Squaring factor copula models

Tight spreads in the credit markets have forced investors to turn to innovative structures in their search for yield. One such structure is the synthetic CDO of CDO tranches, also known as CDO^2. Prasun Baheti, Roy Mashal, Marco Naldi and Lutz Schloegl introduce this contract, and present a quasi-analytical framework for the valuation and hedging of this and other similar 'squared' products

In the past few years, dynamically hedged synthetic collateralised debt obligation (CDO) tranches have had an impact on the credit derivatives market that is difficult to overstate. They have increased the liquidity and changed the dynamics of the default swap market via the 'synthetic bid' for credit, and taken the CDO concept beyond the realm of structured finance into the derivatives arena. Nevertheless, synthetic CDOs are susceptible to arbitrage spreads just like their cashflow counterparts. Given the relentless spread tightening since the end of 2002, it is perhaps not surprising that it has become more difficult to obtain the yields investors had become used to via standard synthetic tranches.

The CDO^2 concept addresses this difficulty by providing an additional layer to the capital structure. In a CDO^2, a portfolio of synthetic CDO tranches is itself tranched into so-called super tranches. This introduces quite a few new variables into the structuring equation. Not only does the composition of the underlying portfolio of individual tranches have to be determined, but also their joint characteristics can be tailored by varying the degree of overlap between the reference credits in the individual pools. Moreover, extra flexibility is provided to the structure by the ability to choose the level of subordination and the width of the super tranche. These additional degrees of freedom make it possible to further fine-tune the risk/return profile of a loss tranche, for example, achieving high leverage while controlling the exposure to idiosyncratic default risk.

CDOs of CDOs have been used for some time in the cashflow world. However, the terms of the purely synthetic CDO^2 we are discussing here are somewhat different, so that it is worth clarifying the structure and the notation, as we do in the following section. Then we turn to the core of this article, and show how to extend the popular factor copula approach to derive quasi-analytical solutions for valuation and hedging of synthetic CDO^2 tranches. Finally, we briefly discuss further applications of the proposed methodology.

The CDO^2 structure

Figure 1 shows a stylised representation of a particular CDO^2 structure, that is, a mezzanine tranche referencing a portfolio of mezzanine tranches.

The fundamental inputs to a CDO^2 trade are a large pool of individual credits. The main constraint on the size of the pool is the number of credits that can be dynamically hedged due to their liquidity in the single-name default swap market. A typical pool might consist of 250–350 credits. We denote the total number of credits in the pool by M. The credits are assigned to different so-called 'mini-portfolios', and we denote the total number of mini-portfolios by N. It is important to note that a given credit can appear in more than one portfolio, and that the weight of a particular credit is specific to each mini-portfolio.

The percentage weight of credit j in mini-portfolio k is denoted by $w_{j,k}$. These weights are constrained to be non-negative and add up to 100% within each mini-portfolio, that is:

$$w_{j,k} \geq 0, \quad j = 1,2,\ldots,M, \quad k = 1,2,\ldots,N$$
$$\sum_{j=1}^{M} w_{j,k} = 1, \quad k = 1,2,\ldots,N$$

The matrix $(w_{j,k})$ is the 'population matrix' of the trade; it encapsulates the information about issuer concentrations and the overlap between different mini-portfolios. As the underlying risk sources of the CDO^2, we consider a tranche linked to each mini-portfolio. The easiest way to describe the kth mini-tranche is by its percentage subordination U^k, and its percentage width V^k. Note that so far we have not fixed any absolute notional amounts. We denote the absolute notional of the kth mini-tranche by N^k, so that the total notional of the corresponding mini-portfolio is N^k/V^k. This will become relevant when we describe how the individual credit losses flow through to the super tranche.

The portfolio underlying the super tranche consists of the N mini-tranches, and its total notional is therefore $\Sigma_{k=1}^{N} N^k$. The super tranche itself is described by its percentage subordination U^{st} and percentage width V^{st}. We also refer to the portfolio of mini-tranches as the super portfolio.

