citigroup

See the Disclosure Appendix for the Analyst Certification and Other Disclosures.

GLOBAL
STRUCTURED CREDIT
RESEARCH

16 SEPTEMBER 2004

Structured Credit Products

GLOBAL

Trading Credit Tranches

Taking Default Correlation out of the Black Box

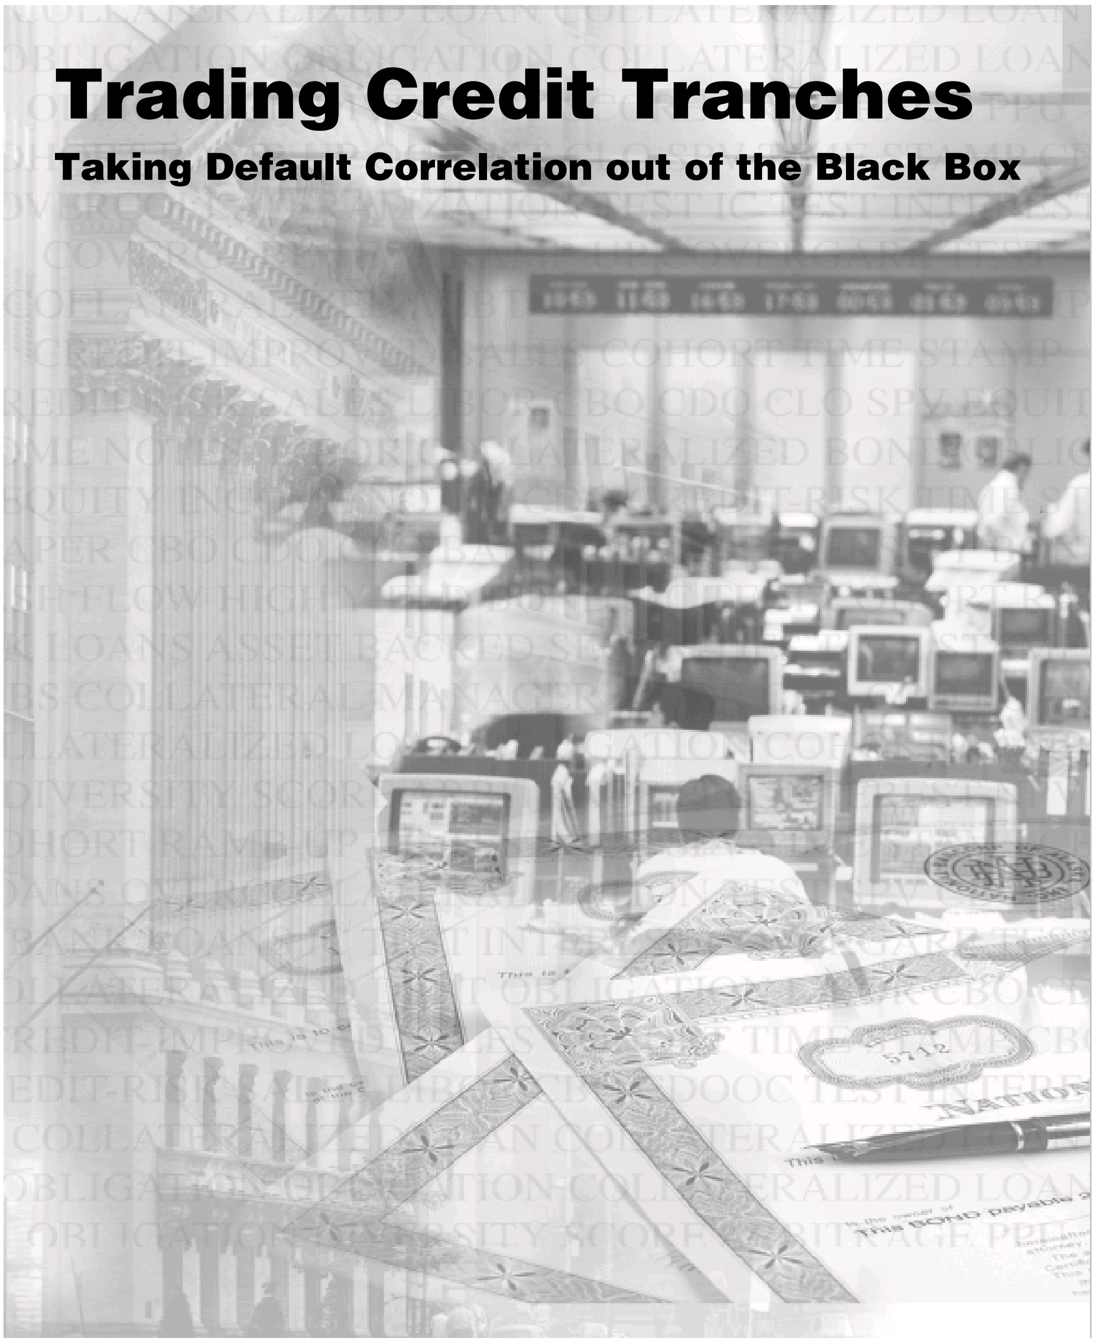

Ratul Roy
(44-20) 7986-9003
ratul.roy@citigroup.com
London

David Shelton
(44-20) 7986 9607
david.shelton@citigroup.com
London

This report can be accessed electronically via:

- FI Direct
- Yield Book
- E-Mail

Please contact your salesperson to receive fixed-income research electronically.

See the Disclosure Appendix for the Analyst Certification and Other Disclosures.

GLOBAL

STRUCTURED CREDIT

STRATEGY

16 SEPTEMBER 2004

Structured Credit Products

GLOBAL

Trading Credit Tranches

Taking Default Correlation out of the Black Box

Ratul Roy
(44-20) 7986-9003
ratul.roy@citigroup.com
London

David Shelton
(44-20) 7986 9607
david.shelton@citigroup.com
London

Contents

Executive Summary

Strong market growth makes understanding correlation critical

Greater market confidence and a search for yield have led to strong growth in the trading and investment of credit tranches. Not just that, we think the degree of leverage now is higher than ever, partly due to more investors getting comfortable with taking equity risk in a low default environment, and partly due to growth in structures leveraging mezzanine tranche risk, e.g. CDO-Squareds. Understanding how tranches behave is therefore even more important.

Traditional tranche-specific correlation has shortfalls

For a tranched investor, default correlation determines what share of the portfolio credit risk stays within a tranche. The street has converged to a common model for portfolio credit risk (Gaussian copula), but no single value within the Gaussian copula framework can explain all index tranche prices. Instead, participants quote tranche-specific correlation. Further, the absence of any pattern among the various implied correlations makes pricing of non-standard tranches difficult. In addition, for many tranches, two correlation values lead to the same price. Many investors are looking for a better model.

A Base Correlation-Skew model breaks portfolio into several equity tranches

A Base Correlation Skew model ("Skew model" in short) breaks a portfolio into a series of equity tranches and treats a mezzanine, e.g. 3-6% tranche, as analogous to spread product, in this case as a spread of two equity (0-3% and 0-6%) tranches. By bootstrapping over losses implied by index equity and mezzanine tranches, we can iteratively find a set of correlations (sometimes expressed as a multiple of a common Base Correlation) that reprices the series of equity tranches. The Skew curve that we obtain expresses the market's risk preference for losses at different attachment points. Supply and demand at various tranches influence the shape of this curve.

The Skew approach has multiple advantages...

We believe there are multiple benefits of the Skew model: the ability to reprice bespoke portfolios and tranches, the ability to capture the market's risk preference, e.g. "fat-tails", and better risk management tools. Price movements in the 3-6% iTraxx tranche have been better predicted through Skew approach, and, non-standard tranches can be correlation-hedged using a combination of liquid iTraxx tranches – which is much more difficult in Tranche Correlation.

...though Tranche Correlation can still occasionally be useful

Tranche Correlation does have some benefits. It is simple. For commoditised index tranches, it provides a quick barometer of periodic price movements. But in our view only the Base Correlation Skew model gives a coherent framework for understanding market movements, and only the Skew model provides a robust method for trading bespoke tranches. Unlike Tranche Correlation, it provides an unambiguous picture of the market's perception of a portfolio's loss distribution and gives investors an opportunity to trade default correlation. In most cases, therefore, we think investors need to start thinking about the Base Correlation Skew approach.

The Credit Tranche Market

Very Strong Growth in Tranche Trading

Investing and trading in credit tranches have seen strong growth catalysed by greater market confidence and a search for yield.

The last few years have witnessed tremendous growth both in credit derivatives, and in tranched credit products referencing pools of corporate credit risk. Usually the reference pools backing these tranched products are pure corporate default risk; more recently, mixed reference pools of corporate and structured finance risk have become popular through synthetic ABS CDO vehicles.

Tight corporate spreads in today's markets have provided a further stimulus to investors' demand for tranched products. Earlier this year, fears were raised that spreads were too tight for investors to take levered corporate risk. If anything, the figure below shows investors did the opposite. Global issuance of synthetic CDOs is running at comparable levels to those of 2003. Investors have been more willing than ever to leverage up, shown by the line in Figure 1 that tracks this measure by taking the ratio of value of tranches distributed to the notional referenced in these structures[1]. This is partly due to more investors getting comfortable with taking equity risk in a low default environment, and partly due to growth in structures leveraging mezzanine tranche risk, e.g. CDO-Squareds.

