Loss in translation

Ben De Prisco, Ian Iscoe and Alex Kreinin introduce a new analytical approach for valuing synthetic collateralised debt obligations. The approach differs from current analytical approaches by focusing on the tranche's loss distribution directly, as opposed to a preliminary determination of the entire pool's loss distribution. Using a compound Poisson approximation, the pool's losses are translated directly into the tranche's losses. The result is a more efficient valuation – both in terms of memory and speed – for investors owning few tranches in the overall pool

In principle, a synthetic collateralised debt obligation (CDO) can be valued by Monte Carlo methods. This is perhaps the most common approach used by practitioners today. The problem is that Monte Carlo methods usually take a significant amount of time to achieve the required accuracy. Hence, while such methods may lend themselves to pricing and structuring, they are not appropriate for risk management where simulation and stress testing on tranche values are required. In fact, nesting a second level of simulation, for pricing, within the risk management simulation represents a performance challenge.

For this reason, many papers have been produced over the past few years describing analytical approaches to the valuation of CDOs. Although such approaches differ slightly based on their underlying techniques and assumptions, almost all approaches to date are similar in having focused on first determining the reference pool's entire loss distribution before valuing the CDO tranches. Such approaches have both memory and performance deficiencies, especially if investors only own a single tranche of the CDO. Our approach addresses these memory and performance issues by focusing directly on the valuation of individual tranches, bypassing the need to determine the complete pool-loss distribution. The end result is a much faster algorithm in the common case where investors own only one or two tranches in a given CDO.

In brief, our approach uses a single Gaussian credit driver to condition pool default behaviour. Using a compound Poisson distribution for estimating conditional pool losses, we are then able to focus strictly on the portion of the pool-loss distribution within the given tranche's attachment and detachment points. In this respect, the conditional pool losses are translated directly into tranche losses, allowing us to efficiently arrive at expected tranche losses. Numerical integration is then used on the single credit driver for unconditioning the expected tranche losses. It should be noted that while a compound Poisson distribution has been used in actuarial science for many years to estimate claim distributions (see, for example, section 6.5 in Panjer & Willmot, 1992), to our knowledge our approach represents the first practical application of this distribution to a credit-related pricing problem involving pool default-loss distributions.

We summarise our assumptions and computational steps in our approach, and indicate some possible extensions, in the box below.

The rest of this article is organised as follows. In the next section, we describe the pricing of CDO tranches, by deriving quantitative versions of their balance equations. Following that, we describe the dependence structure used for modelling correlated default events. To simplify the notation, we restrict attention to a single-factor Gaussian model (extensions to a multi-factor Gaussian model are indicated in the box). Then we present our new analytical method for valuing CDOs. Following that, we present the results of some numerical experiments that compare the accuracy and performance of our new method with conventional analytical and Monte Carlo methods. Finally, we provide an overview of some other analytical methods for valuing synthetic CDOs based on the same model.

Summary of method and extensions

☐ **Assumptions about the pool**
1. Recovery rates are deterministic but can vary from name to name, as can the (unadjusted) notionals.
2. Default correlations can vary from name to name.
3. Default term structure can vary from name to name.

☐ **Assumptions about the default model**
1. Defaults correlated through a single Gaussian factor (credit driver).
2. Conditional independence of defaults, given a credit-driver scenario.

☐ **Assumptions about the CDO contract**
1. There is no replacement of reference names upon default.
2. Compensation for losses is only made at premium dates, so losses are present-valued from those dates and not at the default times.

☐ **Summary of computations**
1. At each time step, for each credit-driver scenario, the pool-loss distribution is approximated, conditional on the scenario, by a (compound) Poisson distribution, using the proposition or theorem, as appropriate.
2. Conditional expected tranche losses are calculated (with the aid of the fast Fourier transform, in the compound case), using (14) or (17), as appropriate.
3. Unconditional results are obtained by numerical integration (8) against the distribution of the credit driver.
4. Unconditional tranche value or premium calculated algebraically by combining the results of the previous step across time periods, as in (5) or (6).