Let us now consider how the credit losses in the pool ultimately flow to the super tranche. Suppose that credit j defaults with a recovery rate of R. The loss to the kth mini-portfolio is then given by $(1 - R)w_{j,k}N^k/V^k$, and the percentage notional lost is equal to $(1 - R)w_{j,k}$. Once the cumulative percentage loss in the kth mini-portfolio is greater than the mini-tranche subordination, the super portfolio starts to take losses and the subordination of the super tranche is reduced. The seller of super tranche protection is obliged to make protection payments once the tranche has been eaten into, just as in a standard synthetic CDO tranche. Similarly, the contractual spread paid to the protection seller is based on the outstanding notional of the super tranche. Note that the synthetic super tranche structure is different from traditional cash CDOs of CDOs in that the mini tranches are not 'physical' assets as such. They have no premium associated with them and only serve to define the subordination structure to resolve losses.

As we explain in greater detail in the next section, valuation can be performed using the concept of 'tranche curve'. To construct this curve, we need to define the cumulative percentage loss of the super tranche L^{st} up to a given time horizon. If the cumulative percentage loss to the kth mini-port-

1. Mezzanine tranche of mezzanine tranches

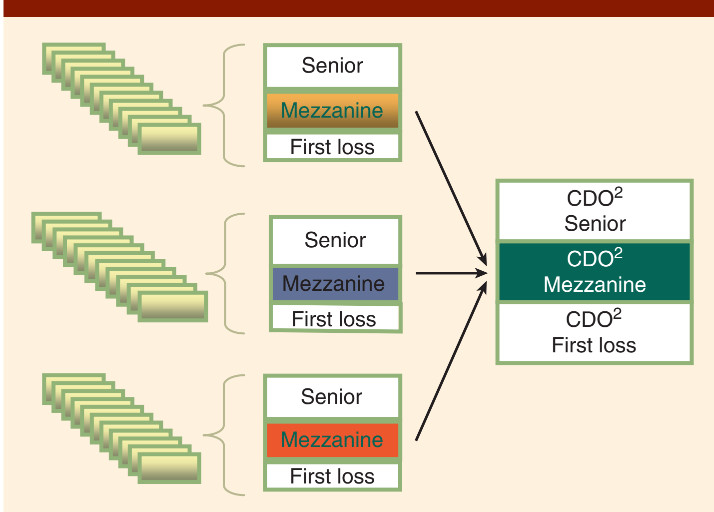

2. Calculating the joint loss distribution of two overlapping portfolios

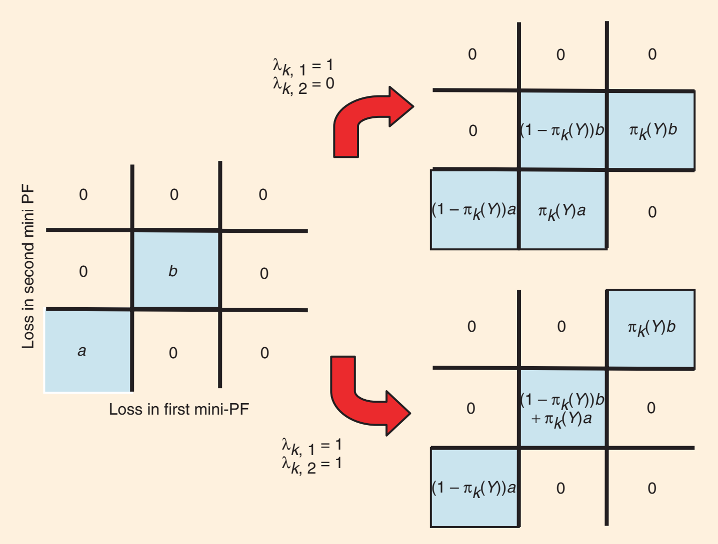

folio is L^k, then the cumulative percentage loss to the kth mini-tranche is:

$$L^{mt(k)} = \frac{\left[L^k - U^k\right]^+ - \left[L^k - \left(U^k + V^k\right)\right]^+}{V^k}$$

The cumulative percentage loss to the super portfolio is therefore:

$$L^{sp} = \frac{\sum_{k=1}^{N} N^k L^{mt(k)}}{\sum_{k=1}^{N} N^k}$$

and the cumulative percentage loss to the super tranche is:

$$L^{st} = \frac{\left[L^{sp} - U^{st}\right]^+ - \left[L^{sp} - \left(U^{st} + V^{st}\right)\right]^+}{V^{st}}$$

While the loss on a standard CDO tranche is a call spread on the underlying portfolio loss, the equations above show that the super tranche loss is given by compound options, which are effectively calls on a basket of vanilla call spreads.