Figure 1. Global Issuance of Synthetic CDOs and Average Leverage — Issuance in Billions of USD and Leverage in Times

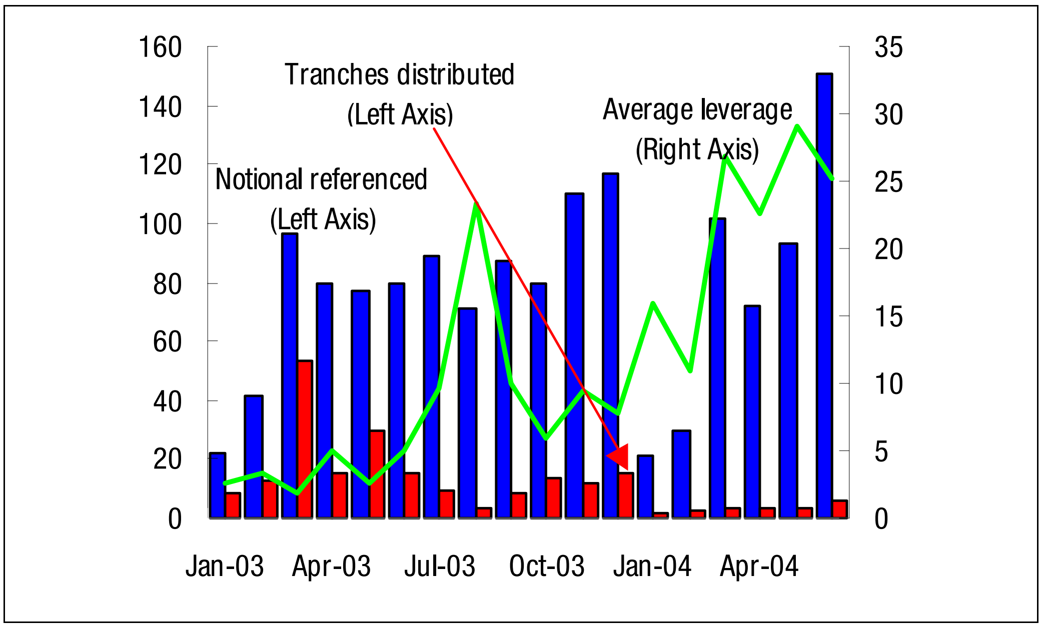

Source: CreditFlux, Citigroup

[1] While this is a simple measure it does have limitations. By only tracking notionals of tranches and pools, a virtually riskless 97-100% tranche would appear to have the same leverage of 33 as a 0%-3% tranche. It would be much better to track the changing market-wide Credit01, the tranche-specific spread-sensitivity, which does differentiate between the 97-100% and 0-3%; however, we do not believe accurate industry-wide data exists.

Index tranche trading is starting to take default correlation out of the black box and bring greater market transparency.

Despite the growth in the market, many investors remain on the sidelines, concerned by the “black-box” nature of the product. In particular, default correlation, which is an important parameter in pricing credit tranches, seems opaque to many investors. Index tranche trading and the prevalence of a common street model for portfolio credit risk (Gaussian copula[2]) has helped to start taking default correlation out of the black box. By quoting prices on standardised tranches of standardised portfolios, participants can trade and hedge default correlation positions separately from their spread and default exposure. Further, as confidence and the number of participants in the product have increased, tranche bid-offer spreads have shrunk.

Importance of Default Correlation in Tranches

Default correlation measures how the average credit risk of a portfolio is distributed across specific tranches.

Why are we so concerned with default correlation? For a tranched investor, default correlation determines what share of the portfolio credit risk stays within a tranche. Using the industry-standard Gaussian copula model, we show below two *excess* loss distributions (i.e. the vertical axis shows probability of loss exceeding values on horizontal axis) for the same average default probability but for two different correlation assumptions (10% and 30%)[3]. The importance of correlation can be illustrated by putting oneself in the shoes of a tranche holder who has 8% subordination (i.e. 8% equity below their investment). The second portfolio, with 30% correlation, would imply a much higher loss for this protection seller, and therefore require a higher premium to be paid in compensation. Higher correlations imply higher losses for senior tranches, and lower losses for equity tranches.

Figure 2. Impact of Correlation on Portfolio Loss Distribution

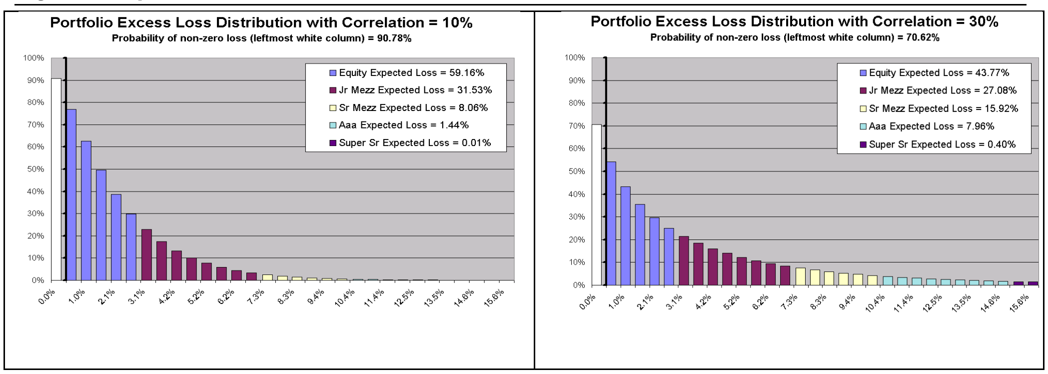

Source: Citigroup

[2] See, for example, “*A Copula Function Approach to Credit Portfolio Modeling*”, David Li, Jerome Connor and Alex Gu, Quantitative Credit Analyst, Citigroup, May 2003.

[3] True default correlation measures the degree to which default of one asset makes the default of another asset more or less likely. These numbers, however, are the asset correlation inputs in a Gaussian copula framework for the construction of a joint distribution of survival times of credits in a portfolio and, although closely related, are not exactly the same. Nonetheless, we will follow the industry standard of calling these parameters default correlation.

Given the importance of this parameter, investors are justifiably concerned about how it should be quantified. Unfortunately, while the index tranche market has brought some welcome transparency, it has also evoked many important questions.

Problems with Traditional Correlation Measure

No single correlation input in the market standard Gaussian copula model can explain all index tranche prices.

It is ironic and concerning that no single default correlation value explains all tranche prices. It is as if different tranche participants at different levels of risk attachment have their own view of the portfolio loss distribution. Figure 3 shows the correlation variable for each tranche that matches the respective tranche premia using a Gaussian copula framework. Senior risk takers, e.g. the 12-22% tranche, are asking for a higher premium, with correlations around 30%, than would be appropriate for more junior tranches. On the 3-6% tranche, meanwhile, correlations are far lower than on the other tranches, at only 4%.

Figure 3. Correlation Factor Fitting Tranche Prices

	5-Year, 20-Sep-09 (42.5bp)		Correlation (%)	
	Bid	Ask	Bid	Ask
0-3%	27*	28.5*	22.4	21.3
3-6%	170	177	4.0	4.1
6-9%	71	75	14.9	15.4
9-12%	42	46	21.9	23.2
12-22%	19.5	22.5	29.7	32.2

Source: Citigroup *Points upfront + 500bp running

We started the discussion with the leveraged position of tranche participants and the impact that default correlation has in the riskiness of any tranche. Yet, we find that the market standard model of calculating the implied Tranche Correlation is raising rather than answering questions. The inconsistencies shown above in correlation levels are to say the least non-intuitive. Supply and demand might cause small differences in traded correlation levels, but there is no reason to think that at different points in the capital structure correlation levels should be more than double the levels for another tranche. We would argue that these inconsistencies are not a problem with the market but rather, evidence of a flaw in the traditional Gaussian copula model.

Skew in Default Correlation

Further Flaws in Tranche Correlation

Tranche Correlation is implied from the specific tranche price but is not unique.

The index tranche market, as we have seen, shows a different implied correlation number for each tranche when one fits the tranche expected loss implied by the traded tranche price as a share of the total overall portfolio loss. We have also seen that there is no simple pattern to this "simple" tranche correlation number. For example, 0-3% is trading at 22%, the next tranche up is at 4% and then correlation rises again to 15% for the 6-9% tranche.

Nor is the inexplicable pattern of correlation the only problem with the traditional Tranche Correlation model. In many cases, this model fails to give a unique correlation value for a given spread level. Two correlation values fit the 3-6% tranche as shown below i.e. it is not clear whether one should use the higher or the lower number. Moreover, as the figure shows, a premium much higher than 300bp for the 3-6% tranche at the portfolio spread of 43bp can never be explained by this approach.

