The Implications of Implied Correlation

Roy Mashal
+1 212-526-7931
rmashal@lehman.com

Marco Naldi
+1 212-526-1728
mnaldi@lehman.com

Gaurav Tejwani
+1 212 526 4484
gtejwani@lehman.com

Implied correlation is increasingly used for relative value considerations when comparing alternative investments in synthetic CDO tranches. Here we show that, by neglecting the heterogeneity of the underlying portfolio, implied correlation may lead to misleading relative value assessments. We argue that a modified implied correlation measure, which we call the "implied correlation bump", may be more appropriate for the relative value analysis of alternative tranched investments.

INTRODUCTION

Synthetic collateralised debt obligations (CDOs) are instruments whose payouts are linked to the performance of a portfolio of synthetic credit exposures. This market has experienced continuous innovation over the past few years. While in the early days synthetic CDOs were mainly used by banks for capital relief, most of the issuance is now generated by the dealer community in the form of one-off, bespoke tranches referencing investment grade credit default swaps. These structures allow for a high degree of customisation, and can therefore be used to tailor specific exposures to the risk preferences of a variety of different credit investors. This customisation has spurred significant growth in the credit derivatives business, as reported last year in a *Risk* survey (February 2003, page 20).

In 2003, default swaps on the TRAC-X and CDX portfolios were introduced, and tranches linked to these reference sets also started to be actively quoted. This portfolio standardisation has allowed for the creation of a more liquid and transparent market for tranched risk, and the recent merger of the two portfolio products is likely to drive this process even further. The new availability of relatively liquid market levels has led to the quotation of tranche prices in terms of "implied correlation" – a practice that is clearly reminiscent of the use of implied volatility in the options markets. Here, we explain what implied correlation is and why this measure alone may be insufficient for comparing two alternative investments.

Synthetic CDO market participants often face investment decisions such as the following: "I can sell protection on one of two equity tranches – which one offers better value?"

Investment A
- *0%-3% (first loss) tranche.*
- 100 name reference pool with avg spread of 60bp and avg historical correlation of 25%
- Investor can sell protection at a spread of 17.4% corresponding to **an implied correlation of 21.5%.**

Investment B
- *0%-3% (first loss) tranche.*
- 115 name reference pool with avg spread of 85bp and avg historical correlation of 26.5%
- Investor can sell protection at a spread of 20.6% corresponding **to an implied correlation of 28.9%.**

Investment A has a lower implied correlation than investment B. Also, investment A is being quoted at an implied correlation that is lower than the average historical correlation on the reference pool, while investment B has an implied correlation higher than the average on its portfolio. What conclusions, if any, can I draw from these observations? Is it fair to say that investment A is cheap to historical correlation while investment B is rich?

WHAT IS IMPLIED CORRELATION?

Increasingly, market participants are quoting the implied correlation rather than the spread or the price of a CDO tranche. The implied correlation of a tranche is the uniform asset correlation[1] number that makes the fair or theoretical value of a tranche equal to its market quote. Currently, the most common models used to price synthetic CDOs are variants of the one-factor Gaussian copula model. In this model, the correlation of default times is determined by the correlation of asset returns, so tranche values are directly related, though in a complex way, to the assumed asset correlations.[2] In the above example (investment A), plugging a uniform correlation of 21.5% into our pricing model (thus ignoring the actual correlation structure) would produce a fair spread on the equity tranche equal to the market quote of 17.4%.

This is somewhat analogous to the equity derivatives market, where quoting implied volatility is equivalent to stating the price, since all other variables are known. Given the model and all the assumptions that go into it, quoting the spread on the tranche is equivalent to stating its implied correlation. Although it is an elegant way of representing a tranche price, there are two key differences in the analogy drawn with the equity market, and investors should be aware of them:

- Unlike the equity options market, where the Black-Scholes model has gained universal acceptance, the models utilised in the CDO market vary across market players and keep evolving over time. Consequently, the implied parameters, which are model-dependent, will be different as well.
- While volatility is a single parameter in equity derivatives models, a typical reference portfolio of a synthetic CDO tranche generally has thousands of pairwise correlation parameters – for example, a 100-name portfolio involves 4,950 parameters in its correlation matrix. Fitting a flat correlation structure is appealing because of its intuitive simplicity, but expressing a complex relationship in one number can often be inaccurate as it does not reflect the heterogeneity of a portfolio. We illustrate this point below with an example and show why investors should be cautious in interpreting the results.

We next use a set of hypothetical investments to demonstrate how tranches on portfolios with identical average characteristics may have significantly different fair prices and, therefore, significantly different fair implied correlations.

