Issues in the pricing of synthetic CDOs

Christopher C. Finger
RiskMetrics Group
chris.finger@riskmetrics.com

In this article, we discuss the standard pricing model framework for synthetic CDOs. Though the standard framework is by now well accepted, how the model is implemented precisely and, importantly, how the model is applied, vary across the marketplace. We discuss some of the outstanding implementation and application issues, and propose a number of questions that further research on the model should seek to address.

1 Introduction

Derivatives markets are typically called mature when they achieve an adequate level of liquidity, and, as either a cause or an effect of that liquidity, standardization. Standardization refers to a number of issues associated with the derivative market, including forms of contracts, documentation, benchmarks and pricing models. Certainly, proprietary pricing models persist in even the most mature markets, but a standard pricing model will always take hold, both to facilitate price communication on the most liquid products and to allow valuation of less liquid products consistently.

While not truly mature, the market for synthetic CDOs and basket credit derivatives has grown to the point that it may at least be considered adolescent. In particular, the standardization of the credit derivatives index contracts and tranches on these indices, coupled with large activity in single-tranche CDOs, has created an accepted standard definition of what a synthetic CDO is. At the same time, the market has coalesced around a specific pricing model, the normal copula model. However, though the majority of market participants rely on this model as a basic framework, there are many points where specific implementations may differ. Furthermore, applications and interpretations of the model differ as well, and a number of debates have arisen as to the best way to quote information from the model.

In this article, we summarize briefly the standard normal copula framework. We then survey the differences that exist in specific implementations and applications, and examine some of these differences through an example. We conclude with some thoughts on what questions we might ask to evaluate the implications of the differences that exist.

2 The standard model

Throughout this article, we will refer to the single-factor normal copula model as the standard pricing model. Li (1999) and Li (2000) are typically cited as the first application of the model to pricing CDOs.

To be clear, the basic structure we are pricing is a simple static unfunded synthetic CDO. The structure references a portfolio of obligors, typically with equal weightings. As defaults occur, losses are assigned to the *protection sellers* sequentially. That is, losses are assigned to the first loss protection seller until they reach a certain threshold, subsequent losses are assigned to the next protection seller until they reach a second threshold, and so on. The portfolio loss level at which a specific tranche begins to experience losses is referred to as the tranche's attachment point, and the level at which the tranche no longer experiences losses as the detachment point. In return, the protection sellers are paid a running premium on the current amount they are protecting. For example, suppose the first loss seller receives 500bp annually for protecting the first three percent of losses on a $10M pool. If no losses have occurred, then the investor is protecting the full $300,000 and is paid 500bp on this amount; once $100,000 of losses occur, the premium is paid on the remaining $200,000 that the investor is protecting.[1] The protection sellers, particularly the first-loss investors, may be paid an upfront premium as well. The standard tranches on the two most common indices (North America and Europe Investment Grade) are listed in Table 1.

Mina and Stern (2003) describe the structure in more detail and point out that to price the structure, it is sufficient to describe the distribution of cumulative portfolio losses at each payment time. This means that it is not necessary to describe the full dynamics of the loss process. If we were to consider structures with path dependence (for example, a structure wherein the second loss position knocked out if the first loss position were exhausted too quickly), then we would need to fully describe the portfolio loss dynamics, but for our purposes, the sequence of loss distributions at each payment time suffices.

[1] A premium that is paid in this fashion is often referred to as a *risky premium* since the investor is not certain about how much premium he will receive. Other conventions, in which the premium is paid on a fixed amount throughout the life of the structure, also exist, but are somewhat less common.

Table 1
Standard index tranches and pricing. May 17, 2004

North America			Europe		
Tranche	Upfront (%)	Running (bp)	Tranche	Upfront (%)	Running (bp)
0-3%	42	500	0-3%	31	500
3-7%	-	331	3-6%	-	255
7-10%	-	126	6-9%	-	105
10-15%	-	54	9-12%	-	58
15-30%	-	16	12-22%	-	22

Source: St. Pierre et al (2004).

An important aspect of copula models in general is that they decouple the specification of the individual stand-alone distributions (in our case, the default probabilities of the obligors in the CDO at a specific time) from the specification of the correlation structure. Thus, we assume that the individual obligor default probabilities are known, most likely derived from the single-name credit default swap market. The correlation structure is then what we specify with our choice of copula function.