Figure 4. Premium of 3-6% Tranche as Function Of Market-Standard Tranche Correlation for iTraxx=43bp

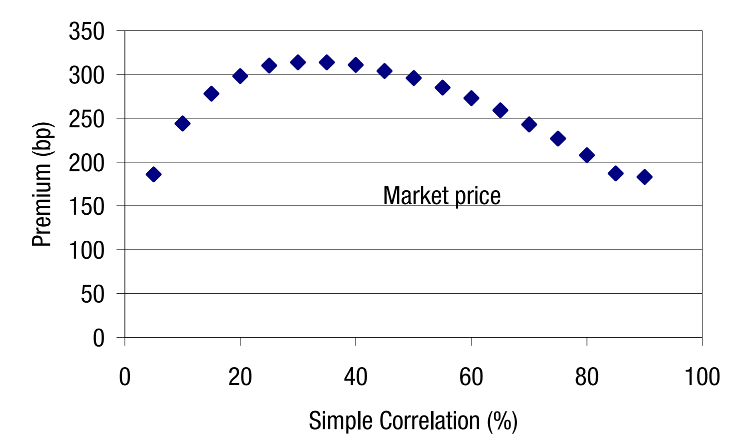

Source: Citigroup

If the Tranche Correlation pattern is really as skewed as Figure 3 shows, it is hard to know what correlation number should be used to price, say, a 4-7% tranche of the index, let alone a tranche from a bespoke portfolio. The pattern in Tranche Correlation does not shed any coherent insight into how investors view risk at different points in the capital structure, and so it is difficult to make comparisons between different portfolios, different maturities, and different risk attachment points within one portfolio. We believe an alternative framework is required.

Correlation Skew is like a Volatility Surface

A Correlation Skew approach looks at mezzanine tranches as spreads between two equity tranches.

Given the problems we have highlighted with Tranche Correlation, we propose a different way of looking at the market's risk appetite at different parts of the loss distribution. We would like to find a model that has an economic explanation and is able to explain the market data, e.g. the relatively high risk premium that we have seen demanded by protection sellers at senior tranches. We take our cue from the equity or FX option markets where two features are widely accepted:

- No single implied volatility in the Black-Scholes framework explains different option prices at different strikes and maturities; however the market accepts the framework as standard and extends it by using a "volatility surface".
- Call/put spread options (options that have a payout between two different strike values) are priced as a difference between call/put options struck at the two strike points. The implied volatilities used at these strike points are given by the "volatility surface" seen in the market.

We therefore propose a 'Base Correlation Skew' model that does just this – it treats a mezzanine, e.g. 3-6%, tranche as analogous to spread product, in this case as a spread of two equity (0-3% and 0-6%) tranches. Equity tranches are analogous to single strike options where the underlying payout to the protection buyer is the portfolio loss.

Quoted index tranche premia provide information on the losses implied by the 0-3% and subsequent mezzanine tranches. As a result, by summing over the losses of mezzanine tranches, we can calculate the loss distribution curve for a series of equity tranches at increasing attachment points (0-3%, 0-6%, 0-9% and so on). We can then iteratively[4] find a correlation that reprices each subsequent equity tranche while satisfying the constraint that the mezzanine tranche is priced and also holding the correlations calculated for each preceding equity tranche fixed. We illustrate the calculation for the 0-6% tranche in Figure 5: as the calculation shows, a correlation value of 27.6% at the 6% attachment point we can reprice the 3-6% tranche. Similarly, we would iteratively calculate the Correlation Skew at the 9% attachment point given the Skew at the 3% and 6% points, and the 6-9% tranche premium.

Rather than looking at the absolute correlation number at each attachment point, investors sometimes look at the correlation as a multiple or fraction of an underlying arbitrarily-chosen 'Base Correlation'. For example, if we chose this number to be 25%, the default correlation at the 3% attachment point would be 79% of this number, and be 110% at the 6% attachment point. Note that the choice of the underlying number would not affect the slope or skew of the default correlation between the 6% and 3% attachment points.

[4] We use a bisection method to perform this calibration. Unlike the multiple solutions seen in Figure 4 for mezzanine tranche correlation, the bootstrapping method produces a unique correlation for each equity tranche attachment point given the monotonic relationship between tranche premium and correlation.

Deriving a Correlation Skew is equivalent to constructing a loss curve for a series of equity tranches.

Figure 5. Correlation Skew and Loss are Inter-linked: Correlation at 6% Attachment Point

Inputs: iTraxx = 43.5bp; 0-3%=28.5points (upfront) plus 500bp running; 3-6%=195bp

Step 1: Calculate 0-3% correlation (equivalent for Skew and Tranche Correlation)

Correlation = 19.7%; LossPV = 13.6mil (1); Annuity01*=101.1mil (2)

Step 2: Calculate 0-6% correlation iteratively

Trial Correlation	Loss (3)	Annuity01 (4)	Premium 3-6%= [(3)- (1)]/[(4)-(2)]
25%	16.6mil	233.9mil	226bp
31.25%	15.6mil	236.2mil	148bp
27.6%	16.2mil	234.8mil	**<u>195bp</u>**

*Annuity01 is the change in present value of the premium leg due to a 1bp change in swap spread. Breakeven premium for zero-upfront tranches is equal to LossPV/Annuity01.

Source: Citigroup

Figure 6 shows an example of the Skew calibration applied to iTraxx and contrasts this with the market-standard Tranche Correlation. Notice most importantly how different the levels are, especially at the 3-6% tranche. We have already highlighted the problem of Tranche Correlation for this risk level. Also note, how for the Tranche Correlation approach, the bid-offer lines cross: this is to do with the different sensitivities of the equity and mezzanine tranches to correlation when measured using this approach. We will revisit this when discussing correlation risk.

Figure 6. Market-Standard Tranche Correlation and Base Correlation Skew Method for iTraxx Tranches

Market-Standard Tranche Correlation	Correlation Skew at various Equity Tranches

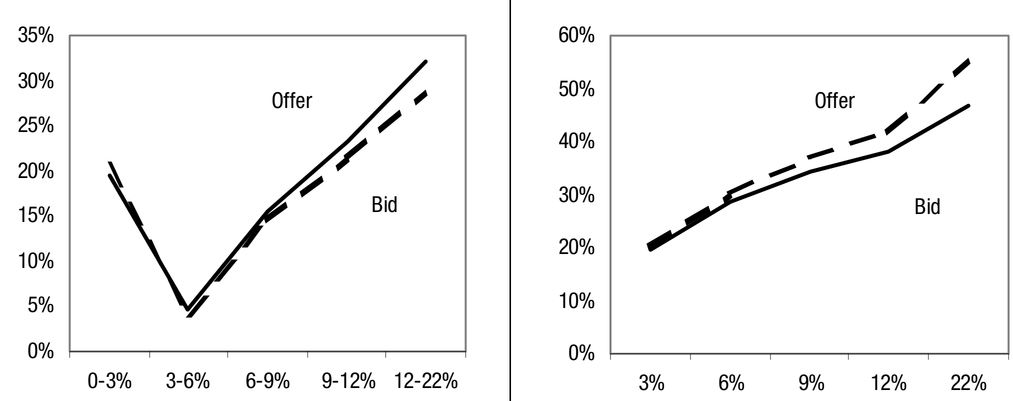

Source: Citigroup, 13 August 2004, iTraxx Europe (mid) =42.5bp

Skew is Market's Risk Preference

The Correlation Skew is an expression of the market's risk aversion at different attachment points of portfolio risk.

The correlation skew, we argue, expresses the market's risk preference for losses at different attachment points. The loss distribution and the correlation curve are intimately linked: our bootstrapping method in Figure 5 explicitly derives the correlation skew curve from the expected loss of the tranche that is implied by its market spread. The method also allows us to build up progressively a loss curve and a correlation skew curve for equity tranches at all index attachment points (3%, 6%, 9%, 12% and 22% for iTraxx). We can then interpolate to price tranches at any attachment point and thickness.

What do changes in the shape of the correlation skew curve mean for participants? We represent this pictorially in Figure 7, which shows the impact of a changing correlation skew curve (right diagram in Figure 6) on the expected loss of each tranche. The circles represent the size of the expected loss of three hypothetical tranches of different seniorities (increasing from left to right), which start off with similar expected losses. The vertical axis shows increasing correlation, and the grey bar shows the correlation in the three tranches. A horizontal bar means that all attachment points have a correlation equal to the base correlation, i.e. there is no skew. We show two types of changes: the top row depicts changes in expected loss for a change in the absolute level of the base correlation across the capital structure while still remaining flat. The bottom row depicts changes in expected loss in situations where the correlation skew curve develops a positive slope for only part or the full height of the capital structure. A lowering of correlation at the junior tranches (either through a parallel shift downwards or a steepening of the grey bar) tends to move out risk into those tranches; in contrast an increase in correlation through a parallel or steepening move increases the expected loss of the senior tranche. When the skew curve is positively sloped throughout, as in the sixth "Smile" scenario, losses are pushed out from the mezzanine into the equity and senior tranches.

If there was no skew, it would be as if there was consensus on the "true" correlation among defaults of individual credits. As we know this is not the case – in particular, today's skew, shown above in the right diagram of Figure 6 resembles the sixth "Smile" scenario, and illustrates the relatively high risk aversion of senior, e.g. 12-22%, tranche holders.