☐ **Extensions**
1. Distributions other than the Gaussian one can be used for the credit driver. The main concern is the efficiency of the choice of numerical integration procedure that must eventually be applied. In the conditional-independence framework, the computational time is essentially proportional to the number of integration nodes.
2. The extension to a multi-factor Gaussian model (multiple credit drivers) is immediate from the theoretical standpoint. Only the argument of Φ in (7) would change slightly, as the numerical integration would be multi-dimensional (see Iscoe & Kreinin, 2004, for details). However, there are performance implications with any analytical method involving multiple credit drivers. Even with only 14 integration nodes per dimension, the computation time rapidly increases with the dimension, so that with four credit drivers, Monte Carlo pricing becomes competitive. Finding an efficient, Gaussian numerical quadrature in multiple dimensions (beyond two or three) would be a prerequisite to a practical implementation of any analytical method. ■

1. Tranche loss function, $\psi^{(j)}$, $j > 0$, as a function of the pool loss, l

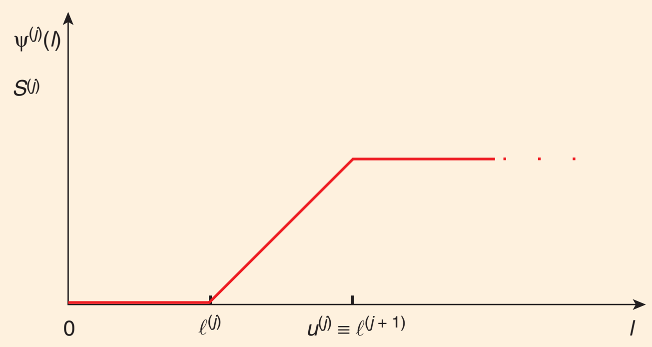

Contract structure and balance equation

A synthetic CDO is a credit derivative whose underlying collateral is a portfolio (or pool) of credit default swaps (CDSs). To offset the pool owner's risk from these default swaps, part of their premiums is allocated to a collection of securities called tranches of the CDO. There is a priority scheme for the tranches to absorb the pool losses, up to a fixed, maximum amount for each tranche. Losses are based on the recovery-adjusted CDS notional values. The buyer of one or more of these tranches sells partial protection to the pool owner, by agreeing to absorb up to the set amount of the pool's losses, in exchange for periodic premium payments.

Let $t_1, t_2, \ldots, t_n = T$ denote the premium dates, with T denoting the maturity of the CDO; t_0 will denote today (it is possible that $t_0 = t_1$, that is, when the contract's effective date is in the past). Let d_i denote the discount factor corresponding to t_i, $1 \le i \le n$.

The pool is comprised of K names. Denote by $N^{(k)}(1 \le k \le K)$ the recovery-adjusted notional associated with each name. In our article, a pool of names having identical recovery-adjusted notionals, identical default correlations and identical risk-neutral default probabilities is termed completely homogeneous. If only the recovery-adjusted notionals are identical, then the pool is termed homogeneous. In all other cases, the pool is termed heterogeneous. We impose the condition that there is no replacement of the underlying reference names for any reason, in particular upon default.

The CDO consists of J tranches, into which the pool's losses are absorbed in accordance with the size and seniority of the tranches. The size of tranche j, $(0 \le j \le J - 1)$, in monetary units, is $S^{(j)} = u^{(j)} - \ell^{(j)}$, where $\ell^{(j)}$ is referred to as the attachment point of the tranche and $u^{(j)} = \ell^{(j+1)}$ is referred to as the detachment point. Denote by $\mathcal{L}_i$ the pool's cumulative losses up to time t_i. Then tranche j absorbs an amount, $L_i^{(j)}$, of those losses:

$$L_i^{(j)} = \min\left(S^{(j)}, \max\left(\mathcal{L}_i - \ell^{(j)}, 0\right)\right) \tag{1}$$

Thus the tranche losses, $L_i^{(j)}$, are non-linear functions, say $\psi^{(j)}$, of the pool losses $\mathcal{L}_i$: $L_i^{(j)}, = \psi^{(j)}(\mathcal{L}_i)$, where:

$$\psi^{(j)}(l) = \min\left(S^{(j)}, \max\left(l - \ell^{(j)}, 0\right)\right) \tag{2}$$

The graph of $\psi^{(j)}$ is sketched in figure 1.