Under the normal copula, the joint probability of two events (in our case, two defaults) is the same as the joint probability of two events with the same individual probabilities that come from normal random variables. Specifically, if two obligors have default probabilities p_1 and p_2, then we represent the default events by the equally likely events $\{Z_1 < \alpha_1\}$ and $\{Z_2 < \alpha_2\}$, where Z_1 and Z_2 are standard normal random variables and α_1 and α_2 are chosen appropriately. We then impose a correlation between Z_1 and Z_2, and calculate the joint probability of the two events.[2] Thus, with the normal copula as our model and the default probabilities derived from the single-name market, we have only to specify the correlations to complete our description of the portfolio loss distribution. Applying the same normal copula and correlations at each payment date gives us the needed sequence of loss distributions.[3]

[2] In the context of structural models of default, the Z variates here can be identified with the obligors' firm asset values; see Li (1999) for a detailed explanation. That the normal copula can be reconciled with structural models gives some "physical" support for the framework, and keeps it from being simply a mathematical device to produce correlations.

[3] It would certainly be possible to consider a term structure of correlations, though this has received little attention in practice.

While the normal copula gives us a framework for correlations, it remains to specify the particular correlation values. While any valid correlation matrix could be used, it has become common practice to restrict the correlation structure to that described by a single common factor. Specifically, the restriction is that the normal variate associated with each obligor can be represented by

$$Z_i = w_i Z + \sqrt{1 - w_i^2}\varepsilon_i,$$

where Z and the ε_i are independent standard normal random variables. Under this framework, the correlation between obligors i and j is $w_i w_j$, so the range of possible correlation matrices is quite limited. In fact, the most common application is to set the weight parameters w_i equal for all obligors and consider just one pairwise correlation value for the entire portfolio.

The benefits of the single-factor restriction are first that it reduces the number of parameters required by the model, but most importantly that it gives rise to effective techniques to calculate pricing in closed-form. Most of these techniques rely on the observation that once the value of the common factor is fixed, then the individual obligor defaults are conditionally independent.[4] Thus, conditionally, the portfolio loss is the sum of independent obligor losses, and the conditional portfolio loss distribution is the convolution of the individual obligor loss distributions. There are numerous standard convolution techniques (Fourier transforms, Laplace transforms, moment generating functions). The pricing calculation then reduces to a numerical integral over the single factor, with each step in the integration involving a convolution to derive the conditional portfolio distribution. There have been numerous articles recently detailing similar pricing frameworks, among which are Mina and Stern (2003), Martin, Thompson, and Browne (2001), Merino and Nyfeler (2002), and Gregory and Laurent (2003). Additionally, most active credit derivatives dealers have provided notes on the model; among these are Ahluwalia, McGinty, and Beinstein (2004), Watkinson and Roosevelt (2004), St. Pierre et al (2004), and O'Kane et al (2003).

A lingering issue is whether the single-factor structure overall will be rich enough to price and hedge all synthetic CDOs in practice. However, to move away from the single-factor framework will require meeting two significant challenges. The first is to develop numerical methods that enable the same level of precision and speed in the calculation of prices and sensitivities that the marketplace has become accustomed to. The second is to educate the marketplace as consumers of richer correlation information. Significant effort has already gone into establishing the single

[4] See Finger (1999) for examples.

correlation parameter as a pricing parameter for synthetic CDOs, and a sense of benchmark values for this correlation now exists. Re-educating the marketplace to a new, more complex correlation scheme would be a significant hurdle.

3 Variations on the standard model

Though the single-factor normal copula has emerged as the standard framework for valuing synthetic CDOs, there are still a number of issues which can produce discrepancies between seemingly similar model implementations. We discuss four of these in this section.

3.1 Correlation structure

As mentioned above, the single-factor correlation structure enables closed-form solutions for CDO prices. While the more restrictive version (that of a single pairwise correlation) is most common for wide communication of pricing, the less restrictive version (where the weights w_i can vary by obligor) is also commonly used. Fortunately, this is a relatively simple issue to communicate, and rarely causes confusion.

3.2 Individual spreads or portfolio average

An exact model of a synthetic CDO accounts for the spreads of each underlying name, and differentiates, for example, between a portfolio with fifty names at 50bp and fifty names at 450bp and a portfolio with one hundred names at 250bp. This exact treatment is certainly possible under the standard model, as long as spreads on all the underlying names are available.