Figure 7. Correlation Skew as Loss Redistribution among Tranches

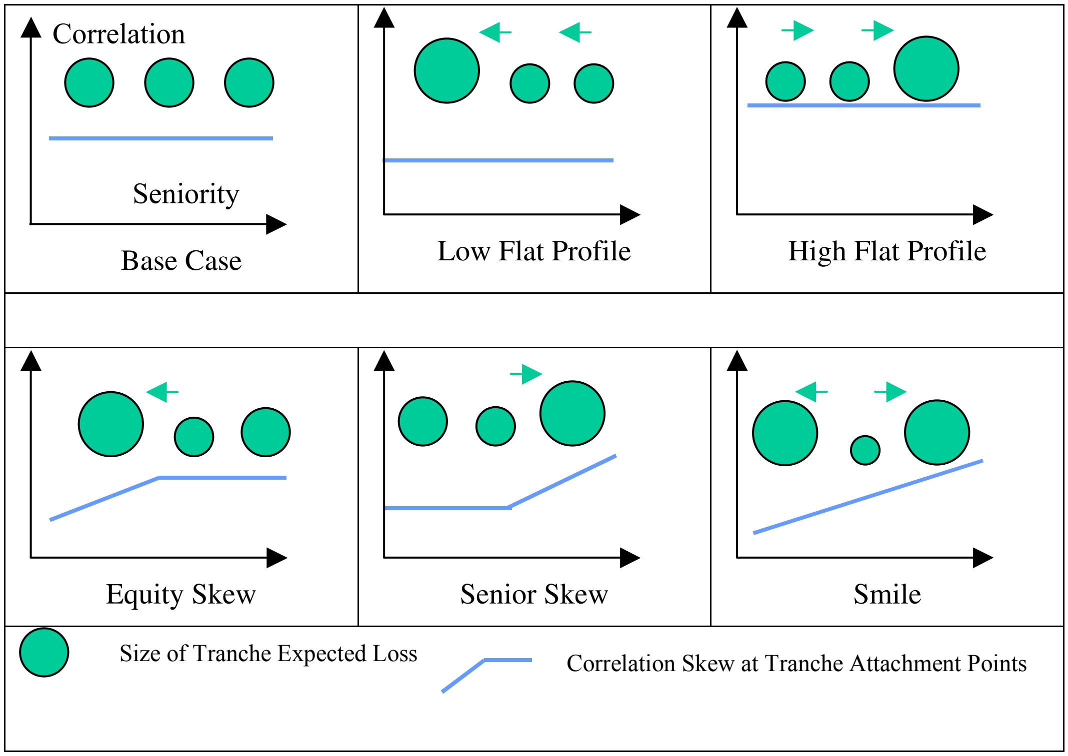

Source: Citigroup

Supply and demand changes will impact the shape of the Correlation Skew.

Supply and demand will change the profile of the skew, and, therefore, the premia for all tranches. Until recently, the market was dominated by substantial demand from ratings-driven institutional investors to sell protection to dealers on tranches typically between 3% and 9% attachment points. These tranches attract investment grade ratings and offer higher spread than similarly rated corporates. This demand for one part of the capital structure puts pressure on tranche premia and creates a "Smile" scenario where mezzanine tranches are priced to lower expected loss. Other factors also influence the skew. Hedged tranche investors wishing to take default risk and positive carry through equity and junior mezzanine tranches but keen to hedge spread movement by buying protection on senior tranches also contribute to the Smile (by bidding up the protection cost). Institutional investors wishing to hedge their cash portfolios, also by buying senior tranche protection, have a similar impact.

Similar to the apparent risk aversion at senior tranches (expressed through higher premium and therefore positively sloped skew curve), investors also seem to demand higher compensation at equity tranches. Investors view portfolios as carrying much more idiosyncratic, and less systemic risk, than senior tranche holders. Broker-dealers too can be part of this group – as a hedge against their long protection position through transacting with investors, many dealers have sometimes been keen to sell protection at these tranches to hedge their correlation risk. In skew terms equity tranches tend to be priced with lower correlation than would be true otherwise: this contributes to the "Smile" effect.

To illustrate the impact of a change in skew on premia, we take two skews – one based on recent iTraxx prices and the other hypothetical, which we call "Thin Tail". We call it so because even though the iTraxx is at the same level (42.5bp in our example) the senior tranches of the Thin Tail pricing are at much lower premia, implying that the probability of high losses is low relative to the risk of high losses implied by recent iTraxx tranche prices. Since the 0-3% in both are at the same level, this means that in the Thin Tail scenario, the expected losses of the senior tranche are now contained in the mezzanine tranche which must therefore demand a higher premium. The pricing and correlation skew of the two scenarios are compared in Figure 8. For completeness, in addition to our Base Correlation Skew levels, we also include the traditional iTraxx Tranche Correlation. Note also that a 364bp premium for the 3-6% tranche is possible in this theoretical scenario, unlike Figure 4.

Figure 8. Premium and Correlation for Equity Tranche at Different Attachment Points for Two Portfolios of Same Spread (42.5bp) But Different Skews: One iTraxx Based, and the Other Hypothetical

Attachment Point	iTraxx Market Skew			Hypothetical Thin Tail Skew	
	Premium	Correlation (Tranche)	Correlation (Skew)	Premium	Correlation (Skew)
3%	28.5*	19.8%	20%	28.5*	20%
6%	177	5.6%	28%	364	13%
9%	75	16.2%	34%	56	13%
12%	46	23.4%	38%	13	13%
22%	22.5	32.1%	46%	2	13%

Source: Citigroup; * Points upfront + 500bp running

We have already shown in Figure 5 that loss and skew are intimately linked (skew is derived from the losses implied by the tranche premia). We should therefore expect that the two different skew curves (Market and Thin Tail) in Figure 8 show different loss distributions. We illustrate two measures of comparison in Figure 9 – one, the cumulative losses at different attachment points (left diagram) and two, the probability of losses exceeding different attachment points (right diagram).

Figure 9. Cumulative Loss and Probability of Excess Loss for Two Different Skew Curves

Cumulative Loss (mil) at Different Attachment Points for iTraxx market based skew, and for hypothetical Thin Tail skew	Probability of Loss Exceeding Attachment Point for iTraxx market based skew, and for hypothetical Thin Tail skew
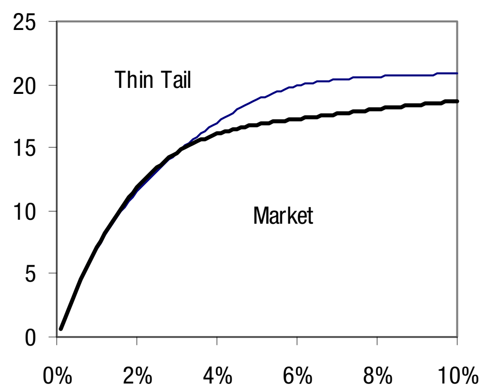	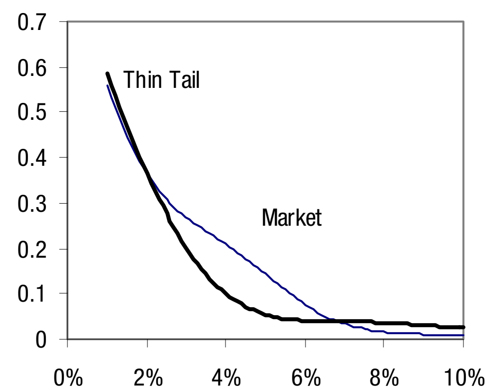

Source: Citigroup

Investor Risk Appetite May Scale Across Markets

The Correlation Skew can be compared to another when one adjusts for the relative riskiness.

If our premise of skew representing risk preference is true, then we should find some relationship among the few index tranche markets in the way investors view risk. Take for example the correlation skew seen in the 5-year and 10-year European iTraxx markets, and in the US CDX markets. At first glance, the skew curves are divergent, as shown in Figure 10 below, which compares the 5-year and 10-year European iTraxx indices, and the 5-year iTraxx and US CDX indices. All three indices show a positively sloped correlation skew curve, much like our 'Smile' scenario, of Figure 7 but there are differences in the levels. For almost all attachment points, correlations on 5-year iTraxx are higher than those on 5-year US CDX and 10-year iTraxx. This seems reasonable, because the same attachment point on 5 years represents a more senior point in terms of expected loss (since expected loss over 5 years is lower than over 10 years). Similarly, expected losses in the US CDX are higher at each attachment point due to the higher spreads, implying more default risk. Hence if risk aversion is greatest at the most senior attachment points (and therefore, willingness to pay for protection for these low-probability default events) then it is fair that a senior tranche holder of the 5-year iTraxx is at a higher correlation position that one in the 10-year iTraxx (because higher correlation means higher premium for the senior tranches).