A quasi-analytical valuation model

At the core of any CDO pricing model is a mechanism for generating dependent defaults. Factor copula models describe default as an event generated by a latent variable – generally interpreted as asset return – falling below a specified threshold, which is in turn calibrated to observable credit default swap spreads of the reference credit. As first noted by Li (1999), the dependence among the default times of different names is naturally determined by the dependence structure (also known as the copula) of these latent variables. To ease the computational effort required for a quasi-analytical implementation of the model, the latent variables are generally represented by a low-dimensional linear factor structure.

The simplest latent variables model combines a Gaussian copula with a single-factor specification. In what follows, we will use these assumptions for notational convenience, but the reader should bear in mind that it is straightforward to adapt everything that follows to different distributional assumptions and/or to a multi-factor setting, although both of these variations will generally come at a cost in terms of computational speed.[1] Suppose the default-triggering latent variable for name j, X_j, is driven by a common factor Y, and an idiosyncratic term E_j:

$$X_j = \beta_j \times Y + \sqrt{1-\beta_j^2} \times E_j$$

Here the variables Y, E_j, $j = 1, 2, \ldots, M$ are taken to be independent standard normal random variables, so that the latent variables X_j are jointly normal with an $M \times M$ correlation matrix given by:

$$\left(C_{j,l}\right) = \left(\beta_j \times \beta_l\right)$$

Within the one-factor Gaussian framework, the dependence structure of the latent variables is fully specified by a vector of betas. Given a particular realisation of the common factor, the probability that the jth credit defaults is now given by:

$$\pi_j(Y) = P\left[X_j < D_j \middle| Y\right] = N\left(\frac{D_j - \beta_j \times Y}{\sqrt{1-\beta_j^2}}\right)$$

where D_j represents the default threshold for a given time horizon calibrated to the jth name's credit curve.

The simple specification of this model implies that, conditional on the realisation of the common factor, the M individual credits are independent. When pricing a standard CDO, this conditional independence greatly facilitates the calculation of the conditional loss distribution of the reference portfolio and, therefore, of any pre-defined loss tranche. This can be done using transform methods as described in Gregory & Laurent (2003) or by means of a fast recursive algorithm first proposed by Andersen, Sidenius & Basu (2003).

In a CDO^2, however, since some of the credits may belong to several mini-portfolios, the loss distributions of the mini-tranches need not be conditionally independent even if the defaults of the individual credits are. The possibility of overlapping credits in the reference mini-portfolios significantly complicates the task of recovering the conditional joint loss distribution of the mini-tranches, which is in turn necessary to calculate the conditional loss distribution of the super tranche. We now show how to overcome this obstacle by means of a multivariate recursive procedure.

First, let us associate each credit with the number of loss units that its default would produce in each of the mini-portfolios. The representative element of the 'loss matrix' ($\lambda_{j,k}$) indicates the integer number of loss units in mini-portfolio k due to the default of name j.

Next, we construct an N-dimensional hyper-cube whose kth side consists of all possible loss levels for the kth mini-portfolio, that is, $(0, 1, \ldots, \Sigma_{j=1}^{M} \lambda_{j,k})$. For ease of explanation, and without loss of generality, we consider here a two-dimensional example ($k = 2$). In this case, our hyper-cube is simply a matrix ($Z_{v1,v2}$) where we can store the conditional joint distribution of the two mini-portfolios. For example, we store in $Z_{3,5}$ the probability of jointly having three loss units in the first mini-portfolio and five in the second. To obtain the correct set of joint probabilities, we first initiate each state (recursion step $j = 0$) by setting:

$$Z^0_{v_1,v_2} = 1 \quad \text{if} \quad v_1 = 0 \quad \text{and} \quad v_2 = 0$$
$$Z^0_{v_1,v_2} = 0 \quad \text{otherwise}$$

