not be a solution for the compound correlation of the mezzanine tranche. Equally, if tranche spreads fall, all of the curves in Figure 3 are shifted downwards. This can cause the mezzanine tranche to lose one, and ultimately both, of its solutions for compound correlation.

Explaining the Smile

The compound correlation curve is shown in Figure 4 for both the CDX and iTraxx tranches. The shape of the compound correlation has become known as the correlation “smile”. This is because the compound correlation is higher for the equity and senior tranches than it is for the mezzanine tranches.

What is interesting is that this smile shape is common to both CDX and iTraxx tranches and has persisted through the period of the last year during which these tranche prices have been available. Although the similarity in the actual values and the shapes is apparent, care must be taken when comparing CDX and iTraxx since the standard tranches have different attachment points and widths.

Figure 4. The compound correlation curve for CDX Investment Grade NA Series 3 and iTraxx Europe Series 2 indices, for 13 October 2004



Source: Lehman Brothers

Compound correlation is clearly not the same for all tranches. This simply tells us that a Gaussian copula does not capture the dependency structure implied by market CDO tranche prices. This is not a surprise – indeed, it would be amazing if we could exactly fit the market-

implied dependency structure of a portfolio of 125 different credits with a Gaussian copula characterised by a single correlation number.

What is interesting is the smile shape of the compound correlation for CDX IG and iTraxx tranches. This smile is driven by market prices – prices at which buyers and sellers of tranche protection are willing to trade. They therefore contain a mixture of effects, including concerns about systemic versus idiosyncratic credit risk, fear of principal versus mark-to-market losses, liquidity effects, and supply and demand for certain tranches.

Starting with the equity tranche, we note that the compound correlation is typically about 20%. This is actually lower than the 25-30% correlations found using historical equity returns, and since the equity tranche investor is long correlation, this means that the equity investor receives a higher spread than historical correlations would imply. One reason for this effect may be dealer correlation desks paying above the theoretical model price in order to hedge the risks in their correlation books created by selling mezzanine tranches.

At the mezzanine tranche, we see the compound correlation fall below the compound correlation of the equity tranche. As the mezzanine investor has a short correlation position this is simply reflecting the fact that the market is paying a lower spread than historical correlations would imply. This is probably due to the considerable demand for mezzanine tranches in the market. The size of the decline in the compound correlation to values in the range of 5-10% is due mainly to the low correlation sensitivity of the mezzanine tranche, ie, a large reduction in the compound correlation is required in order to fit this lower spread.

The senior tranche compound correlation is the simplest to explain. We see that it has a value similar to the historical average of 25-30%.

The CDO tranche market is segmented, with banks and hedge funds buying equity tranches, retail investors buying mezzanine tranches, and insurance companies focusing on senior tranches. This may explain why there has been little relative movement of tranche compound correlations, ie, few market players are willing or able to put on significant trades across the capital structure in an attempt to take advantage of any perceived richness or cheapness.

We list the advantages and disadvantages of compound correlation in Figure 5 below.

Figure 5. Advantages and disadvantages of compound correlation

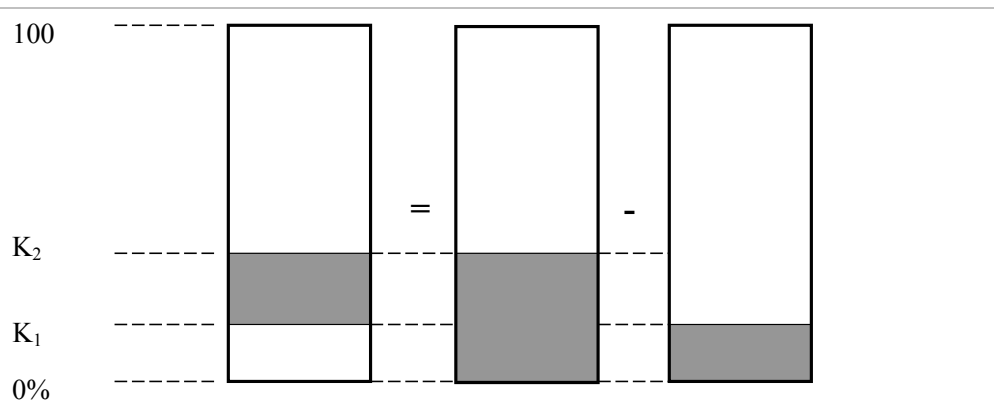
Advantages	Disadvantages
<ul style="list-style-type: none"> • Only one number per tranche. • Can be compared directly to estimates of historical asset correlation and can also be mapped easily to default correlation. Therefore it is quite intuitive. • We can calculate the compound correlation for a tranche without information about the pricing of the other tranches. 	<ul style="list-style-type: none"> • Sometimes there are two solutions for the mezzanine tranche. One must be chosen as the more economically sensible solution. • Not arbitrage-free across the capital structure, ie, the sum of the protection legs of the tranches does not equal the sum of the protection legs of the underlying CDS portfolio. • Not possible to extend compound correlation to the pricing of tranches on standard indices with non-standard strikes

BASE CORRELATION

The fundamental idea behind the concept of base correlation is that we decompose all tranches into combinations of base tranches, where **a base tranche is simply another name for an equity tranche**. The word “base” comes from the fact that the subordination of a base tranche is always zero, ie, it is attached to the base of the loss distribution.