Figure 10. Correlation of 5Y Eur iTraxx Skew with 5Y US CDX, and with 10Y iTraxx, 8 July 2004

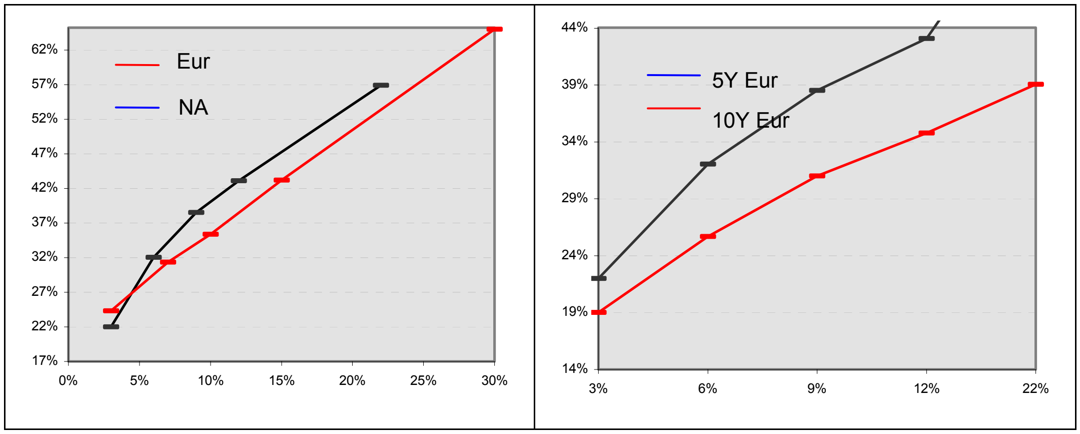

Source: Citigroup

One way to test whether skew corresponds to risk aversion is to see if the differences in the correlation skew curves narrow if one somehow adjusts for the differences in default risk between the portfolios. We do this in two ways – first, by scaling the attachment point by the ratio of spread (as a proxy for default risk), and second, by the ratio of expected loss. The correlation seen at a specific attachment point for a high-risk portfolio must be compared with the correlation seen for a lower attachment point for a lower-risk portfolio. Only by doing this are the two positions placed at the same level of risk aversion along one common portfolio. For example, the 6% attachment point for a 10-year iTraxx should be compared with 4.2% point for 5-year iTraxx if the 10-year spread is 35% more than the 5-year (the common portfolio being the 5-year iTraxx for this comparison). On doing this, we find some convergence, as shown in Figure 11 below.

Figure 11. Correlation of 5Y and 10Y Eur iTraxx and 5Y US CDX, Rescaled to Spread (left) and Loss (right)

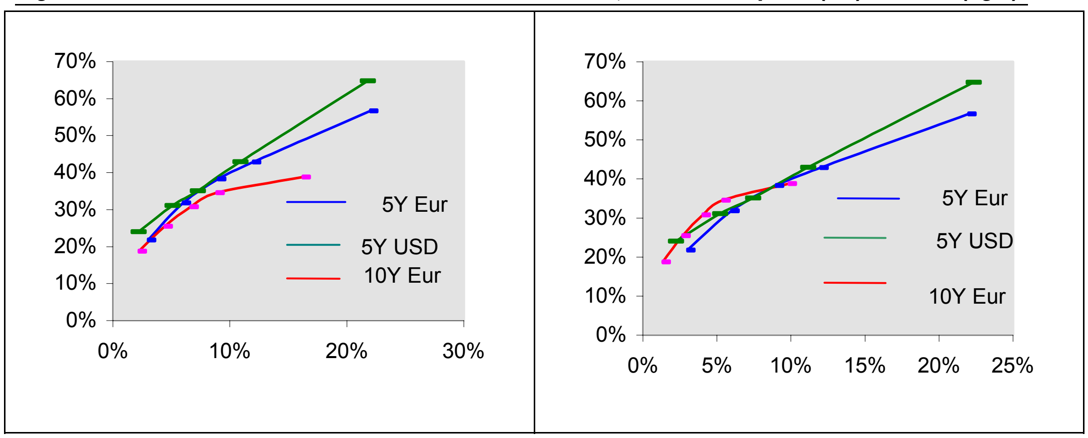

Source: Citigroup, 8 July 2004

The convergence that we have shown above when scaling portfolios based on their risk is certainly very appealing and leads to credibility to our view that Base Correlation Skew represents the market's risk preference at different risk attachment points. Historical data on index tranche trading is relatively short, however, and we will comment on this relationship in future publications.

Greeks: Managing Correlation and Delta Risk

Now that we have established that the Tranche Correlation and Base Correlation Skew methods are different ways of looking at the loss distribution of the mezzanine tranches (and we have argued for the more rational stance of the Skew approach), it will come as no surprise to see that risk measures from the two approaches will show differences.

Skew Model gives more reliable spread sensitivities

The Tranche Correlation and Correlation Skew models give different sensitivity measures for tranches.

The first dramatic result concerns spread sensitivity especially at the junior mezzanine 3-6% tranche. Even though both the Skew and the Tranche Correlation models reprice to the same tranche price, the risk measures can be different. Tranche loss is a function of both portfolio spread and correlation. Moreover, the ratio of the tranche expected loss to that of the portfolio determines the "at-the-money"-ness of the tranche protection. For example, as the probability of portfolio losses recede because of tightening spreads, the mezzanine tranche starts resembling the senior tranche. With increasing spreads, the behaviour is more like equity. The 3-6% tranche sits at the crossroads of two very large jumps in Tranche Correlation (see left diagram of Figure 6). When one uses the Tranche Correlation approach, a change in expected loss of the tranche can shift it into very different correlation territory. In contrast, the Skew approach creates a smooth curve (right diagram of Figure 6) without these jumps. Figure 14 shows that the Tranche Correlation and Skew methods give very different profiles of change in present value of the 3-6% tranche with respect changes in portfolio spread. Since Credit01 is the slope of the curve in Figure 12, one can see that the two approaches present different numbers. As a consequence a risk taker who wants to position for convexity by doing a long tranche-short single name Credit01 neutral trade is presented with two different hedge ratios.

Figure 12. 3-6% Tranche Present Value with Respect to Spread Using Base Correlation Skew and Tranche Correlation Methods

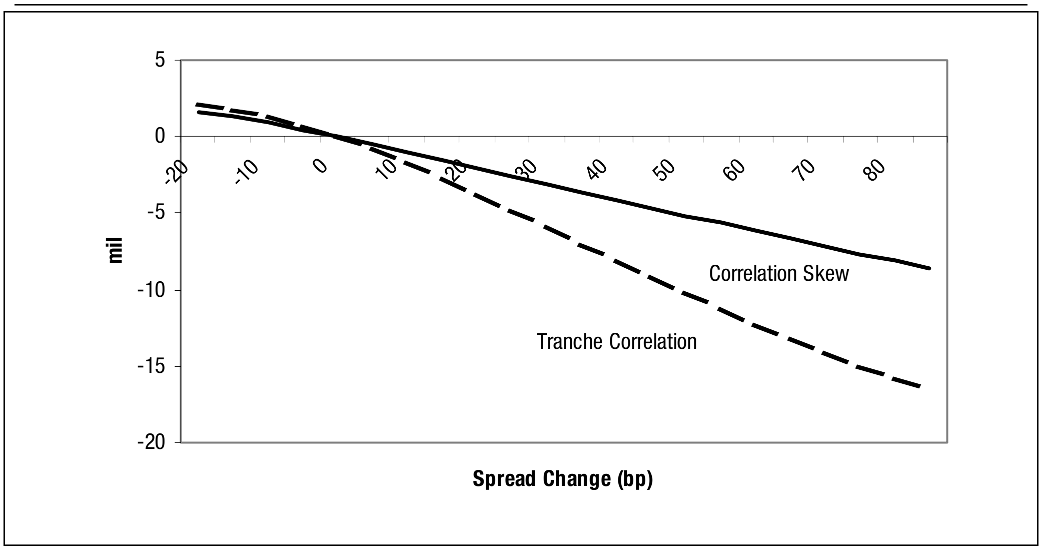

Source: Citigroup

Price movements in the 3-6% iTraxx tranche have been better predicted through the Correlation Skew approach.

We believe the sensitivity predicted by the Correlation Skew is more correct. The market, in fact, quotes Credit01s in relation to index tranche trading, which are closer to the theoretical values calculated by the Skew framework. In Figure 13 below we compare the spread-sensitivity (Credit01) of the two methods by predicting the change in the daily 3-6% tranche premia from the change in the iTraxx premia, knowing the tranche duration[5]. A hedged investor would have performed better in general with the Skew model. The very recent divergence has, however, shown that model improvements can be made in our understanding of portfolio credit risk.