Then, we feed the M credits, one at a time, through the following recursion:

$$Z^j_{v_1,v_2} = \left(1-\pi_j(Y)\right) \times Z^{j-1}_{v_1,v_2} + \pi_j(Y) \times Z^{j-1}_{\left(v_1-\lambda_{j,1}\right),\left(v_2-\lambda_{j,2}\right)} \quad \text{if} \quad v_1 \geq \lambda_{j,1} \cap v_2 \geq \lambda_{j,2}$$
$$Z^j_{v_1,v_2} = \left(1-\pi_j(Y)\right) \times Z^{j-1}_{v_1,v_2} \quad \text{otherwise}$$

where $\pi_j(Y)$ indicates the conditional probability that issuer j defaults. Each credit can either survive, and every state then 'keeps' its position, or default, in which case every state 'moves' in the direction $[\lambda_{j,1}, \lambda_{j,2}]$.

Figure 2 shows how to update the joint loss distribution when we add the kth credit, whose default is assumed to produce exactly one loss unit. Two possible updates of the matrix are shown, the first one corresponding to the case where the kth credit belongs to the first mini-portfolio only,

[1] *It is well known that a Gaussian copula model calibrated to historical correlations cannot replicate the prices of tranches of standard portfolios such as CDX and i-Traxx. To anchor the valuation of synthetic CDO² to market-implied default correlations, one can apply the methodology described in this article to an underlying factor copula model that is able to explain these observable market prices*

the second one corresponding to the case where the kth name appears in both mini-portfolios.

After including all the issuers, we set:

$$\left(Z_{v_1,v_2}\right)=\left(Z^M_{v_1,v_2}\right)$$

The matrix $(Z_{v1,\,v2})$ now holds the joint loss distribution of the two mini-portfolios conditional on the realisation of the market factor. It is then straightforward to recover the conditional joint distribution of losses on the mini-tranches, the conditional loss distribution of the super portfolio and the conditional loss distribution of the super tranche.

Once we know how to calculate the loss distribution of the super tranche for a given realisation of the common factor, it is straightforward to integrate against its probability distribution and recover the unconditional loss distribution of the super tranche. Repeating the entire procedure for a grid of horizon dates, and interpreting the expected percentage loss up to time t as a cumulative default probability, we can price a tranche using exactly the same analytics as in a single-name default swap. More precisely, define the 'survival probability' of the super tranche up to time t as:

$$Q^{st}(t)=1-E\left[L_t^{st}\right]$$

Then the two legs of the CDO^2 swap can be priced using:

$$PV(\mathit{protection\ leg})=N^{st}\sum_{i=1}^{S}B(s_i)\left(Q^{st}(s_{i-1})-Q^{st}(s_i)\right)$$

$$PV(\mathit{premium\ leg})=c^{st}N^{st}\sum_{i=1}^{T}\Delta_i Q^{st}(t_i)B(t_i)$$

where c^{st} is the coupon paid on the super tranche, N^{st} is the notional of the super tranche, t_i, $i = 1, 2, \ldots, T$ are the coupon dates, Δ_i, $i = 1, 2, \ldots, T$ are accrual factors, s_i, $i = 1, 2, \ldots, S$ discretise the timeline for the valuation of the protection leg, and $B(t)$ is the risk-free discount factor for time t.

To summarise this section, we have seen how to extend a factor copula model for pricing the super tranche of a CDO^2 after specifying:

- The CDO^2 structure, that is, the population matrix $(w_{j,k})$, the mini-tranches triple (U^k, V^k, N^k) and the super tranche triple (U^{st}, V^{st}, N^{st}).
- The issuer curves of the underlying credits (used to calibrate the thresholds D_j, $j = 1, 2, \ldots, M$ for each horizon date).
- The (risk-neutral) dependence structure of the default-triggering latent variables.