Some model implementations utilize only the portfolio average spread and assume that all underlying names are characterized by this spread. Clearly, this approximation results in sacrifices to pricing accuracy. Additionally, this approximation prevents us from being able to distinguish between the different underlying names for hedging purposes; the simplified model can only produce a "portfolio hedge" consisting of equal weights on all the underlying names. On the other hand, the average spread approximation is quite attractive in that it significantly reduces the amount of information necessary to describe the structure. In practice, the average spread approximation appears adequate, at least for investment grade CDOs, where the range of underlying spreads is relatively small. However, there is little information available to quantify the impact of the approximation, and it remains an open topic how this approximation fares as spreads widen and disperse. We take up this issue further in the example in Section 5.

3.3 Fully granular model or large pool approximation

In a full model implementation, we distinguish between portfolios of twenty, fifty or one hundred names, and account for the fact that portfolio losses can only occur in discrete increments. We will refer to this as the fully granular model. A second typical model simplification is to approximate the portfolio by a limiting case with infinitely many small positions. To characterize the limit, start with the original portfolio. Then create a second portfolio with twice as many obligors, the same overall characteristics,[5] but position sizes half as large as in the original. Repeating this process brings us to a limiting portfolio with the same characteristics as the original, but whose distribution is continuous. This limiting distribution is actually quite simple to express, and is now used commonly to describe credit portfolios. Vasicek (1997) is the first description of the limiting distribution.[6] Finger (1999) applies the model as a shortcut to simulating large credit portfolios. Most recently, the new Basel capital guidelines[7] implement a version of the distribution in their Internal Ratings-Based Approach to minimum regulatory capital.

The difference between the fully granular model and the large pool approximation is that in the full model, there is still some idiosyncratic risk: once we condition on the single factor, there is still some variability in the portfolio, and this variability depends, among other things, on the number of obligors in the portfolio. In contrast, in the large pool limit, the realization of the single factor uniquely determines the realization of portfolio loss. As with the average spread approximation, the large pool approximation involves a sacrifice in modeling accuracy. Here, though, the benefit is yet another gain in model simplicity: the resulting implementation involves only a numerical integral and no need for a convolution step.

JP Morgan has been the most vocal dealer promoting the large pool model, and supplies a spreadsheet implementation to its clients.[8] It is important to note that the JP Morgan implementation utilizes both the large pool and the average spread approximations. There is nothing inconsistent with applying both, but this has resulted in some confusion in the marketplace, that the

[5] For instance, if the 10% of the obligors in the original portfolio have spreads over 100bp, then 10% of the obligors in the new portfolio do as well.
[6] Vasicek (1997) modestly calls the distribution the normal inverse, though many practitioners today refer to it as the Vasicek distribution.
[7] Basel Committee on Banking Supervision (2004).
[8] McGinty et al (2004).

two approximations are inextricably linked. In fact, it is possible to run the fully granular model using the average spread approximation, as it is to run the large pool model using the individual obligor spreads. Thus, any evaluation of the JP Morgan implementation should seek to isolate the effects of the two approximations, and not treat them as one choice.

As with the average spread approximation, the adequacy of the large pool approximation relative to its benefits is still an open question. The JP Morgan model has performed well for the investment grade index tranches (one hundred obligor portfolios), while no practitioner would ever apply the large pool approximation to five- or ten-name first-to-default baskets. In a different context, there has been some research[9] on the viability of the large pool model for producing economic capital figures, but this has treated much larger and more diverse portfolios than those that underlie a typical CDO. There is a need for more research on this topic, applied particularly to the CDO pricing case and isolating the large pool from the average spread approximation. We provide further discussion in Section 5.

3.4 Default timing

A final implementation question is how we account for the timing of defaults. As discussed previously, the standard model requires the computation of the cumulative loss distribution at each payment date of the CDO. From this point, however, it remains to specify when, during a payment period, defaults occur. This impacts pricing in two ways. First, if losses are paid at the moment a default occurs (as opposed to being paid at the end of the period), then the timing of the default within the period influences the time period over which the loss is discounted. Earlier defaults then result in losses with a larger discounted value. Second, the timing of a default can influence the premium paid for the period, since the default (if it impacts the tranche in question) reduces the size of the tranche. An earlier default thus implies that the size of the tranche is reduced earlier, meaning that the premium paid for the period is lower. Thus, other things being equal, earlier defaults during each payment period increase the present value of the loss payments and reduce the present value of the premium payments, resulting in a reduced valuation for the seller of protection.