Consider a mezzanine tranche with lower and upper attachment points K_1 and K_2 . This is equivalent to being short the equity tranche with subordination K_1 and long the equity tranche with subordination K_2 . This is shown in Figure 6.

Figure 6. A mezzanine tranche with strikes K_1 and K_2 decomposed into long a K_2 strike “base” tranche and short a K_1 strike “base” tranche



Source: Lehman Brothers

Whereas for compound correlation we calculate the flat correlation required for each tranche to match the market spreads, **for base correlation we value any tranche as the difference between two base tranches**. We then calculate the flat correlation required for each base tranche so that we match the observed market spreads.

Calculating Base Correlation

For a tranche with lower and upper attachment points K_1 and K_2 with a tranche market spread, S , we can compute the net present value of the protection and the premium legs. At inception, this must by definition equal 0.

We can do this by expressing the mezzanine tranche as the difference between two base tranches, as shown in Figure 6, where we allow each base tranche to be priced *using a different flat correlation*. This means that the base correlation is only a function of one parameter, the width of the base tranche. Compare this with compound correlation, which is therefore a tranche-specific function of both the lower and upper strike of the tranche.

We solve for the base correlation using a recursive technique called bootstrapping, ie, we use the information from the first tranche to solve for the second tranche, and so on. Contrast this with the case of compound correlation, where each tranche implied correlation is solved for independently of the other tranches. The procedure for base correlation is:

1. We solve for the equity tranche first by finding the value of ρ_{K_1} which solves the equation:

$$0 = PV_{tranche}(0, K_1, S_{0,K_1}, \rho_{K_1})$$

where

$$PV_{tranche}(0, K_1, S_{0,K_1}, \rho_{K_1}) = U_{0,K_1} + S_{0,K_1} PV_{premium}(0, K_1, \rho_{K_1}) - PV_{protection}(0, K_1, \rho_{K_1})$$

$$= U_{0,K_1} + S_{0,K_1} \sum_{n=1}^N Q_{0,K_1}(t_n) \Delta_n Z(t_n) - \sum_{m=1}^M (Q_{0,K_1}(t_{m-1}) - Q_{0,K_1}(t_m)) Z(t_m)$$

and the base tranche survival probability is given by:

$$Q_{0,K_1}(t) = 1 - \frac{\mathbf{E}_{\rho_{K_1}}^{LHP} [\text{Min}(L(t), K_1)]}{K_1}.$$

This is equivalent to what we did in the calculation of the compound correlation for the equity tranche. Indeed, the base and compound correlation measures are identical for an equity tranche.

2. For the next tranche, we solve for the value of ρ_{K_2} that solves:

$$\begin{aligned} 0 &= PV_{tranche}(K_1, K_2, S_{K_1, K_2}, \rho_{K_1}, \rho_{K_2}) \\ &= PV_{tranche}(0, K_2, S_{K_1, K_2}, \rho_{K_2}) - \underbrace{PV_{tranche}(0, K_1, S_{K_1, K_2}, \rho_{K_1})}_{\text{known and fixed as } \rho_{K_2} \text{ varies}} \end{aligned}$$

The second term is the PV of the first base tranche calculated using the spread from the second tranche. Breaking the tranche PV equation into the premium and protection legs, we have:

$$\begin{aligned} 0 &= (S_{K_1, K_2} PV_{premium}(0, K_2, \rho_{K_2}) - PV_{protection}(0, K_2, \rho_{K_2})) \\ &\quad - (S_{K_1, K_2} PV_{premium}(0, K_1, \rho_{K_1}) - PV_{protection}(0, K_1, \rho_{K_1})) \end{aligned}$$

which can be written in terms of tranche survival probabilities as before. We have:

$$S_{K_1, K_2} \sum_{n=1}^N Q_{K_1, K_2}(t_n) Z(t_n) - \sum_{m=1}^M (Q_{K_1, K_2}(t_{m-1}) - Q_{K_1, K_2}(t_m)) Z(t_m) = 0$$

where

$$Q_{K_1, K_2}(t) = 1 - \frac{\mathbf{E}_{\rho_{K_2}}^{LHP} [\text{Min}(L(t), K_2)] - \mathbf{E}_{\rho_{K_1}}^{LHP} [\text{Min}(L(t), K_1)]}{K_2 - K_1} \tag{4}$$

This tranche survival probability is fundamentally different from equation (2), the tranche survival probability for the compound correlation. In equation (2) we use the same correlation for both strikes. In equation (4) we take the expectation for the different base tranches at different correlations. The base correlation is linked to the strike of the base tranche. Since we already know ρ_{K_1} from the previous step, we have one equation with one unknown ρ_{K_2} and we can solve for this by using a one-dimensional root search.