EXAMPLE 1: PRICING A FIRST LOSS (0-3%) SYNTHETIC CDO TRANCHE

In this example, we construct four hypothetical reference portfolios, each with the same average characteristics but different relationships among the reference assets within the portfolio. Each portfolio consists of 100 names with an average spread of 60 basis points and an average observed asset correlation of 25% with other assets.[3] The key characteristics of the portfolios are the following:

- Homogenous or base case – flat (constant) spreads (60bp) and uniform correlation (25%).
- Portfolio 1 – flat (constant) spreads, but variable asset correlations.

---

[1] *The cashflows to a tranche are determined by the default realisation in the reference pool. Most models utilise asset correlation (or its proxy - equity correlation) as a means of generating correlated defaults.*

[2] *See D. O'Kane, M. Naldi, S. Ganapati, A.M. Berd, C. Pedersen, L. Schloegl and R. Mashal (2003),* The Lehman Brothers Guide to Exotic Credit Derivatives*, Risk, October.*

[3] *As shown in Figure 1, the average spread of the investment-grade CDX portfolio on 22 April 2004 was 59bp, while the average historical correlation was close to 25%. The hypothetical portfolios we choose here are therefore similar to the CDX portfolio in terms of their average characteristics.*

- Portfolio 2 – spreads positively related to asset correlations.
- Portfolio 3 – spreads negatively related to asset correlations.

Next, we compare the prices of the four equity (0-3%) tranches[4] using a typical one-factor Gaussian copula model. The results of our analysis are shown in Figure 1.

Figure 1. Comparing the prices of four equity (0-3%) tranches

	Portfolio Description				Equity Tranche	
	Spreads (bp)	Correlation	Avg. Spread (bp)	Avg. observed Correlation	Fair Spread	Implied Correlation
Base Case	Flat 60	Flat 25%	60	25%	16.11%	25.00%
Portfolio 1	Flat 60	Variable[5]	60	25%	16.09%	25.03%
Portfolio 2	30 to 90	Higher correlation amongst higher spread names	60	25%	15.02%	28.88%
Portfolio 3	30 to 90	Higher correlation amongst lower spread names	60	25%	17.43%	21.49%
CDX.IG.NA 4/22/04	Varies		59	25.1%	16.77%	24.28%

Figure 1 shows the fair spread of the tranche using the observed pairwise correlations of the reference pool, as well as the implied correlation that can be backed out from the fair spread. For example, an investor who sells protection on the first loss piece of portfolio 2 should receive a spread of 15.02% as a fair compensation for his risk exposure. Although the average observed correlation on the underlying portfolio is 25%, the implied correlation is 28.88%; that is, we can substitute the actual correlation structure with a flat number (28.88%) and achieve the same price (fair value spread) on the equity tranche.

The example clearly shows how fair implied correlations of the tranches referencing the four hypothetical portfolios are different, even though the average characteristics of the reference pools are the same. For example, investors considering the first loss piece from portfolio 3 should receive a fair implied correlation of 21.49% compared with 28.88% for portfolio 2 – a difference of 7.39%! This demonstrates our first premise – that a portfolio with *uniform* pairwise correlations of X (25% in this case) does not necessarily have the same risk profile as a portfolio with an *average* pairwise correlation of X (25%). If the above hypothetical tranches traded at their fair values, investors should be indifferent in their choice. However, an investor choosing on the basis of a comparison between implied correlation and average historical correlations would be mistakenly attracted to portfolio 3.

A careful look at the characteristics of the reference pools and their heterogeneity helps explain the previous results. Because of the positive relation between spreads and correlation of the underlying names, the loss distribution of portfolio 2 is more volatile than that of the other three portfolios. This means portfolio 2 has a greater probability of realising a very low number of defaults as well as a higher likelihood of realising a very large number of defaults, with average default realisations being less likely. Equity investors benefit from the extra

[4] A similar analysis can be performed using mezzanine or senior tranches, but the relatively high sensitivity of the equity tranche makes it a better candidate for illustrative purposes

[5] In the context of a one-factor model, correlations are generated by the product of the sensitivities to the common market factor, generally called the "betas". More formally, we produce a heterogeneous correlation matrix by taking the cross product of a vector of betas with its transpose, and varying the elements of this vector from 25% to 75%. The average correlation is thus 50% x 50% = 25%

probability of very benign default scenarios, while they are not particularly hurt by the extra probability of a large number of defaults. Consequently, the fair spread to be paid to the equity investor on portfolio 2 is lower than the fair spread to be paid to the equity holders of the other three portfolios. A symmetric argument explains why the fair spread to be paid to the equity investor on portfolio 3 is higher than the other three fair spreads.