Conceptually, it is possible to have the default timing be purely model-driven, either by implementing a Monte Carlo scheme or by discretizing the closed form solution at a daily

[9] Gordy (2000).

frequency. In practice, however, it is common to make an assumption. Perhaps the most neutral assumption is to assume that all defaults happen midway through the payment period. For valuing the protection seller's position, a conservative assumption is to assume defaults occur at the beginning of the period, while for the protection buyer, the conservative approach is to assume defaults occur at the end of the period. In any case, the effect is subtle, particularly for CDOs with quarterly frequency, but this can explain pricing discrepancies when two model implementations appear to be identical. We touch on this issue in the example.

4 Application issues

More active than the implementation questions recently have been discussions on how the model is applied, particularly regarding the correlation parameter. The discussion of appropriate data sources for credit correlation is actually fairly mature, in that this has been an active topic among users of economic capital models since 1997. The CDO market brings a new element to this discussion, however, since the market provides a mechanism to infer correlation from actual pricing levels. The most active debate in this arena centers on what method to use for this inference. Finally, the choice of correlation data influences how the model is used for calculating hedge ratios and other sensitivities. We treat each of these topics in more detail in this section.

4.1 Correlation data sources

Absent a liquid market from which to draw information, earlier research centered on choosing a data source and estimation technique that could provide reasonable correlation parameters for pricing. Correlations of illiquid credits are also critical to economic capital models, and so investigations into this area date to the release of a number of public versions of these models in 1997.[10]

Under the structural model,[11] defaults are driven by the evolution of a firm's assets. Consequently, dependencies between defaults are driven by dependencies between firm assets. Asset values being unobservable directly, there are two approaches to estimating their correlation: one is to derive time series of asset values from the equity market and compute correlations from the derived series[12];

[10] The year 1997 saw the release of three well publicized credit portfolio models: CreditMetrics, CreditRisk+, and Credit Portfolio View. See respectively Gupton, Finger, and Bhatia (1997), CSFP (1997), and Wilson (1997). While implementations of all three persist in economic capital applications, the structure of CreditMetrics is most similar to the standard pricing models for CDOs.

[11] Merton (1974).

[12] Kealhofer (2001).

the other is to use the equity correlations directly. The advantage of either approach is that it allows for a very granular assessment of correlation; we may use equity information on precisely the names in the portfolio we are pricing. The disadvantage is that we must rely on information from outside the credit markets, and a model (albeit one with considerable academic and empirical support) to make the connection. Mashal and Naldi (2002) offer a recent examination of both asset and equity information as drivers of credit correlations.

A less granular, but more direct, approach is to infer correlations from historical default experience. The standard pricing model, given default probabilities and correlations, forecasts a distribution of possible default experiences. With appropriate assumptions, we can consider the yearly outcomes in a default database as realizations of such a distribution. If we assume knowledge of the default probabilities for the names in the database, it is then possible to infer the correlation level that best describes the realized outcomes. De Servigny and Renault (2003) take this approach to provide correlation estimates which they compare to equity market information. Demey et al (2004) provide a more rigorous treatment in a maximum likelihood context, and investigate the bias and small sample properties of their estimator. While promising, particularly for economic capital problems, small sample issues restrict the correlation estimates to a sector level, and thus, the application of this method to CDOs may be limited to providing benchmark correlation levels.

Another approach using credit data is to simply correlate moves in the spreads for individual names. There has been less historical work in this area, since it is not applicable in most economic capital applications. One drawback of this approach is that the standard CDO pricing model does not describe movements of spreads, but rather only default events.[13] Thus, the correlation in the model is not a spread correlation, but rather is better interpreted, as discussed previously, as an asset correlation. Structural models, which link asset moves and spread moves, provide a mechanism to convert spread correlations to asset correlations, which in turn can be used in the standard model. The link is complex, however, and to first order, it is reasonable, as many practitioners do, to ignore this somewhat academic argument and simply use spread correlations in the model.

[13] Duffie (2004) notes that this fact also keeps us from using the same model framework to price derivatives whose payoffs are spread dependent, such as options on credit indices.

4.2 Correlation as a market factor

While investigations into empirical sources of correlation data continue, the discussion has shifted in the last year to the use of the synthetic CDO market itself as a source for correlation information. Many dealers now provide regular pricing on standard tranches of the larger credit indices. This pricing is the most direct indicator of the correlation of the names in the collateral pool. It is desirable, then, to leverage this information to value non-standard tranches on the index portfolios, as well as synthetic CDOs backed by other portfolios. Of course, there is a basis risk involved with applying correlations for the index portfolio to bespoke CDOs, but this risk appears preferable to any of the more indirect methods of establishing correlations.