Figure 13. Actual versus Predicted 3-6% Tranche Premium Change with Tranche Correlation and Skew

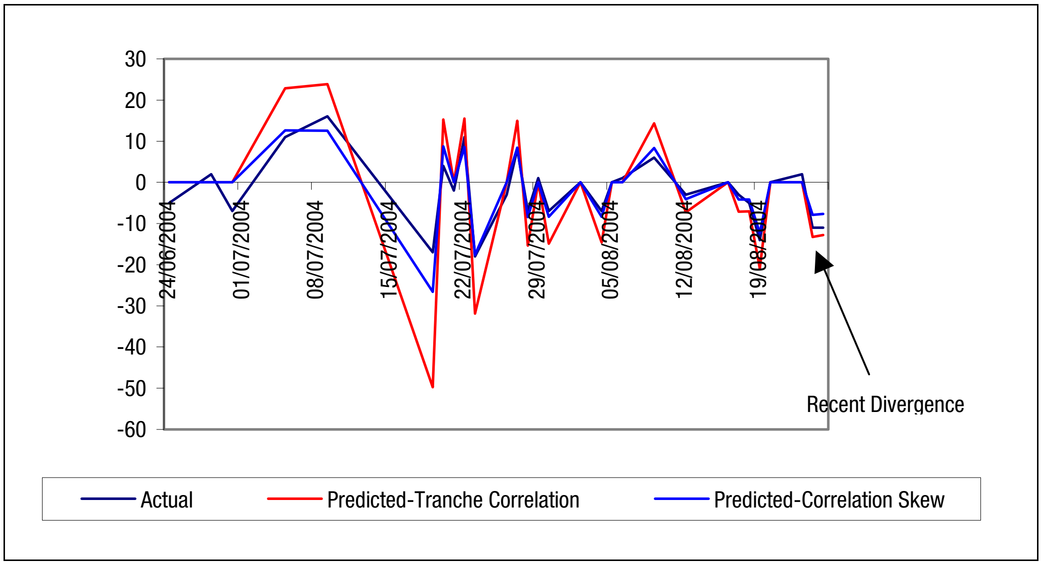

Source: Citigroup

Skew Model gives unique Loss Exposure

The Tranche Correlation model gives two solutions for the probability of loss exceeding any attachment point.

As a tranche risk participant, one is anxious to know about the probability of loss exceeding various attachment points – in particular, the lower attachment point of the tranche. Here too, the two approaches provide different results. We note that the implied probability may be obtained by differentiating with respect to either the high or the low attachment point of each mezzanine tranche. However we have two possible observations for each attachment point – the tranche below and the one above the point. For example, if we consider the 3% attachment point, we can obtain the probability of losses exceeding this level by either differentiating the 0-3% tranche or the 3-7% tranche. Since in the Tranche Correlation approach the two tranches are priced using very different correlations (see Figure 3 for example), we obtain 2 sets of different results for this and higher attachment points (see Figure 14 below). We also see quite different values to those we get based on using a Base Correlation Skew approach of single correlation per equity tranche. Notice how the Skew approach produces the widest probability of extreme losses by having the fattest tail for losses exceeding 30%.

---

[5] Predicted change in 3-6% Tranche = (Change in iTraxx Spread) Multiplied by (Credit01) Divided by (Tranche Duration)

Figure 14. Implied Probability of Losses Exceeding Various Attachment Points

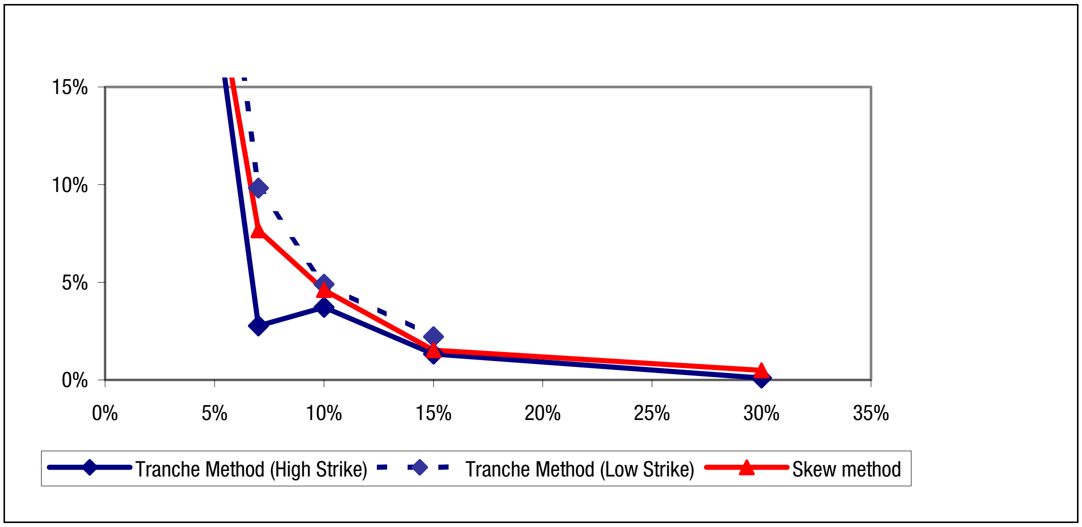

Source: Citigroup

Skew Model allows accurate Correlation Sensitivity hedging

Non-standard tranches can be priced using the Correlation Skew curve and correlation-hedged using a combination of liquid iTraxx tranches.

The street is used to thinking of correlation sensitivity as change of tranche value by changing the simple Tranche Correlation by 1%. This is a relatively substantial change for the 3-6% tranche (which trades at a low Tranche Correlation) and much less for the others. In the Base Correlation Skew framework, the correlation sensitivity of each mezzanine needs to be considered in terms of two or more buckets of equity tranches whose attachment points correspond to the upper and lower attachment points of the mezzanine. A seller of 6-9% mezzanine tranche protection is exposed to the short the 0-6% and long 0-9% part of the correlation skew curve. If the skew curve moves in parallel then the net sensitivity is just the sum of the two numbers.

Correlation bucketing is a powerful way of looking at correlation risk. It aggregates a portfolio of tranched products – index and bespoke portfolios – into various parts of the correlation skew curve and suggests suitable correlation hedges. Tranche Correlation is not as robust because there is no underlying relationship to describe the correlation behaviour of different tranches. For example, it would be unclear as to what hedges would be appropriate for a non-index tranche, e.g. 5-9%. In contrast, as shown in Figure 15, which illustrates the sensitivities to the different buckets, this tranche has sensitivities to the 0-3%, 3-6% and 6-9% tranche correlations.

Figure 15. Sensitivities to Correlation Using Two Approaches for 4 iTraxx Tranches (each EUR 30 mil)

	Tranche Correlation	Skew (Parallel Move)	Correlation Skew Bucketed Risk			
			0-3%	0-6%	0-9%	0-12%
0-3%	214,723	219,467	219,467			
3-6%	(131,000*)	(27,087)	(204,301)	177,214		
6-9%	(73,952)	(4,137)		(172,678)	168,541	
9-12%	(41,606)	(48,652)			(167,465)	118,813
5-9%	N/a	(17,014)	(63,025)	(125,957)	172,171	

* Average of 188,000 and (450,000)

Source: Citigroup

Likewise, one can derive a term structure of correlation. Since indices are increasingly quoted for several maturities, it is possible (and consistent) to imply a separate correlation skew for different maturities, thus ensuring that the pricing of a seven-year trade (for example) is consistent with both the skew at 5-years and the skew at 10-years. One can think of this as a line somewhere between the two curves in the right diagram of Figure 10.

Finally, remember that a Skew approach is merely a framework for looking at mezzanine tranches as a payoff between two default strikes. It makes no assumptions about the Copula model used (Gaussian or otherwise). We can also take a general correlation matrix, which may include different sectoral and subsectoral correlations and apply the skew to the entire matrix: the most simple choice is to rescale all correlations by the same skew factor. In this way we can still calibrate the market skew whilst capturing more of the name-specific detail.

In Summary: Why Skew is A Better Model

There are multiple benefits of the Skew model.

We think the Base Correlation Skew Model is a more robust way of looking at the well-established investor risk profile that the Tranche Correlation model shows. We have four main reasons. The first advantage is purely practical, and relates to the pricing of non-standard tranche attachments. Given the jumps in Tranche Correlation, we have no insight into the value to be used for a tranche that spans, for example, part of two index attachment points. As we have illustrated, by being able to relate skews across a range of portfolios through their risk characteristics and maturities, one can price and hedge customised tranches of bespoke portfolios.

The second advantage is that the Skew approach captures the market's risk preference – an example being the risk aversion at senior tranches manifested by relatively high premium for low risk probability. These risk preferences, commonly termed "Fat-Tail" or "Smile", can be represented, however, by other analytical approaches, e.g. the Marshall-Olkin copula, which will exhibit different loss distributions from the Gaussian copula.

The third advantage is the uniqueness and range of correlation values particularly for the 3-6% tranche. We have seen that at today's iTraxx index and tranche levels, the 3-6% tranche has two solutions for Tranche Correlation. Further, the maximum allowable premium for this tranche given today's index spread is limited irrespective of the Tranche Correlation used (see for example Figure 4 where the maximum allowed premium for the 3-6% tranche was a little over 300bp). In contrast, the Skew approach only has one solution and as we have shown in Figure 8 can have solutions higher than this limit.

The last and important benefit relates to risk measures. It has long been known that the spread sensitivity of the 3-6% tranche as predicted by the Tranche Correlation approach over-predicted risk. The market also quotes the spread sensitivity for this and other tranches in addition to premium – it does so by quoting the Delta of the tranches as a multiple of the underlying iTraxx index. The quoted Delta for the 3-6% tranche (currently 6.7) is close to what is predicted by the Skew approach. Likewise, the correlation sensitivity for the tranche is better expressed by the Skew approach for the reasons we have described above.