Since our focus is on the valuation of single tranches, we will henceforth omit the superscript j in our notation, and write simply ℓ and u for the attachment and detachment points, S for the tranche thickness, etc.

In general, valuation of a CDO tranche is based on a balance equation that can be expressed, at a high level, as follows:

$$\begin{aligned}&\textit{Expectation of discounted premium payments}\\ &= \textit{Expectation of discounted credit losses}\end{aligned} \tag{3}$$

Denoting risk-neutral expectation by $\mathbb{E}$, the balance equation can be represented quantitatively as:

$$\mathbb{E}\left[\sum_{i=1}^{n} s_i\left(t_i - t_{i-1}\right)\left(S - L_i\right)d_i\right] = \mathbb{E}\left[\sum_{i=1}^{n}\left(L_i - L_{i-1}\right)d_i\right] \tag{4}$$

Since L_i represents the cumulative tranche losses, the difference $L_i - L_{i-1}$ is the losses absorbed by the tranche in the interval $(t_{i-1}, t_i]$. Therefore, the right-hand side of (4) represents the mean discounted losses of the tranche during the lifetime of the contract. The credit risk premium, s_i, paid by the tranche seller, depends on the monetary value of the pool assets for the names surviving up to time t_i, $(1 \le i \le n)$. Therefore, the left-hand side of (4) represents the mean discounted premium payment over the lifetime of the contract.

We can use (4) for pricing the CDO tranche (that is, calculating the required constant premium) by:

$$s = \frac{\sum_{i=1}^{n}\left(\mathbb{E}L_i - \mathbb{E}L_{i-1}\right)d_i}{\sum_{i=1}^{n}\left(t_i - t_{i-1}\right)\left(S - \mathbb{E}L_i\right)d_i} \tag{5}$$

Alternatively, in the case where premiums are already set, we can use (4) for valuing a CDO tranche as the difference between the left- and right-hand sides:

$$V_{buy} = \sum_{i=1}^{n} s_i\left(t_i - t_{i-1}\right)\left(S - \mathbb{E}L_i\right)d_i - \sum_{i=1}^{n}\left(\mathbb{E}L_i - \mathbb{E}L_{i-1}\right)d_i \tag{6}$$

In either case, the problem is reduced to the computation of the mean absorbed losses, $\mathbb{E}[L_i]$. For that, we must specify the default model, which is the topic of the next section.

Conditional independence framework

To calculate the expected values in (5) and (6), we have to specify a default process and explicitly define the correlation structure of the default events.

A multi-factor, Gaussian copula model has become the most popular among practitioners. It was introduced first in the credit risk framework by Vasicek (1987) in the one-factor setting and later generalised by different authors in Andersen, Basu & Sidenius (2003), Gordy & Jones (2003), Hull & White (2003), Iscoe, Kreinin & Rosen (1999), Laurent & Gregory (2003), Schönbucher (2003) and Bielecki & Rutkowski (2002). We specify the default process using this one-factor model in which the correlation of the names is effected through a single, standard Gaussian credit driver. Default events are simulated, independently in each credit-driver scenario, as the first crossing time of a barrier (see Hull & White, 2003, and Iscoe & Kreinin, 2004).

In this framework, the risk-neutral default probabilities describing the default-time distribution of the kth name can be bootstrapped from the credit spreads of CDSs (see, for example, pages 207–208 in Lando, 2004). We denote the unconditional, risk-neutral, default-time (cumulative) distribution of the kth name by $\hat{\pi}_i^{(k)} : \hat{\pi}_i^{(k)} = \mathbb{P}(\tau^{(k)} \le t_i)$.