Arguably the most active discussion at this point centers on the details of how to represent correlations on the standard tranches and to apply these correlations to other CDOs. If we have pricing information for a tranche and its underlying names, we can solve for the single correlation that under the standard model reproduces the tranche pricing. This notion of correlation is typically referred to as the tranche's *implied correlation* or *compound correlation*. Typically, the implied correlations vary across the different tranches on a single portfolio, and indeed a profile where implied correlation falls between the first and second loss tranches and then rises for subsequent tranches is common. While quite intuitive, the implied correlation approach has a number of drawbacks, chief among which is that it frequently fails when applied to the second loss tranche of the major indices.[14] The failure is either that there are two correlation levels that reproduce the tranche pricing, or more seriously, that no correlation level reproduces the tranche pricing.[15] This failure can be an indication that the constant correlation, or single factor framework, is not adequate to describe all price possibilities, but can also be an indication that we should utilize the constant correlation framework differently.

To address this and other shortcomings of the implied correlation framework, a number of dealers[16] now utilize *base correlations*, sometimes referred to as *detachment correlations*. Though specific implementations vary slightly, the idea behind base correlations is to infer correlation levels that price a succession of first loss tranches, rather than the specific standard tranches. The key observation in the approach is that any arbitrary tranche can be represented as the difference between two first loss tranches. For instance, selling protection on the 3-7% tranche of the North

[14] The standard second loss tranches are 3-7% on the Dow Jones CDX North America Investment Grade, and 3-6% on the iBoxx European Investment Grade.
[15] McGinty et al (2004) and St. Pierre et al (2004).
[16] McGinty et al (2004) and St. Pierre et al (2004).

America index is equivalent to buying protection on the 0-3% tranche and selling protection on a hypothetical 0-7% tranche. Under the base correlation approach, we fix a single correlation to price the 0-3% tranche, then look for a second correlation to price the 0-7% tranche such that the price difference is consistent with the observed 3-7% tranche price. By construction, this approach resolves the problem of multiple correlation solutions, since first loss tranche prices are always monotonic with correlation. In practice, the base correlation approach is successful for pricing the problematic second loss tranches. It should be noted, however, that calibration problems persist with this approach, as it is possible that no set of base correlations can be chosen to calibrate to the most senior tranche.[17] Depending on our disposition toward the standard model, we can see such a situation as either an arbitrage opportunity (though one that is difficult to realize) or a failure of the model framework.

Beyond improving our ability to calibrate the model to market prices, the base correlation approach also provides a more natural framework to translate calibrated correlations to other tranches or CDOs. For instance, suppose we are pricing a 2-8% tranche on the North American index. Under the implied correlation framework, we would have correlation levels for the standard 0-3%, 3-7%, and 7-10% tranches. Since each correlation is associated with both an attachment and detachment point, it is difficult to see how we might interpolate across the implied correlation levels. Under the base correlation framework, we associate correlation levels with the sequence of first loss tranches: 0-3%, 0-7%, and 0-10%. Since these correlations depend only on a detachment point, it is natural to interpolate to obtain correlations for 0-2% and 0-8% tranches, and then to use these tranches to value the 2-8%.

This same interpolation approach can be utilized to apply a correlation curve calibrated to the index tranches to price tranches on a bespoke CDO portfolio. As mentioned before, there is a basis risk here, though at present, there is little information to quantify this. One other issue that arises is how to normalize the correlation curve. For instance, if the bespoke CDO portfolio has an expected loss that is twice that of the index, then it may not make sense to use the same correlation to price similar tranches on the respective portfolios.

JP Morgan[18] proposes normalizing the credit correlation according to multiples of tranche expected loss, with the base correlation associated precisely with the portfolio expected loss indicated as the At-the-Money (ATM) correlation. Thus, to price the 0-6% tranche of the bespoke CDO, we would

[17] St. Pierre et al (2004).
[18] Ahluwalia, McGinty, and Beinstein (2004).

use the base correlation associated with the 3% detachment point on the index. Similarly, if we wish to assess whether a particular index tranche is expensive relative to historical levels, we would compare the base correlation today with the historical base correlation corresponding to a detachment point at the same multiple of expected loss. The JP Morgan study demonstrates that this approach does explain much of the difference between the correlation profiles of the North America, Europe, and Japan indices, though it recognizes that an examination over a longer historical period will be necessary to be more conclusive. Interestingly, this approach does not appear to explain correlations on the newer high yield indices, where correlations are significantly higher, even accounting for the higher expected loss.[19]

4.3 Hedging conventions

An important consequence of how we represent the correlation structure is how we utilize the model to compute hedge ratios. Simply put, we compute the hedge ratio for a tranche by shifting the spread on one or more of the underlying names, and then revaluing the tranche, holding all other pricing parameters equal. There is a difference, however, between holding implied correlations constant, and holding the base correlation curve constant, meaning that the hedge ratios that result will depend on which representation of correlations is used.