That is not to say that the Skew model is the final word on the subject. Part of the success, or otherwise, of the model is the level of detail in the assumptions that are used. Common instances where greater detail may be useful are the use of individual spreads for all credits instead of an average spread, and characterisation of the default correlation between credits as due to several, and not just one, i.e. systemic, parameter. As presented in this report, we have not found it necessary to implement a multi-parameter model for default correlation between credits, but this is one of several adjustments that can be accommodated within the Skew framework. In many instances, e.g. less diversified portfolios, we would recommend such additions. We have also just shown instances, e.g. in spread sensitivity of the 3-6% iTraxx, where the Skew model had good, but not perfect, predictive power. And, finally, as in any statistical model of portfolio loss, complete reliance on credit spreads and market-implied default correlation is not the best strategy if other information on individual credits, e.g. bottom-up credit analysis, is available.

Trading Opportunities for Investors

Tranche Correlation can still provide Insight

For standard index tranches, Tranche Correlation can still be useful in providing valuable trading insight.

Having spent most of the last section singing the praises of the Correlation Skew approach, it may come as some surprise to hear that we think Tranche Correlation can be useful in identifying relative value opportunities. Where there is little change in portfolio credit quality, e.g. composition and spread, Tranche Correlation provides an attractive shorthand for changes in risk appetite across different attachment points. Thus, for example, as shown in the left chart in Figure 16, Tranche Correlation moved out significantly on the 0-3% tranche from its low levels in late 2003 as greater participation from the hedge fund community was able to satisfy the natural axe of broker-dealers to buy protection at this tranche. Conversely, as market participants have started to view the senior (9-12% and 12-22%) tranches as an efficient and levered way to go short[6], the Tranche Correlation for these tranches have also moved out. Deriving the same intuition from the correlation skew is less convenient – as the schematic representation Figure 7 shows, increasing risk aversion for the senior attachment points (i.e. higher premia) is a function of both the level and steepness of the curve. The right chart in Figure 16 shows that this is the case by looking at three dates. Two features can be observed – first, the absolute correlation level at low attachment points was indeed the lowest late last year, and second, the curve is indeed at its steepest now, i.e. risk has moved out currently from the mezzanine tranches to the equity and senior tranches (i.e. the sixth scenario in Figure 7). From a Tranche Correlation perspective, this would mean a lowering in implied correlation of the 3-6 tranche, which is what we see in the left diagram as well.

Figure 16. Evolution of Tranche Correlation and Base Correlation Skew

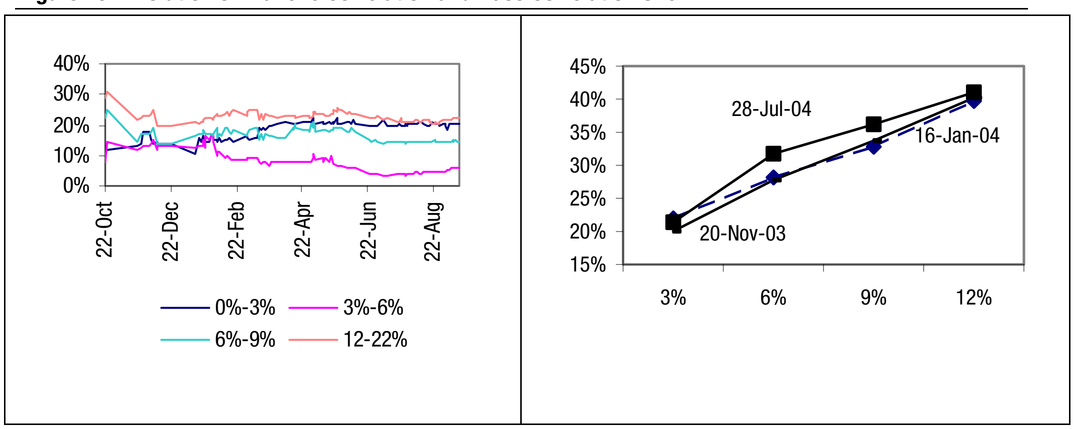

Source: Citigroup

Often, though, the relationship between Skew and Tranche Correlation is not that obvious. As a quick illustration, we show in Figure 17 how a change in tranche premia by the model Credit01 (for a 1bp move in iTraxx, Scenario B versus A) has left the Skew curve unchanged, but altered each of the Tranche Correlations.

[6] See "Bull and Bear in a Boxx", Arvind Rajan et al., Citigroup, February 2004.

Likewise, a change in the premia of a specific tranche, e.g. the 3-6% without any change in iTraxx (comparing Scenario C with B) will change only the Tranche Correlation of this specific tranche, but leave the others unchanged. In contrast, since the Skew curve at each attachment point is bootstrapped from all the junior tranches below this point, such a change will affect the Skew curve across all attachment points.

Figure 17. Tranche and Skew Correlation

	Actual (A) (iTraxx=36.5bp)			Full Delta Move (B) (iTraxx=37.5bp)			3-6% offer down (C) (iTraxx=37.5bp)		
	Prem (bp)	Tranche (%)	Skew (%)	Prem (bp)	Tranche (%)	Skew (%)	Prem (bp)	Tranche (%)	Skew (%)
0-3%	24*	18	18	25*	17.9	17.9	25*	17.9	17.9
3-6%	115	5.4	28.3	122	5.2	28.3	117	4.8	28.9
6-9%	52	16.4	34.5	56	16.5	34.5	56	16.5	35.1
9-12%	33	24.1	38.7	34	23.8	38.7	34	23.8	39.4

Source: Citigroup *Points upfront + 500bp running

Convergence in Correlation Skew in Markets

Following differences in Correlation Skew between markets can lead to profitable convergence trading opportunities.

If investors share our view that the skew curve represents relative risk preference then participants can put on trades that will profit from greater convergence between two markets. We made this case in an earlier publication[7] suggesting that investors position for convergence in correlation skew on the 10-year Eur iTraxx and US CDX converge. This convergence had already happened in the 5-year Eur iTraxx and US CDX, but had yet to be seen in the 10-year maturity. The recommended trade was to position for convergence in the most liquid tranches – in particular to buy protection on the 6-9% iTraxx versus selling protection on 7-10% CDX (the same idea could be expressed through other tranches e.g. 3-6% tranches of US CDX and Eur iTraxx). Figure 18 shows one skew convergence scenario and the P&L impact on the trade. Clearly the most obvious downside for the trade would be if the European tranche became even more expensive relative to the US tranche, causing the spread difference to widen. Since then, however, the skew did converge resulting in a profit.

We think Base Correlation Skew framework will continue to be useful in identifying further such opportunities as investors now have a common metric to compare various tranches and portfolios.

[7] See "*US Europe 10 Year Correlation Trade*", Matt King and Antoine Pain, Citigroup, 29 July 2004.

Figure 18. Scenario for Skew Convergence to US-Europe Average, and P&L for Trade, 28 July 2004

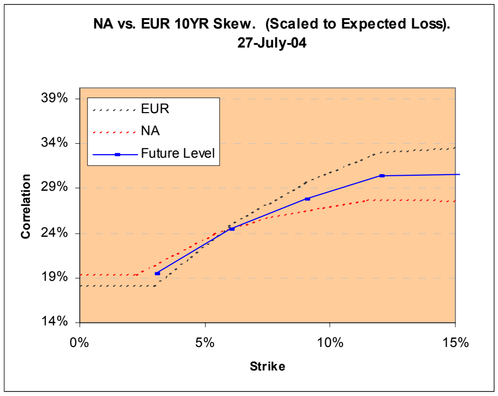

	Current	Scenario	P&L
Eur			
3-6	460	476	16.1
6-9	216	240	**23.4**
NA			
3-7	662	649	13.8
7-10	327	297	**29.5**

Source: Citigroup

Pricing Off-Market Tranches

Investors and the street may view the value of bespoke portfolio tranches differently, and execute trades accordingly.

The way dealers price tranches of bespoke portfolios necessarily reflects observables in the tranched market, i.e. 5-year and 10-year Europe iTraxx and US CDX markets. This represents only two portfolios, five tranche attachment points, and two maturities (a total of twenty combinations). Broker-dealers have to use this data set to price tranches of a wide range of subordinations and thicknesses belonging to bespoke portfolios of all hues. Two common assumptions that are made are that default correlation between individual credits is uniform and independent of the specific sector, and that correlation skew between portfolios scale to the expected loss (as we have illustrated with some success in Figure 11). A third assumption is also sometimes made, often as part of the so-called Homogeneous Large Pool Model, which treats each credit at the average portfolio spread and ignores any barbelling in spreads. The set of steps that dealers often follow is

1. Calculate portfolio expected loss using individual spreads, using the bid or offer side of the market depending on whether protection is bought or sold.

2. Decide on a correlation skew proxy, using either the US or European markets based on the regional portfolio composition.

3. For trade maturities that do not correspond to the index, interpolate a skew based on the 5-year and 10-year points.

4. Scale up or down the index skew based on expected loss differences between the index and the transaction portfolios following the argument in Figure 11.