Conditional on a value x of the credit driver, the risk-neutral default-time distribution is given by:

$$\hat{\pi}_i^{(k)}(x) = \Phi\left(\frac{\Phi^{-1}\left(\hat{\pi}_i^{(k)}\right) - \beta^{(k)}x}{\sigma^{(k)}}\right) \tag{7}$$

for which $\int_{-\infty}^{\infty}\hat{\pi}_i^{(k)}(x)d\Phi(x) = \hat{\pi}_i^{(k)}$. Here Φ denotes the standard normal cumulative distribution function and Φ^{-1} denotes its inverse function. We set $\bar{\pi}_i^{(k)}(x) = 1 - \hat{\pi}_i^{(k)}(x)$.

In our framework, the default events of the names are, by fiat, conditionally independent. (Together with (7), this suffices for the specification of the correlated default model.) This assumption allows us to reduce computations to the case of independent names. Also, the mean tranche loss satisfies the relation:

$$\mathbb{E}\left[L_i\right] = \int_{-\infty}^{\infty}\mathbb{E}_x\left[L_i\right]d\Phi(x) \tag{8}$$

where $\mathbb{E}_x$ denotes mathematical expectation, conditional on $X = x$. Therefore, it suffices to calculate the conditional tranche losses.

Poisson pricing

It was observed earlier, from the formulas (5) and (6), that the basic quantities that must be calculated are $\mathbb{E}[L_i]$. We begin the description of our approach with some general, preliminary results and important observations, which do not require that the pool be large. We will then treat a succession of cases of increasing complexity, starting with the completely homogeneous case and moving towards the heterogeneous case.

By conditioning on a credit-driver scenario, x, we reduce the calculation of $\mathbb{E}L_i$ to $\mathbb{E}_xL_i$: $\mathbb{E}L_i = \int_{-\infty}^{\infty}\mathbb{E}_xL_i d\Phi(x)$. Having made this reduction, we will henceforth suppress the fixed scenario, x, as well as the subscript i (indexing the fixed premium date), from our notation.

Now, $L = \psi(\mathcal{L})$, where the function ψ was defined in (2). So, since $\psi(l)$ is zero for $l \le \ell$ and has constant value, S, for $l \ge u$:

$$\mathbb{E}[L] = \mathbb{E}\left[\psi(\mathcal{L})\right] \tag{9}$$

$$= \sum_{\ell<l<u}[l-\ell]\mathbb{P}(\mathcal{L}=l) + S\left[1-\sum_{l<u}\mathbb{P}(\mathcal{L}=l)\right] \tag{10}$$

The first term in the last line of the derivation is the contribution to the tranche's mean loss from the 'ramp' in the graph of ψ (see figure 1). The second term is the contribution from the 'plateau'. This method of evaluating $\mathbb{E}[L]$ makes its dependence on the (conditional) distribution of L very explicit. It also stresses that it is not necessary to pre-compute the entire pool-loss distribution, since the minimal, pertinent range of values l depends only on the notional values in the pool and on the particular tranche attachment and detachment points, ℓ and u. An expression similar to (10), was used in Anderson, Basu & Sidenius (2003).

□ **Completely homogeneous case.** To motivate the large-pool approximation, we begin with a very simple case: a completely homogeneous pool. Although somewhat artificial, it does serve as a vehicle for explaining the basic idea. It also provides a benchmark for comparison with other methods, both for accuracy and performance.

For a completely homogeneous pool, conditional on a credit-driver scenario, the number, ν, of defaults in the interval $(t_0, t_i]$, is binomially distributed with parameters K and $\hat{\pi}^{(1)}$. With all recovery-adjusted notionals being equal, say to $N^{(1)}$, the pool losses are given by $\mathcal{L} = \nu N^{(1)}$. The expected tranche losses can then be written directly as:

$$\begin{aligned}\mathbb{E}[L] &= \mathbb{E}\left[\psi\left(N^{(1)}\nu\right)\right]\\ &= \sum_k \psi\left(N^{(1)}k\right)Bin\left(k;K,\hat{\pi}^{(1)}\right)\\ &= \sum_{\ell/N^{(1)}<k<u/N^{(1)}}\left(N^{(1)}k-\ell\right)Bin\left(k;K,\hat{\pi}^{(1)}\right)\\ &\quad + S\left[1-\sum_{k<u/N^{(1)}}Bin\left(k;K,\hat{\pi}^{(1)}\right)\right]\end{aligned} \tag{11}$$

where:

$$Bin\left(k;K,\hat{\pi}^{(1)}\right) = \binom{K}{k}\left[\hat{\pi}^{(1)}\right]^k\left[\bar{\pi}^{(1)}\right]^{K-k}$$