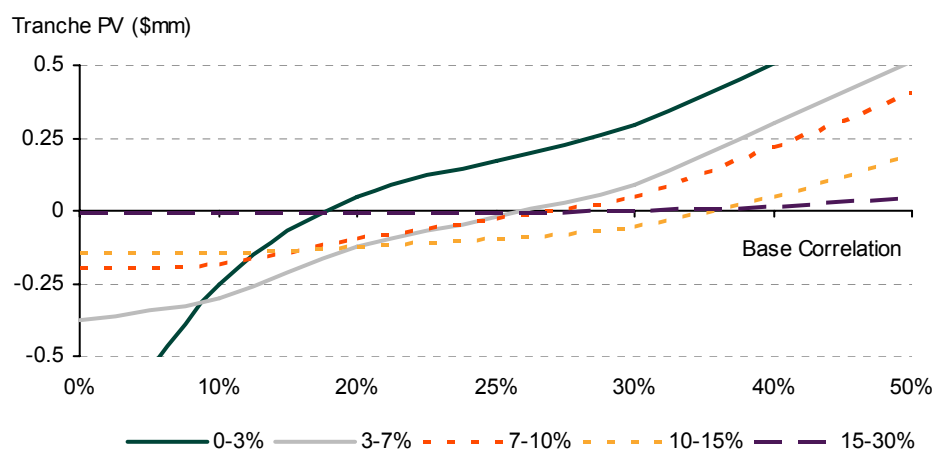
3. We then continue in the same manner through the higher tranches.

Figure 7 shows the tranche PVs for the five standard CDX tranches as a function of base correlation. To generate this graph, we first plotted the PV of the equity tranche by varying the 3% strike base correlation. We then selected the solution for the 3% strike base correlation at which the 0-3% tranche PV is zero. We then calculated the PV of the 3-7%

tranche using the solution for the 3% strike base correlation and for different values of the 7% strike base correlation. This produced the line for the 3-7% tranche. The solution for the 7% strike base correlation is the value that gives a 3-7% tranche PV of zero. We then plotted the 7-10% tranche PV using the solution for the 7% tranche by varying the 10% strike base correlation, and so on up the capital structure.

What we find is that all of the tranche PVs are a monotonic and increasing function of the base correlation. This means that there is only one solution, or in certain cases no solution. We therefore avoid the two-solution problem that we had when determining the compound correlation on the 3-6% mezzanine tranche.

Figure 7. PVs of the five standard CDX tranches as we change the base correlation of the upper base tranche, while using the correct value of base correlation for the lower base tranche



Source: Lehman Brothers

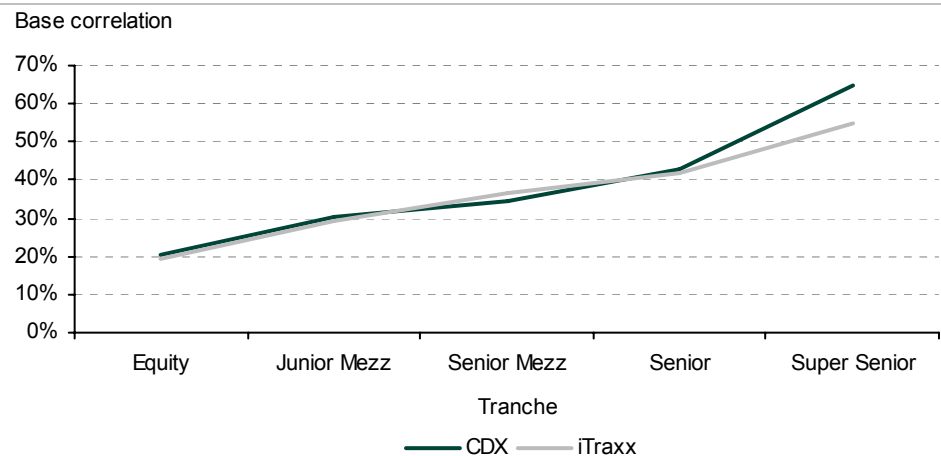
The reason why all the tranche PVs increase as base correlation increases is that for each tranche PV, we are fixing the lower strike base correlation and increasing only that of the upper strike. Hence the only changing component of the tranche PV is due to changes in the PV of the upper base tranche. As all base tranches are equity tranches which are long correlation, all the tranche PVs are increasing functions of their upper strike base correlation.

In some cases, base correlation can have trouble finding a solution for the senior tranche. This may reflect an inconsistency between the market spreads paid on the tranches and the spread paid on the underlying CDS index within the base correlation modelling framework. However it may also reflect a more serious violation of no-arbitrage constraints as discussed later.

Base correlation produces a skew

Figure 8 shows the base correlation calculated for the CDX and iTraxx linked tranches. As with compound correlation, we see considerable similarity between the indices, in both the shapes and levels of the implied base correlations. Instead of a smile shape, we get a “skew”. As a result, we find that people who prefer compound correlation speaking of a “smile” and those who favour base correlation speaking of a “skew”.