5. Calculate the tranche premia and risk measures using a portfolio loss model.

Of course, some participants may continue to price using the simple Tranche Correlation model. Equally some participants may choose to use average instead of individual spreads. What this means for the risk taker is that there may be some relative value arguments either for or against certain bespoke transactions. Take, as an illustration in Figure 19, differences in risk measures that emerge if one uses the individual credit spreads for a portfolio (the curved lines) versus modelling each credit at the same average spread (the horizontal lines). Depending on model choice, someone executing the trade on a delta (credit spread) neutral basis will place somewhat different hedges.

Figure 19. Credit01 for Tranches Using Actual 225-Name Portfolio Spreads versus Using Average Spread

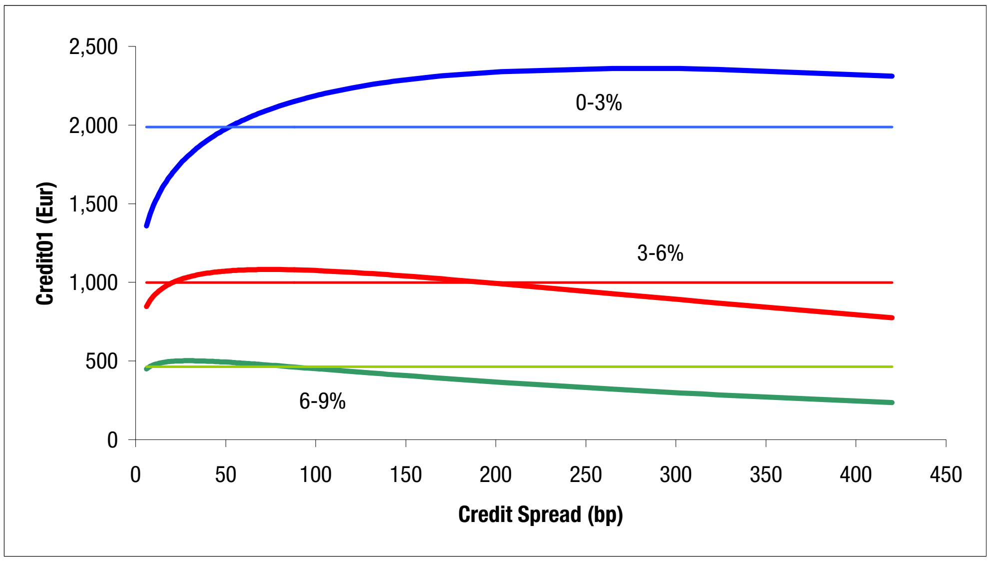

Source: Citigroup.

Looking at the correlation skew in bespoke portfolios approach can lead to other trading opportunities, including exploiting any cheapness in default correlation. For example, if a participant is able to buy protection on a lumpy, low diversity portfolio at levels that are cheap because the quoted spreads (and implied correlation) are more in line with the implied correlation observed for tranches of the higher-diversity CDX and iTraxx indices, then there are relative value arguments to do the trade. The participant can hold the position outright, or hedge the position against spread movements by selling protection on other tranches or single name credit swaps. By taking the latter route, the position would then be primarily on default correlation; any mark-to-market gain in future can be monetised by unwinding the trade.

Conclusion and Future Agenda

The growing liquidity and emergence of indices in the credit tranche market has opened up numerous trading opportunities for participants. Key among the new developments is a better understanding of one of the important factors driving tranche price and risk – default correlation. We believe the approach that we describe in this report – a Base Correlation Skew approach similar to the pricing of currency and equity spread options – is better able to explain observable tranche prices and risk compared to the previous simple Tranche Correlation approach. The Skew approach also has advantages in being able to make comparisons across markets.

We urge investors, however, to use both approaches as they seek to exploit trading opportunities in the market. Tranche Correlation does have some benefits. It is simple. For commoditised index tranches, it provides a quick barometer of periodic price movements. But in our view only the Skew model gives a coherent framework for understanding market movements, and only the Skew model provides robust method for trading bespoke tranches. Unlike Tranche Correlation, it provides an unambiguous picture of the market's perception of a portfolio's loss distribution and gives investors an opportunity to trade default correlation.

While the approach discussed here, a combination of Gaussian copula with correlation skew, has enabled a better understanding of market's risk preference at various tranche attachment points, this is not the final word. We have shown instances, e.g. in spread sensitivity of various tranches, where neither model has perfect predictive power. Alternative copula expressions for joint distributions of credit default, as well as introduction of additional risk parameters, e.g. global catastrophic shock, are analytical variations that we are currently exploring. We will come back to you with these in future publications.

Disclosure Appendix

ANALYST CERTIFICATION

We, Ratul Roy, hereby certify that all of the views expressed in this research report accurately reflect my personal views about any and all of the subject issuer(s) or securities. I also certify that no part of my compensation was, is, or will be directly or indirectly related to the specific recommendation(s) or views in this report.

Other Disclosures

ADDITIONAL INFORMATION AVAILABLE UPON REQUEST

Citibank, N.A., London Branch and Citigroup Global Markets Inc, including its parent, subsidiaries, divisions and/or affiliates ("the Firm"), may make a market in the securities discussed in this report and may sell to or buy from customers, as principal, securities recommended in this report. The Firm may have a position in securities or options of any issuer recommended in this report. The Firm may be regular issuers of, and trade in (including position taking), financial instruments linked to securities, which may have been reported on in this research report. The Firm may perform or solicit investment banking or other services from any issuer recommended in this report. An employee of the Firm may be a director of an issuer recommended in this report. Within the past three years, the Firm may have acted as manager or co-manager of a public offering of the securities of any issuer recommended in this report. Securities recommended, offered, or sold by the Firm : (i) are not insured by the Federal Deposit Insurance Corporation; (ii) are not deposits or other obligations of any insured depository institution (including Citibank); and (iii) are subject to investment risks, including the possible loss of the principal amount invested. Past performance is not a guarantee of future results. This report does not take into account the investment objectives, financial situation or particular needs of any particular person. Investors should obtain advice based on their own individual circumstances before making an investment decision. Investing in non-U.S. securities, including ADR's entails certain risks. The securities of non-U.S. issuers may not be registered with, nor be subject to the reporting requirements of, the U.S. Securities and Exchange Commission. There may be limited information available on foreign securities. Foreign companies are generally not subject to uniform audit and reporting standards, practices and requirements comparable to those in the U.S. Securities of some foreign companies may be less liquid and their prices more volatile than securities of comparable U.S. companies. In addition, exchange rate movements may have an adverse effect on the value of an investment in a foreign securities and its corresponding dividend payment for U.S. investors. Net dividends to ADR investors are estimated, using withholding tax rates conventions, deemed accurate, but investors are urged to consult their tax advisor for exact dividend computations. Although information has been obtained from and is based upon sources the Firm believes to be reliable, we do not guarantee its accuracy and it may be incomplete or condensed. All opinions and estimates constitute the Firm 's judgement as of the date of the report and are subject to change without notice. This report is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security. This research report does not constitute an offer of securities. Any decision to purchase securities mentioned in this research must take into account existing public information on such security or any registered prospectus. Investing in non-US securities by US persons may entail certain risks. Investors who have received this report from the Firm may be prohibited in certain US States from purchasing securities mentioned in this report from the Firm; please ask your Financial Consultant for additional details. This report is distributed in the United Kingdom by Citibank, N.A. London Branch or Citigroup Global Markets Limited, Citigroup Centre, Canada Square, Canary Wharf, London E14 5LB, UK. This material is directed exclusively at market professional and institutional investor customers in the United Kingdom and is not for distribution to private customers in the United Kingdom, as defined by the rules of the Financial Services Authority, who should not rely on this material. Moreover, any investment or service to which the material may relate will not be made available to such private customers. This material may relate to investments or services of a person outside of the United Kingdom or to other matters which are not regulated by the Financial Services Authority and further details as to where this may be the case are available upon request in respect of this material. If this publication is being made available in certain provinces of Canada by Citigroup Global Markets (Canada) Inc. ("The Firm Canada"), the Firm Canada has approved this publication. If this report was prepared by the Firm (excluding Nikko Citigroup Limited) and distributed in Japan by Nikko Citigroup Limited, it is being so distributed under license. This report is made available in Australia, to non-retail clients through Citigroup Global Markets Australia Pty Limited (ABN 64 003 114 832 and AFSL No. 240992), a participant of the ASX Group and a participant of the Sydney Futures Exchange Limited and to retail clients through Smith Barney Citigroup Australia Pty Ltd (ABN 10 009 145 555 and AFSL No. 240813), a participant of the ASX Group. In New Zealand it is made available through Citigroup Global Markets New Zealand Limited, a member firm of the New Zealand Stock Exchange. Citigroup Global Markets (Pty) Limited is incorporated in the Republic of South Africa (company registration number 2000/025866/07) and its registered office is at Citibank Plaza, 145 West Street, Sandown, Sandton, 2196, Republic of South Africa. The investments and services contained herein are not available to private customers in South Africa. This publication is made available in Singapore through Citigroup Global Markets Singapore Pte Ltd, a Capital Markets Services license holder. This report is being distributed in Hong Kong by or on behalf of, and is attributable to Citigroup Global Markets Asia Limited, 20th Floor, Three Exchange Square, Hong Kong. Citigroup Global Markets Inc. is a member of the Securities Investor Protection Corporation (SIPC).