Discussion

The recursive algorithm we have described in the previous section is subject to a dimensionality problem: as we increase the number of mini-portfolios, the memory requirements grow exponentially. For example, a squared structure with five mini-portfolios, each referencing 100 names and detaching after 10 defaults, will require the calculation of joint loss probabilities distributed over $11^5 = 161{,}051$ states. Adding a sixth mini-portfolio will increase the state space to $11^6 = 1{,}771{,}561$ elements. Moreover, the computational requirements are quite sensitive to the number of states in each of the mini-tranches' loss distributions. For example, an equally weighted mini-portfolio of 100 assets with equal recoveries generates a loss distribution with 101 states. If we now consider the same portfolio, but let one of the assets have half the notional of the others, then the number of possible losses immediately rises to 200. Notice also that, in most cases, our algorithm can be further optimised by means of a compactification of the state space. Consider, for example, the case where one of the mini-tranches is a 7–10% mezzanine. Then, once all the names in the corresponding mini-portfolio have been entered into the recursion, we no longer need to pay attention to states where this tranche loses less than 7%.

In terms of speed, pricing a super tranche referencing five mini-tranches will take four to eight minutes on a standard PC, depending on the degree of overlap, the weight matrix, the recovery assumptions and the coarseness of the time discretisation. A naive calculation of the deltas, implemented by perturbing the input curves and repricing the super tranche, will therefore be very time-consuming.

Although the limitations of our methodology can be partially addressed by allocating extra computing power, one should also consider using 'smart' Monte Carlo techniques when dealing with high-dimensional CDO^2. Recent contributions in this direction include Joshi & Kainth (2004), Joshi (2004), as well as Glasserman & Li (2004), who propose an importance sampling method that is specifically tailored to factor copula models.

The 'dimensionality curse' described above is less of a problem when we consider some low-dimensional variations of the synthetic CDO^2 structure. One example is the recently introduced 'basket of baskets'. This is simply a super-nth-to-default basket referencing a portfolio of mini-m_ith-to-default baskets. The super basket is triggered when n of the underlying baskets are triggered, and the ith mini-basket is triggered when the m_ith name in the ith reference portfolio has a default event. As an example, consider a second-to-default super basket on 10 first-to-default mini-baskets. This structure will be triggered once two of the 10 mini-baskets have been triggered. The complete state space of the joint distribution of the 10 mini-baskets is of size 2^{10}, allowing for fast pricing and risk calculations. Similarly, it is possible to combine tranches and baskets in the same squared instrument. Another recent innovation in the market is a loss tranche referencing a portfolio of m_ith-to-default baskets. The pricing and hedging of these contracts can generally be handled very efficiently using the algorithm described above.

Finally, our methodology can also be applied to account for counterparty risk when pricing multi-name credit derivatives. In particular, the recursive algorithm described earlier can easily be used to construct the joint distribution of, say, the loss on a CDO tranche and the default of the protection seller, while taking into full account the dependence between the default time of the latter and the default times of the issuers in the reference portfolio of the CDO. This dependence is crucial, since the states in which the counterparty defaults are likely to be the states in which it would have to protect losses on the tranche.

Summary

The market for exotic correlation products such as CDO^2 has grown significantly over the past couple of years. At the same time, factor copula models have become the *de facto* standard for pricing plain-vanilla synthetic CDOs. In this article, we have shown how to extend this class of models to price and hedge CDO^2 structures. Although the workload of the proposed algorithm grows exponentially with the number of underlying sub-tranches, its applicability is immediate when we consider a few low-dimensional variations of the synthetic CDO^2 structure that have recently appeared in the credit derivatives markets. ■

Prasun Baheti is an associate, Roy Mashal is a director, Marco Naldi is senior vice-president and Lutz Schloegl is a director at Lehman Brothers. **They would like to thank two anonymous referees for their suggestions. The usual disclaimer applies.** **Email: prbaheti@lehman.com, rmashal@lehman.com, mnaldi@lehman.com, luschloe@lehman.com**

REFERENCES

Andersen L, J Sidenius and S Basu, 2003
All your hedges in one basket
Risk November, pages 67–72

Glasserman P and J Li, 2003
Importance sampling for portfolio credit risk
Working paper, Columbia University

Gregory J and J-P Laurent, 2003
I will survive
Risk June, pages 103–107

Joshi M, 2004
Applying importance sampling to pricing single tranches of CDOs in a one-factor Li model
Working paper, available at *www.quarchome.com*

Joshi M and D Kainth, 2004
Rapid computation of prices and Greeks for nth to default swaps in the Li model
Quantitative Finance 4(3), pages 266–275

Li D, 1999
On default correlation: a copula function approach
Journal of Fixed Income 9(4), pages 43–54