Figure 8. The base correlation curve for CDX Investment Grade NA Series 3 and iTraxx Europe Series 2 indices, for 13 October 2004 for each tranche



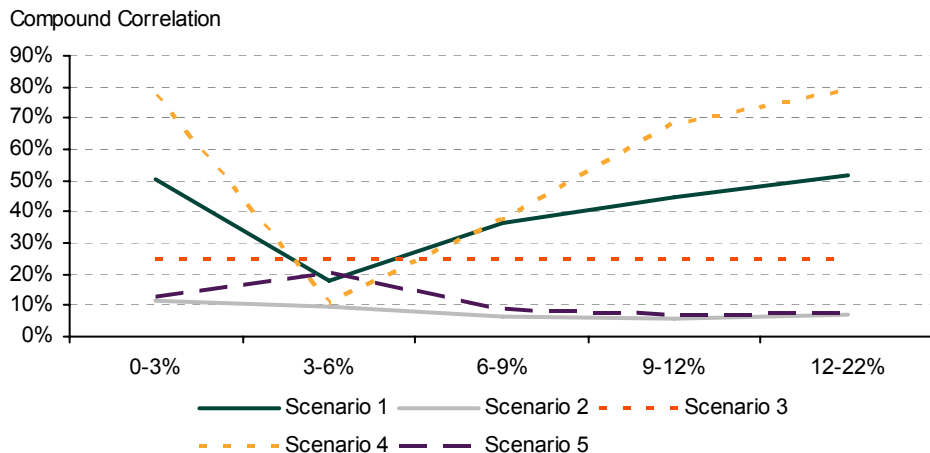
Source: Lehman Brothers

The base correlation and compound correlation for the 0-3% equity tranche are the same. This follows from the definition of both implied correlation measures.

The next tranche is the 3-7%. If we price the 3-7% junior mezzanine tranche with the 3% strike base correlation implied by the 0-3% tranche spread, then we find the tranche PV is generally negative, ie, the PV of the spread being received on the premium leg is not sufficient to cover the PV of the protection leg. This is due to the apparent low spread paid on mezzanine tranches, as discussed earlier. We can only reduce the value of the 3-7% protection leg by lowering the value of the 0-7% base tranche. As this is an equity tranche, we need to increase the 7% strike base correlation above the 3% strike base correlation. The base correlation for subsequent tranches tends to be higher still because it has to compensate for the previous high values of base correlation on which it depends. As a result, we have an upward sloping base correlation curve.

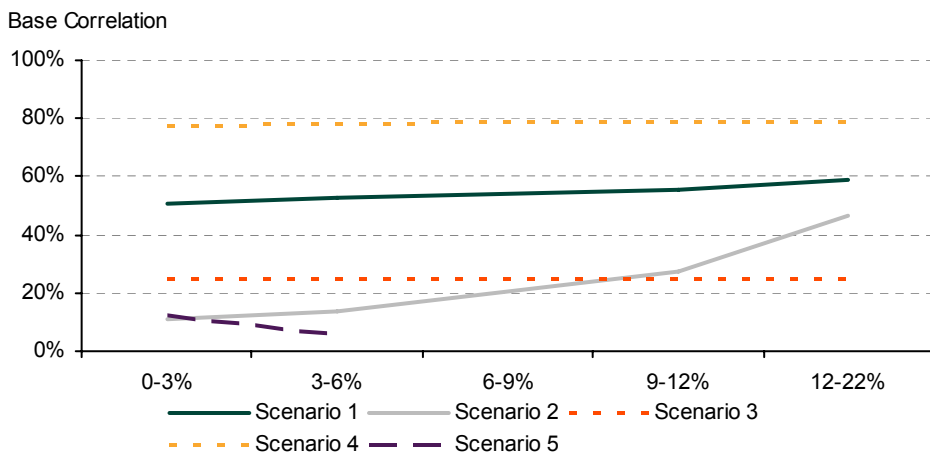
While Figure 8 shows the current shape of the base correlation curve in the market, other shapes are very possible. To demonstrate this, and the mapping to the equivalent compound correlation curve, we generated a number of different tranche spread scenarios. This was done in a way that guaranteed that the resultant spreads were arbitrage-free, see later for a discussion of arbitrage-free skew models. All these scenarios were created for a reference portfolio in which we held the spreads of all of the underlying assets fixed. By changing the underlying dependency structure, we were able to produce different shapes for compound and base correlation. Figures 9 and 10 show the corresponding compound and base correlation for the set of tranche spreads shown in Figure 11. This illustrates the relationship between the two and gives an intuitive feel for how base correlation varies depending on the shape of compound correlation and vice versa.