A second issue regarding how the model is used to compute hedges is how granular the hedges should be. It has become common practice, particularly with tranches on the indices, to describe the leverage of a tranche by providing the tranche's delta to the index, that is, the tranche's sensitivity to a parallel move in the spreads of all of the index names. In fact, many dealers now quote the index delta when they quote tranche pricing, and provide preferential pricing if the counterparty is willing to "exchange delta," or take on the other side of the dealer's index hedge, at the same time they enter a position on the tranche. Clearly, there is a risk to hedging only against a parallel spread move, when the same overall move in the index could be achieved by larger moves in a smaller number of names.

If we take the opposite extreme, and examine spread moves in only one name at a time, we see that the resulting hedge ratios depend significantly on the spread level of each name.[20] Moreover, the names with the largest hedge ratios can vary depending on the tranche attachment point, particularly when there is a significant dispersion of spreads in the portfolio. Thus, the practice of

[19] Mahadevan, Polanskyj, and Ambardar (2004).
[20] Mina and Stern (2003).

Table 2
Pricing on North America Investment Grade index tranches. September 10, 2004

Tranche	Upfront (%)	Running (bp)
0-3%	38.9	500
3-7%	-	266
7-10%	-	106
10-15%	-	39
15-30%	-	12

Source: Mahadevan, Polanskyj, and Ambardar (2004).

hedging only with a position in the index (that is, a set of equally weighted positions on the underlying index names) might not be effective when the spreads on the index names are diverse. In practice, with spreads tight and compressed today, this issue is of only marginal importance. A further examination under stressed spread scenarios would be prudent.

5 Example

To illustrate the effects of some of the implementation and application choices, we provide an example using tranches on the North American investment grade index (CDX.NA.IG.3). The index is composed of 125 investment grade credits. For each of the underlying credits, we calibrate default probabilities to the five-year credit default swap spread as of September 10, 2004, as provided by Mark-It Partners. On this day, the average five-year spread across the index credits is 57 basis points; the tightest spread in the index is 17 basis points (AIG), while the widest is 493 basis points (Bombardier). We will examine the five standard tranches on the index, using the prices in Table 2. Throughout the example, we will assume a recovery rate of 40% for all names and a risk-free annual discount rate of 4%.

Our first exercise is to investigate the effect of the implementation choices described in Sections 3.2 and 3.3. The two choices – using individual spreads or the portfolio average, and using the fully granular model or the large pool approximation – give us four distinct model implementations. For this exercise, we assume in all cases that defaults occur at the end of each payment period.

Under each of the four model implementations, we calibrate a curve of base correlations to the tranche prices. Recall that base correlations are the correlations used to price of a succession of first loss tranches (0-3%, 0-7%, and so on) in a way that is consistent with the market pricing of the actual tranches. The four base correlation curves are presented in Table 3. We see that the correlations calibrated using the average spread implementations are slightly less than those calibrated with the individual spread implementations. For the first loss (0-3%) tranche, this result is consistent with our intuition. The first loss tranche (which, assuming 40% recovery, is completely exhausted by just seven defaults) is mostly influenced by the names with the widest spreads (such as Bombardier at 493bp). Other things being equal, replacing the handful of such names with the overall benign average spread of 57bp, reduces the expected loss on the 0-3% tranche. Recalling that expected losses on equity tranches are inversely related to correlation, we must therefore decrease correlation to compensate for the reduction in expected loss to recover the market price.

By contrast, the calibrated correlations rise when we move from the granular to the large pool model. Recall that the large pool model has the same correlation structure as the granular, but that given the common factor return, it is assumed that there is no conditional variance in the portfolio. Thus, other things being equal, the portfolio loss volatility under the large pool model will be lower. To compensate for this lower volatility, we calibrate to a higher correlation.