The above examples have shown how changing the heterogeneity of a portfolio, even while maintaining a similar average profile, changes the theoretical value of a tranche. Therefore, in the real world, comparing the implied correlations of two tranches referencing different portfolios carries even lesser meaning. Let us now revisit the investment decision we posed at the beginning of this note, and compare investments A and B – two equity tranches referencing different portfolios.

EXAMPLE 2: COMPARING INVESTMENTS A AND B

Investment A has been used as portfolio 3 in the earlier example, and given our assumption that it is being shown at its fair value spread, the present value or mark-to-market on the tranche is equal to zero on day one. Investment B, on the other hand, references a higher spread portfolio. As shown in Figure 2, at a market quote of 20.6% (which corresponds to an implied correlation of 28.9%), and using historical pairwise correlations to price the tranche, the investment is worth, on day one, $94,613 for a $10 million notional. Clearly, investment B would be a better choice here. Yet, a naive comparison between historical and implied correlations would suggest that investment A is a better value.

Figure 2. Comparing two alternative investments

	Investment A (0-3% Tranche) $10m Notional	Investment B (0-3% Tranche) $10m Notional
Reference Pool		
Number of Names	100	115
Average Spread	60bp	85bp
Average Correlation	25%	27%
Market Quote		
Spread	17.4%	20.6%
Implied Correlation	21.4%	28.9%
Present Value	$0	$94,613

AN ALTERNATIVE: THE “IMPLIED CORRELATION BUMP”

One of the main reasons for which the notion of implied correlation has gained so much ground in the market is the simple fact that it is just one number, making it easy to quote. But, as we have just seen, implied correlation is a poor measure for relative-value analysis because it neglects the heterogeneity of a portfolio. As an alternative, we suggest using an “implied correlation bump”, ie, a single number that multiplies all historical pairwise correlations in order to re-price a tranche.

Revisiting the comparison between investment A and investment B, there is evidently no need to bump the correlations for investment A since we assumed the equity piece trades at fair value. The implied correlation bump for investment B, on the other hand, is 92.5%, that is, all the elements in the historical pairwise correlation matrix would have to be scaled down by 7.5% to price the tranche at its current market (spread) quote. We can therefore say that the tranche is "cheap to historical correlation" – an observation that cannot be made by looking at implied correlation alone. This bump in correlation is worth $94,613 in terms of present value, as discussed earlier.

The implied correlation bump has the advantage of respecting the specific diversification of the portfolio, while retaining the convenient feature of fitting just one number. The main drawback, however, is that contrary to implied correlation, it is based on a historical estimate of the correlation matrix. Consequently, the implied correlation bump is likely to be more useful for relative-value analysis than for quoting the price of a tranche.

SUMMARY AND CONCLUSION

The increased liquidity of standardised bespoke tranches has brought more transparency to the market, as well as the ability to better calibrate proprietary models. One of the recent trends in the synthetic market is the quotation of tranche prices by means of their implied correlations. However, implied correlations are increasingly being used for other purposes such as relative-value analysis. In this article, we have highlighted some of the potential drawbacks of using implied correlations for evaluating the relative attractiveness of alternative tranched investments. In particular, we have shown that using a flat implied correlation number does not properly account for the cross-sectional variability of pairwise correlations between individual credits, and this can lead to misleading investment choices. Consequently, we do not recommend drawing conclusions about the relative attractiveness of two tranches based on the comparison of their implied correlations, especially if the two tranches refer to different portfolios. As an alternative, we have suggested looking at the "implied correlation bump", a measure that in our opinion can be more meaningfully employed to detect relative value across tranches.

It must be pointed out that our previous examples were all based on the popular one-factor Gaussian framework. Using this model, the correlations implied by the observable prices of junior and senior tranches are generally higher than those needed to match the prices of mezzanine slices – a phenomenon known as the "correlation smile". As recently observed by Duffie (2004)[6], the inability of the model to explain observable market prices across the capital structure using the same correlation matrix raises doubts about the appropriateness of the underlying distributional assumptions.

Looking further into the future, we believe that identifying models that are able to fit observable prices is more than an academically interesting exercise: it responds to practitioners' and regulators' growing demand to somehow "anchor" the valuation of a large notional amount of illiquid, customised exposures to the aggregate opinion of the market-place, thereby increasing transparency and promoting the growth of these products even further.

[6] *See D Duffie (2004), Time to adapt copula methods for modelling credit risk correlation,* Risk *April 2004, page 77.*