Figure 9. Compound correlation for the different tranche spread scenarios, based on the iTraxx Europe Series 2 index (see Figure 11 for details)



Source: Lehman Brothers

Figure 10. Base correlation for the different tranche spread scenarios, based on the iTraxx Europe Series 2 index (see Figure 11 for details)



Source: Lehman Brothers

Figure 11. The tranche spreads associated with the five different scenarios for iTraxx Europe Series 2, which lead to the different shapes for the compound and base correlation curves shown in Figures 9 and 10

Tranche	Lower-Upper strike	Running Spread (bp)				
		Scenario				
		1	2	3	4	5
Equity ³	0-3%	646	1396	1101	337	1360
Junior Mezzanine	3-6%	224	160	234	174	233
Senior Mezzanine	6-9%	123	8	86	126	18
Senior	9-12%	82	0.3	36	106	0.6
Super Senior	12-22%	44	0.02	8.2	72	0.03

Source: Lehman Brothers

While the values calculated for compound and base correlation for the equity tranche are the same, the values for the other tranches vary. For different tranches the change in base correlation tends to be in the opposite direction to that in compound correlation. Let us discuss the different scenarios:

1. The low equity spread and high senior spread mean there is a lot of “smile” in the compound correlation curve. The base correlation curve is upward sloping, starting at a higher correlation due to the low spread of the equity tranche.
2. The equity tranche spread has increased and senior spreads are close to zero. As a result the compound correlation for equity falls and the low senior spreads mean that it does not rise again. In base correlation, we see a steep upward sloping curve starting from a low level.
3. Both compound correlation and base correlation curves are flat as the tranche spreads were generated by the LHP model, using the same correlation for all tranches.
4. The compound correlation curve is extremely “smiley”. This is because of the extremely low equity spread combined with significant senior spreads. The base correlation curve has to start off at the high equity compound correlation. The slope of the curve is then squeezed into the remaining gap.
5. High equity spreads and high mezzanine spreads make the compound correlation curve inverted. Base correlation is unable to find a solution for the senior tranches, which illustrates how base correlation can fail to find a solution to reprice the senior tranches even though compound correlation can.

More information: each tranche is characterised by two correlations

The way that base correlation is defined means that the base correlation calculated for each tranche is linked to the base correlation of the tranche below. This is clear from the bootstrapping methodology that we have to employ to calculate base correlation. A 7-10% tranche “knows” the value of the base correlation we computed from the 3-7% tranche, which in turn “knows” the value of the base correlation we computed from the 0-3% tranche.

Indeed, the fact that each tranche survival probability curve is explicitly a function of two base correlations implies that each tranche “knows” more about the shape of the underlying market implied loss distribution than when using compound correlation.

³ Here we are quoting the equity tranche with purely a running spread and no upfront payment.

This does not mean that a base correlation curve contains more information than a compound correlation curve. Taking all points together, both actually contain the same amount of information, ie, the market prices of the different tranches. The point is that individually, each tranche is represented by two base correlation numbers, and so the modelling of that tranche embeds more information about the pricing of the other tranches than in the compound correlation framework.

Base correlation conserves expected loss and delta

A basic arbitrage-free requirement of any CDO pricing model should be that the sum of the protection legs of all of the CDO tranches should equal the sum of the protection legs of all of the underlying credit default swaps.

The easiest way to see why this is an arbitrage-free requirement is to consider what would happen if all the tranches and underlying CDS were quoted in upfront premium terms. A trade in which we sold protection on the whole capital structure of tranches and bought protection on each name in the reference portfolio would give us a risk-free position. Any default on a name in the reference portfolio would incur a loss on a tranche which would be exactly offset by a payment on the CDS linked to that name. Arbitrage requirements mean that the initial cost of this strategy should be zero. As the upfront price of a CDS or tranche is simply the present value of the protection leg, the result follows.

A simple way to show that this arbitrage-free requirement holds for base correlation is to take a simple capital structure consisting of a 0-5% equity tranche, 5-20% mezzanine tranche, and a 20-100% senior tranche. For simplicity we assume that interest rates are zero and that all losses are taken at the maturity date T. Hence, the expected loss on the individual tranches is given by:

$$EL(0,5\%) = E_{\rho(5\%)}^{LHP} [Min(L(T),5\%)] - E_{\rho(0\%)}^{LHP} [Min(L(T),0)]$$

$$EL(5\%,20\%) = E_{\rho(20\%)}^{LHP} [Min(L(T),20\%)] - E_{\rho(5\%)}^{LHP} [Min(L(T),5\%)]$$

$$EL(20\%,100\%) = E_{\rho(100\%)}^{LHP} [Min(L(T),100\%)] - E_{\rho(20\%)}^{LHP} [Min(L(T),20\%)]$$

Summing these, we find that all the intermediate terms cancel out and we get:

$$\begin{aligned} &EL(0,5\%) + EL(5\%,20\%) + EL(20\%,100\%) \\ &= E_{\rho(100\%)}^{LHP} [Min(L(T),100\%)] - E_{\rho(0\%)}^{LHP} [Min(L(T),0)] \end{aligned}$$

Since

$$0\% \leq L(T) \leq 100\% ,$$

we have

$$EL(0,5\%) + EL(5\%,20\%) + EL(20\%,100\%) = E[L(T)]$$

We see that the sum of the expected tranche losses is equal to the expected loss of the underlying portfolio which is not a correlation sensitive quantity. It can be shown that this result holds for base correlation even when interest rates are non-zero and losses are taken as they occur.