While interesting, the differences in levels of the calibrated correlations are not of themselves crucial, as long as each of the four model implementations can in fact be calibrated to the current market. What is more relevant are consequent outputs and applications of the model. As we have discussed, a common practice is to hedge index tranches with a position in the index itself. Once we have calibrated our model to prevailing tranche prices, we can compute the mark-to-market impact on the tranches and the index of a change in the underlying credit spreads, assuming correlation stays constant. In our example, we examine a shift of one basis point in each of the underlying index names. We then define a tranche's hedge ratio as the size of an index position required to offset the tranche mark-to-market change under our one basis point shift. The hedge ratios under each of the four model implementations are displayed in Table 4. Here, we see that the granular and large pool models produce almost identical hedge ratios, and that the two spread implementations produce discrepancies of five to ten percent.

Our second exercise is to investigate the effect of the assumed default timing. We return to the granular model using the individual spreads, and compare our previous results, where defaults were assumed to occur at the end of each payment period, to the case where defaults occur at the

Table 3

Base correlations (in %) for different modeling choices

Tranche	Granular Ind. spreads	Granular Avg. spread	Large pool Ind. spreads	Large pool Avg. spread
0-3%	19.9	18.5	23.0	21.5
3-7%	29.2	27.8	30.8	29.4
7-10%	33.0	31.9	34.3	33.2
10-15%	40.7	40.0	41.6	40.9
15-30%	61.4	61.1	61.6	61.7

Table 4

Hedge ratios for different modeling choices

Tranche	Granular Ind. spreads	Granular Avg. spread	Large pool Ind. spreads	Large pool Avg. spread
0-3%	12.50	13.08	12.38	13.06
3-7%	3.82	3.77	3.86	3.78
7-10%	0.97	0.91	0.97	0.90
10-15%	0.46	0.41	0.46	0.43
15-30%	0.20	0.18	0.22	0.21

Table 5
Effect of default timing on correlations and hedge ratios

	Correlation (%)		Hedge ratio	
Tranche	End	Beginning	End	Beginning
0-3%	19.9	21.4	12.50	12.75
3-7%	29.2	31.3	3.82	3.88
7-10%	33.0	35.7	0.97	0.98
10-15%	40.7	44.4	0.46	0.46
15-30%	61.4	66.7	0.20	0.20

beginning of each payment period. The results for base correlation and hedge ratio appear in Table 5. As we discussed earlier, earlier defaults lead to lower valuations for the protection seller. To compensate for this, then, on first loss tranches, we expect to calibrate to a higher correlation. This is in fact what we observe, though the differences are minor, and the discrepancies in hedge ratios are comparable to those we saw earlier from switching to the large pool model.

Our last exercise is to examine the impact of a switch in correlation convention. In Table 6, we display the implied (or compound) correlations, and resulting hedge ratios, that we calibrate to the same set of tranche prices. By construction, the results for the 0-3% tranches are identical, and not surprisingly, the implied correlation for the second loss tranche is anomalous. Of more concern are the hedge ratios, where we see discrepancies of a factor of three for most tranches. Thus, the choice of correlation convention has significantly more practical impact than any of the modeling issues we examined previously.

6 Conclusion – what questions should we be asking?

The single-factor normal copula model has taken hold as the standard pricing framework for synthetic CDOs. It is premature, however, to state that this is more than a temporary phenomenon. Certainly, there are alternate models, but just as importantly, as we have surveyed in this article, there are numerous implementation alternatives, even within the standard framework. The resolution of these alternatives, and ultimately the longevity of today's standard model, still depends on the model's performance through a variety of market scenarios. Fortunately, the existence of an active

Table 6

Effect of correlation convention on correlations and hedge ratios

	Correlation (%)		Hedge ratio	
Tranche	Base	Implied	Base	Implied
0-3%	19.9	19.9	12.50	12.50
3-7%	29.2	5.2	3.82	11.07
7-10%	33.0	19.7	0.97	3.98
10-15%	40.7	21.8	0.46	1.70
15-30%	61.4	30.4	0.20	0.52

market in synthetic CDO tranches provides for a true test of the model. This is in stark contrast to the discussion of credit portfolio models for economic capital, where there is no market, and little empirical data, with which to adequately judge the model assumptions.

We conclude with some thoughts on some of the dimensions along which the model and its implementations will be judged.

Link to the empirical data. Do any of the empirical methods of estimating correlation provide reliable results, or is correlation in the standard model purely a technical factor, unrelated to anything that is truly observable? It will be difficult to continue to use the standard model if there is no empirical way to at least estimate a range for the correlation parameter.

Calibration of the full range of prices. As discussed before, there are occasions where no level of correlation can explain tranche pricing. That this occurs frequently under the simple implied correlation representation gave rise to the discussion of base correlations. Base correlations, while providing for broader coverage, can fail at times to describe a set of market prices that do not violate any strict arbitrage constraints. Should these situations become common, unless they can be resolved through a judicious set of implementation choices, it will call into question either the single-factor structure, or the normal copula, or both.