The important observation, for our approach, is that for a large pool (that is, large K), we can approximate the binomial probabilities, $Bin(k; K, \hat{\pi}^{(1)})$, with Poisson probabilities:

$$\mathbb{P}(\nu = k) \doteq e^{-\hat{\lambda}}\hat{\lambda}^k/k! \tag{12}$$

where $\hat{\lambda} = K\hat{\pi}^{(1)}$. Thus, we can write down an expression similar to (11), replacing the binomial probabilities with the Poisson ones. We will give examples indicating the accuracy of the Poisson approximation, for values of K even as low as 10. It is this Poisson approximation, for large pools, that can be generalised to the heterogeneous case.

□ **Homogeneous case.** Before proceeding to the general case, in order to make use of the Poisson approximation, we return to the derivation (9), inserting an intermediate conditioning on ν to bring it explicitly into the picture:

$$\begin{aligned}\mathbb{E}[L] &= \mathbb{E}\left[\psi(\mathcal{L})\right]\\ &= \sum_m \mathbb{E}\left[\psi(\mathcal{L})\middle|\nu=m\right]\mathbb{P}(\nu=m)\\ &= \sum_m \mathbb{P}(\nu=m)\sum_N \psi(N)\mathbb{P}(\mathcal{L}=N|\nu=m)\end{aligned} \tag{13}$$

where the symbol N is used as a possible recovery-adjusted notional amount (individual or aggregate), in this and the next subsection.

The nature of the function ψ is simple enough. The work in making (13) more concrete consists of finding approximations for $\mathbb{P}(\nu = m)$ and $\mathbb{P}(\mathcal{L} = N | \nu = m)$. We can do that fairly quickly, without introducing any new constructs, for an intermediate case: a homogeneous pool (the recovery-adjusted notionals are common but default probabilities may vary from name to name).

Denoting:

$$\hat{\lambda} = \sum_{k=1}^{K}\hat{\pi}^{(k)}$$

we have the following generalisation of the previous Poisson approximation.

■ **Proposition.** For a large pool (K large) with $N^{(k)} \equiv N^{(1)}$ for all k: for fixed $i = 1, 2, \ldots, n$, the pool's cumulative losses, $\mathcal{L}$, up to time, t_i, are given by $\mathcal{L} = N^{(1)}\nu$ where ν, the number of defaults in the pool up to time t_i, is approximately conditionally Poisson distributed under $\mathbb{P}$ (that is, conditional on a credit-driver scenario):

$$\nu \overset{\mathcal{D}}{\approx} Pois\left(\hat{\lambda}\right)$$

The conditional expected cumulative losses absorbed by a tranche, up to time t_i, is:

$$\mathbb{E}[L] \doteq S\left(1-e^{-\hat{\lambda}}\right) - e^{-\hat{\lambda}}\left\{S\sum_{1\le m\le \ell/N^{(1)}}\frac{\hat{\lambda}^m}{m!} + \sum_{\ell/N^{(1)}<m<u/N^{(1)}}\frac{\hat{\lambda}^m}{m!}\left[u-mN^{(1)}\right]\right\} \tag{14}$$

□ **Heterogeneous case.** To extend this result to the general heterogeneous case, we need the following additional notation:

$$N_* = \min_k N^{(k)}, \quad N^* = \max_k N^{(k)}$$

$$f(N) = \sum_{k:N^{(k)}=N}\hat{\pi}^{(k)}/\hat{\lambda}, \qquad N_* \le N \le N^*$$