This important result would not hold for compound correlation because the strike of 5% would be priced at the 0-5% compound correlation in the equity tranche and at the 5-20%

compound correlation in the mezzanine tranche. Therefore, the intermediate terms with the 5% strike and the 20% strike would not cancel out and so the sum of the expected tranche losses would not be the same as the expected loss of the underlying portfolio.

The delta of a tranche to any credit is simply the sensitivity of the value of the tranche with respect to a shift in the CDS curve for that specific credit. Since the expected loss is conserved, we would also expect the delta to sum correctly as we look across the entire capital structure, ie, a dealer wishing to hedge a correlation book by building the capital structure around a sold tranche would have more confidence in the base correlation delta than the compound correlation delta.

Interpolating Non-Standard Strikes

Base correlation is a powerful concept because it transforms the compound correlation ρ_{K_1, K_2} , which is a two-dimensional correlation, into the one-dimensional base correlation measure ρ_{K_1} . This reduction of the dimensionality of the correlation parameters enables a simple mechanism for pricing non-standard strikes on the standard indices. Consider the following example.

Suppose we want to price a 6-9% tranche of the CDX portfolio. Using base correlation the price will be the difference between the price of a 0-9% tranche and that of a 0-6% tranche. These base tranches are priced using the base correlation for the 9% and 6% strikes. However, the market information only gives us the base correlations for the 3%, 7% and 10% strikes. It is worth noting that if we plot these base correlations as a function of the strikes, the resulting shape is monotonic, increasing with strike, and close to linear. This suggests that we may be able to interpolate values for the 6% and 9% strike with a reasonable degree of confidence.

For example, suppose the market implied that the base correlation for the 3% strike is 20%, for the 7% strike is 28% and for the 10% strike is 34%. Then using linear interpolation between these values, the base correlation for a 6% tranche is:

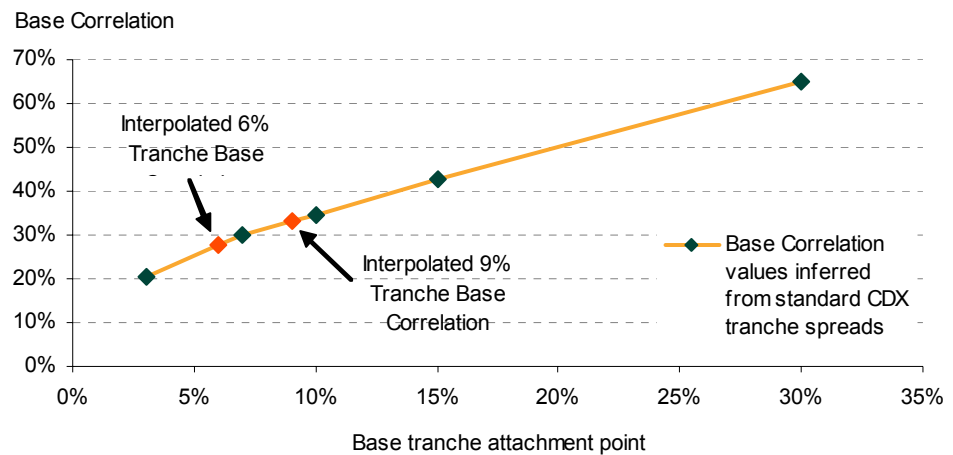
$$\frac{1}{4} \times 20\% + \frac{3}{4} \times 28\% = 26\%$$

and the base correlation for the 9% tranche is:

$$\frac{1}{3} \times 28\% + \frac{2}{3} \times 34\% = 32\%$$

We can then price the 6-9% CDX tranche. Figure 12 shows the values of base correlation that we infer from the market for CDX tranches on 13 October 2004 and the values of base correlation we would interpolate for 6% and 9%.

Figure 12. Base correlation for the five standard CDX tranches and the interpolated base correlations for a 6-9% tranche



Source: Lehman Brothers

The spread implied by these interpolated base correlations for the 6-9% tranche is 118bp, which compares with 260bp for the 3-7% tranche and 101bp for the 7-10% tranche.

Care must be taken to ensure that the interpolation methodology does not introduce any arbitrage.

Not a proper model of the correlation skew

Although base correlation clearly has more attractive properties than compound correlation in terms of its conservation of the expected losses and tranche deltas, it is not a proper model of the correlation skew. By “proper model”, we mean a model that allows us to price and risk-manage all of the tranches on the same CDS index using the same underlying portfolio loss distribution which has been generated in an arbitrage-free manner. An arbitrage-free model satisfies the following conditions:

1. The tranche survival probability, ie, the expected outstanding notional of a tranche must be a monotonically decreasing function of the horizon date. This must be true for all tranches.
2. The absolute value of the expected loss of an equity tranche must be a monotonically increasing function of the width of the equity tranche. This must be true at all horizon dates.
3. The expected loss of a tranche with strikes K_1 and K_2 plus the expected loss of a tranche with strikes K_2 and K_3 must equal the expected loss of a K_1 to K_3 tranche. This must be true at all horizon dates.