Stability of calibrated parameters. Even if the model can describe market prices consistently, it will prove unsuitable if the correlation parameters are not stable through time, or if movements in the correlation parameters cannot be rationalized with broader events. Across the capital structure of a given CDO, a single correlation is not sufficient to price all the tranches, and a correlation structure exists. While an indication that the standard model does not perfectly describe prices, the correlation structure is not in itself a problem. However, if we are unable to rationalize the shape of the structure, or if the structure changes in ways that are not at least intuitive in hindsight, we may judge the standard model to be deficient.

Hedge performance. As the market for the most liquid tranches evolves, pricing levels will derive from supply and demand effects, rather than from model outputs. At this point, the primary use of the pricing models will be for hedging positions in the tranches. The ability of a specific model implementation to describe an adequate hedging program will then be the ultimate model test. St. Pierre et al (2004) investigate the historical performance of index hedges under the implied and base correlation representations, and conclude that the base correlation framework produces lower, though still at times unacceptably high, tracking error. Extensions of this study should seek to cover a longer historical period, hedging strategies other than the simple index hedge, and a range of collateral portfolios (for instance, a diverse high yield portfolio in addition to the typical compressed investment grade).

References

Ahluwalia, R., McGinty, L., and Beinstein, E. (2004). A relative value framework for credit correlation, Credit Derivatives Strategy, JP Morgan, April 27.

Basel Committee on Banking Supervision (2004). International convergence of capital measurement and capital standards: A revised framework, Bank for International Settlements, June.

CSFP, CreditRisk+: A credit risk management framework, Credit Suisse Financial Products.

de Servigny, A. and Renault, O. (2003). Correlations evidence, *Risk*, July: 90–94.

Demey, P. et al (2004). Maximum Likelihood Estimation of Default Correlations. *Risk*, November: 104–108.

Duffie, D. (2004). Time to adapt copula methods for modelling credit risk correlation, *Risk*, April: 77.

Finger, C. C. (1999). Conditional approaches for CreditMetrics portfolio distributions, *CreditMetrics Monitor*, April: 14–33.

Gordy, M. (2000). Credit VaR models and risk-bucket capital rules: A reconciliation, *Proceedings of the 36th Annual Conference on Bank Structure and Competition*, Federal Reserve Bank of Chicago, pp. 406–417.

Gregory, J. and Laurent, J.-P. (2003). I will survive, *Risk,* June: 103–107.

Gupton, G., Finger, C., and Bhatia, M.(1997). *CreditMetrics™—Technical Document*. Morgan Guaranty Trust Co.

Kealhofer, S. (2001). Portfolio Management of Default Risk, KMV Corporation.

Li, D. X. (1999). The valuation of basket credit derivatives, *CreditMetrics Monitor*, April: 34–50.

Li, D. X. (2000). On default correlation: a copula function approach, *Journal of Fixed Income,* **9**(4): 43–54.

Mahadevan, S., Polanskyj, P., and Ambardar, A. (2004). Streetwise correlation — taking the high road, Credit Derivatives Insights, Morgan Stanley, September 10.

Martin, R., Thompson, K., and Browne, C. (2001). Taking to the saddle, *Risk*, June: 91–94.

Mashal, R. and Naldi, M. (2002). Extreme events and default baskets, *Risk*, June: 119-–122.

McGinty, L. et al (2004). Credit correlation: A guide, Credit Derivatives Strategy, JP Morgan, March 12.

Merino, S. and Nyfeler, M. (2002). Calculating portfolio loss, *Risk*, August: 82–86.

Merton, R. (1974). On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* **29**: 449–470.

Mina, J. and Stern, E. (2003). Examples and applications of closed-form CDO pricing, *RiskMetrics Journal*, Fall: 5–34.

O'Kane, D. et al (2003). The Lehman Brothers guide to exotic credit derivatives, Risk Waters Group.

St. Pierre, M. et al (2004). Valuing and hedging synthetic CDO tranches using base correlations, Bear Stearns, May 17.

Vasicek, O. (1997). The loan loss distribution, KMV Corporation.

Watkinson, L. and Roosevelt, D. (2004). The layman's guide to implied correlation, Product Note, Morgan Stanley, May 15.

Wilson, T. (1997). Credit Portfolio Risk, *Risk*, October: 61–72.