In general, the function $f(N)$ is a probability mass function with respect to N. (In the case where $\hat{\pi}^{(k)}$ does not depend on k, f is simply the relative frequency of the notional values and does not depend on x. Also, $\hat{\lambda} = K\hat{\pi}^{(1)}$, in this case.) The relevance of f to our pricing problem is that it approximates (assuming that all $\bar{\pi}^{(k)}$ are very close to one) the conditional probability that the pool loss is of size N, given that there has been only one default:

$$\mathbb{P}\left(\mathcal{L}=N\middle|\nu=1\right) \approx f(N)$$

Given that there has been exactly one default, the pool loss amounts to a single notional, but as we do not know who defaulted, in the heterogeneous case there is still some randomness left; f captures (approximately) that randomness.

More generally, it can be shown that:

$$\mathbb{P}\left(\mathcal{L}=N\middle|\nu=m\right) \approx f^{*m}(N)$$

under a similar assumption on $\bar{\pi}^{(k)}$, where f^{*m} denotes the m-fold convolution of f with itself, as a probability mass function. As such, given that

A. Risk-neutral cumulative default probabilities ($\hat{\pi}_i^{(k)}$) of the instruments

Rating	1 yr	2 yr	3 yr	4 yr	5 yr
Baa2	0.0007	0.0030	0.0068	0.0119	0.0182
Baa3	0.0044	0.0102	0.0175	0.0266	0.0372

B. CDO tranche structure (% of total notional)

Tranche	Attachment	Detachment
Super-senior	12.1%	100%
Senior	6.1%	12.1%
Mezzanine	4%	6.1%
Mezzanine junior	3%	4%
Equity	0%	3%

C. CDO premiums: Poisson approximation versus binomial analytics

K	Tranche	Baa2 premiums (bp)		Baa1 premiums (bp)	
		Poisson	Binomial	Poisson	Binomial
200	Senior	0	0	0	0
	Mezzanine	0	0	0	0
	Mezzanine jr	7	6	0	0
	Equity	978	978	246	246
100	Senior	0	0	0	0
	Mezzanine jr	43	41	0	0
	Mezzanine	4	3	0	0
	Equity	958	958	246	246
50	Senior	1	1	0	0
	Mezzanine	29	27	0	0
	Mezzanine jr	118	115	5	5
	Equity	898	898	244	244
25	Super-senior	0	0	0	0
	Senior	9	8	0	0
	Mezzanine	114	112	11	11
	Mezzanine jr	144	141	15	11
	Equity	787	790	231	232
10	Super-senior	1	1	0	0
	Senior	71	70	17	17
	Mezzanine	342	344	100	100
	Mezzanine jr	342	344	100	100
	Equity	342	344	100	100

there have been exactly m defaults, the pool loss amounts to a sum of m notional amounts, but we do not know who defaulted. This notional randomness is captured (approximately) by f^{*m}.

As we have seen, in the limiting case of a large pool with relatively small default probabilities, the limiting distribution of ν is Pois($\hat{\lambda}$). From (13), we then expect, in the general heterogeneous, large-pool case, to obtain a result of the form:

$$\mathbb{E}[L] \approx \exp\left(-\hat{\lambda}\right) \sum_m \frac{\hat{\lambda}^m}{m!} \sum_N \psi(N) f^{*m}(N)$$

Indeed, we have the following rigorous result, a proof of which can be found in Iscoe & Kreinin (2004).