Provided that the same correlation is used for different time horizons, compound correlation guarantees that (1) holds, whereas base correlation does not. In practice, however, base correlation only violates (1) very slightly, even when the tranches are relatively thin.

Neither compound nor base correlation guarantees that (2) holds and so both fail the requirements for an arbitrage-free model.

However, base correlation does guarantee (3), which compound correlation does not. The fact that compound correlation violates (3) can lead to very strange looking deltas, which suggests that it is not a reliable method for pricing.

In addition to these arbitrage-free requirements, a proper model would also allow us to extend the pricing of standard tranches to the pricing of tranches on non-standard portfolios. For example, with base correlation it is not clear how one would price, say, a CDO on a mixed portfolio of CDX and iTraxx indices. Nor is it clear how base correlation should change as the spread of the CDS index changes. Capturing this cross-dynamic in line with empirical observations would definitely be a desirable feature of a proper model.

Since it satisfies requirement (3), we believe that base correlation is a safer approach than compound correlation. However, we do not believe that it fulfils the requirements of those looking for a model of the correlation skew. This remains an active area of research.

We list the advantages and disadvantages of base correlation in Figure 13.

Figure 13. Advantages and disadvantages of base correlation

Advantages	Disadvantages
<ul style="list-style-type: none"> • Each tranche is characterised by two base correlations. This embeds more information about the market-implied loss distribution. • There is always either one solution or no solution. The situation of having two solutions never arises. • Conserves expected loss and delta across the capital structure. • Extends to the pricing of tranches with non-standard strikes. 	<ul style="list-style-type: none"> • Calculating the base correlation for any tranche requires that the market prices for all of the more junior tranches are known. • Difficult to build intuition about base correlation and to relate it directly to historical estimates of correlation. • Not an arbitrage-free model of the correlation skew.

CONCLUSION

The Gaussian Copula LHP model has become the Black-Scholes model of the CDO tranche market. However, there are currently two ways of quoting the implied correlation – compound and base correlation. We have defined, discussed and compared these approaches, and shown that base correlation has a number of distinct advantages over compound correlation. These include the fact that it extends naturally to the pricing of non-standard tranches on the liquid indices, and correctly conserves the expected loss and delta across the capital structure. It is also simple to implement. However, we emphasise that base correlation is not a proper model of correlation skew.

REFERENCES

- Mashal, Naldi and Tejwani (2004), “The Implications of Implied Correlation”, Lehman Brothers Fixed Income Research.
- O’Kane, Naldi, Ganapati, Berd, Pedersen, Schloegl, and Mashal (2003), “The Lehman Brothers Guide to Exotic Credit Derivatives”, *Risk*, October.
- O’Kane and Schloegl (2001), “Modelling Credit: Theory and Practice”, Lehman Brothers Fixed Income Research.

APPENDIX A: LHP FORMULA FOR PRICING A CDO TRANCHE

Under the Gaussian copula LHP model, the assets of the n issuers are modelled by standard normal random variables with a common correlation ρ , probability of default p , notional N and recovery rate R . Default is said to occur if the asset value of an issuer i , Z_i falls below the default threshold C , which is given by $C = \Phi^{-1}(p)$. Under this model we can write the asset value of each issuer as a market factor and a specific component, as follows:

$$Z_i = \sqrt{\rho} Z + \sqrt{1 - \rho} \varepsilon_i$$

where Z and the ε_i s are independent standard normal random variables. We can then write the conditional probability of asset i defaulting as:

$$\Pi(Z) = \Phi\left(\frac{C - \sqrt{\rho}Z}{\sqrt{1 - \rho}}\right)$$

The conditional loss will therefore be the sum of n independent identically distributed random variables with an expected value of $\Pi(Z)(1 - R)N$. Under the LHP model, the assumption is that the number of issuers is sufficiently large so that the law of large numbers causes the conditional loss to be exactly $\Pi(Z)(1 - R)N$. This is equivalent to saying that all the specific risk of default has been diversified away.

So under the Gaussian copula LHP model, the probability of the loss being greater than some level K is:

$$P[L > K] = \mathbf{E}[1_{\{\Pi(Z)(1-R)N > K\}} | Z] = P[Z < A] = \Phi(A)$$

where

$$A = \frac{1}{\sqrt{\rho}} \left(C - \sqrt{1 - \rho} \cdot \Phi^{-1}\left(\frac{K}{N(1 - R)}\right) \right)$$

Furthermore, we can calculate $\mathbf{E}[\min(L, K)]$ in a similar way:

$$\mathbf{E}[\min(L, K)] = \mathbf{E}[K 1_{\{L > K\}} + L 1_{\{L < K\}}] = K\Phi(A) + \mathbf{E}[L 1_{\{L < K\}}]$$

and

$$\begin{aligned} \mathbf{E}[L 1_{\{L < K\}}] &= \mathbf{E}[\mathbf{E}[L 1_{\{L < K\}} | Z]] \\ &= \mathbf{E}[\mathbf{E}[\Pi(Z)(1 - R)N \cdot 1_{\{L < K\}} | Z]] \end{aligned}$$

$$\begin{aligned}
&= (1-R)N \mathbf{E}[\mathbf{E}[\Pi(Z)1_{\{Z>A\}} | Z]] \\
&= (1-R)N \int_A^\infty \Phi\left(\frac{C-\sqrt{\rho}z}{\sqrt{1-\rho}}\right)\phi(z)dz \\
&= (1-R)N \int_\infty^{-A} \Phi\left(\frac{C-(-\sqrt{\rho})z}{\sqrt{1-\rho}}\right)\phi(z)dz \\
&= (1-R)N\Phi_{2,-\sqrt{\rho}}(C,-A)
\end{aligned}$$

So we have:

$$\mathbf{E}[\min(L, K)] = K\Phi(A) + (1-R)N\Phi_{2,-\sqrt{\rho}}(C,-A)$$

This allows us to calculate the tranche survival probability:

$$Q_{K_1, K_2}(t) = 1 - \frac{\mathbf{E}[\text{Min}(L(t), K_2)] - \mathbf{E}[\text{Min}(L(t), K_1)]}{K_2 - K_1}$$

and the PV of a CDO tranche, which is:

$$V_{K_1, K_2}(t) = S_{K_1, K_2} \sum_{n=1}^N Q_{K_1, K_2}(t_n) \delta_n Z(t_n) - \sum_{m=1}^M (Q_{K_1, K_2}(t_{m-1}) - Q_{K_1, K_2}(t_m)) Z(t_m)$$

where

K_1 and K_2 are the attachment point and detachment point of the tranche,

S_{K_1, K_2} is the tranche contractual spread at issuance,

δ_n is the accrual period between times t_{n-1} and t_n ,

$Z(t)$ is the LIBOR discount factor to time t .

The views expressed in this report accurately reflect the personal views of Dominic O'Kane and Matthew Livesey, the primary analyst(s) responsible for this report, about the subject securities or issuers referred to herein, and no part of such analyst(s)' compensation was, is or will be directly or indirectly related to the specific recommendations or views expressed herein.

Any reports referenced herein published after 14 April 2003 have been certified in accordance with Regulation AC. To obtain copies of these reports and their certifications, please contact Larry Pindyck (lpindyck@lehman.com; 212-526-6268) or Valerie Monchi (vmonchi@lehman.com; 44-(0)207-102-8035).

Lehman Brothers Inc. and any affiliate may have a position in the instruments or the Company discussed in this report. The Firm's interests may conflict with the interests of an investor in those instruments.

The research analysts responsible for preparing this report receive compensation based upon various factors, including, among other things, the quality of their work, firm revenues, including trading, competitive factors and client feedback.

This material has been prepared and/or issued by Lehman Brothers Inc., member SIPC, and/or one of its affiliates ("Lehman Brothers") and has been approved by Lehman Brothers International (Europe), authorised and regulated by the Financial Services Authority, in connection with its distribution in the European Economic Area. This material is distributed in Japan by Lehman Brothers Japan Inc., and in Hong Kong by Lehman Brothers Asia Limited. This material is distributed in Australia by Lehman Brothers Australia Pty Limited, and in Singapore by Lehman Brothers Inc., Singapore Branch. This material is distributed in Korea by Lehman Brothers International (Europe) Seoul Branch. This document is for information purposes only and it should not be regarded as an offer to sell or as a solicitation of an offer to buy the securities or other instruments mentioned in it. No part of this document may be reproduced in any manner without the written permission of Lehman Brothers. We do not represent that this information, including any third party information, is accurate or complete and it should not be relied upon as such. It is provided with the understanding that Lehman Brothers is not acting in a fiduciary capacity. Opinions expressed herein reflect the opinion of Lehman Brothers and are subject to change without notice. The products mentioned in this document may not be eligible for sale in some states or countries, and they may not be suitable for all types of investors. If an investor has any doubts about product suitability, he should consult his Lehman Brothers representative. The value of and the income produced by products may fluctuate, so that an investor may get back less than he invested. Value and income may be adversely affected by exchange rates, interest rates, or other factors. Past performance is not necessarily indicative of future results. If a product is income producing, part of the capital invested may be used to pay that income. Lehman Brothers may, from time to time, perform investment banking or other services for, or solicit investment banking or other business from any company mentioned in this document. © 2004 Lehman Brothers. All rights reserved. Additional information is available on request. Please contact a Lehman Brothers' entity in your home jurisdiction.

LEHMAN BROTHERS