■ **Theorem.** For a large pool (K large), conditional on a credit-driver scenario, we have the approximate equality in distribution:

$$\mathcal{L} \overset{\mathcal{D}}{\approx} \sum_{m=1}^{\nu} \mathcal{N}^{(m)} \qquad (15)$$

where $(\mathcal{N}^{(m)})_{m=1}^{K}$ is a sequence of independent identically distributed random variables, with common probability mass function f, and independent of ν, the number of defaults in the pool up to time t_i, which is approximately conditionally Poisson distributed:

$$\nu \overset{\mathcal{D}}{\approx} Pois\left(\hat{\lambda}\right)$$

More precisely:

$$\max_N \left| \mathbb{P}(\mathcal{L} = N) - \mathbb{P}\left(\sum_{m=1}^{\nu} \mathcal{N}^{(m)} = N\right) \right| = \mathcal{O}\left(\sum_{k=1}^{K} \left(\hat{\pi}^{(k)}\right)^2\right) \qquad (16)$$

Therefore, the expected cumulative losses absorbed by a tranche up to time t_i is:

$$E[L] \doteq S\left(1 - e^{-\hat{\lambda}}\right) - e^{-\hat{\lambda}} \left\{ S \sum_{1 \le m \le \ell / N_*} \frac{\hat{\lambda}^m}{m!} \sum_{mN_* \le N \le l} f^{*m}(N) + \sum_{1 \le m < u / N_*} \frac{\hat{\lambda}^m}{m!} \sum_{l < N < u} [u - N] f^{*m}(N) \right\} \qquad (17)$$

The form of the right-hand side of the representation at (15) is called a compound Poisson random variable. Thus, the accumulated pool-loss distribution, at a fixed time, is approximated by a compound Poisson distribution. (In Iscoe & Kreinin, 2004, a stronger result is established: over time, the pool-losses, $\mathcal{L}$, are approximately a compound Poisson process.)

Accuracy and performance

Here, we report on three tests: the accuracy of the Poisson approximation in the completely homogeneous setting; a comparison of the results of Monte Carlo estimation with that obtained using the analytical solutions, in the heterogeneous setting; and some performance experiments, pitting our analytic method against an industry standard method. In the first two tests, the results of the computation of the premiums are rounded to the nearest integer number of basis points. Unless specific mention is made otherwise, our numerical experiments share the following common features.

The time horizon $T = 5$ years and the premium dates are at one, two, three, four and five years from today ($t_0 = 0$). The interest rates are $r_1 = 4.6\%$, $r_2 = 5\%$, $r_3 = 5.6\%$, $r_4 = 5.8\%$ and $r_5 = 6\%$ (continuously compounded actual/365). The corresponding discount factors are $d_1 = 0.955$, $d_2 = 0.905$, $d_3 = 0.845$, $d_4 = 0.792$ and $d_5 = 0.741$.

The recovery rate of the instruments in the pool is 30%. The risk-neutral cumulative default probabilities, corresponding to two credit ratings, will be taken from table A. (These probabilities were bootstrapped from the credit spreads of CDSs.) The CDO structure is described in table B.

□ **Accuracy of the Poisson approximation: completely homogeneous case.** We compare the results of premium computations using the closed-form solution (11) and its Poisson approximation using (12), for pools of various sizes. All credit-driver weights, $\beta^{(k)}$, are taken to be zero, so that defaults among the names in the pool occur independently. Therefore, there is no (implicit) x dependence in (11) and (12), and there is no need for any numerical integration.

In table C, the premiums for the completely homogeneous pool consisting of K names having common initial credit ratings Baa2 or Baa1 are displayed. All notionals are equal to 100.

The premium for the super-senior tranche is 0bp in the table. The latter is explained by the fact that all the results of the premium computation and simulation are rounded to the nearest integer basis-point value.

The key result is that, in the completely homogeneous case, the exact solution presented in the 'Binomial' column of the table and the approximate one, presented in the 'Poisson' column, only differ by about 2–3bp, or less. This result holds even for pools as small as 10 names. Statistical estimation of the confidence intervals for the premiums shows that the number of Monte Carlo scenarios required to obtain a confidence interval of this size is equal to 6×10^7 when the number, K, of names in the pool equals 100. In this case, the simulation time is more than 2,800 times greater than the computation of the solution based on the Poisson approximation.

Accuracy of the Poisson approximation: heterogeneous case. We compare the simulation results of a heterogeneous pool consisting of names having non-zero correlation. Due to the heterogeneous nature of the pool, the binomial solution is no longer applicable. In this case, we use standard Monte Carlo valuation as a benchmark for assessing the reliability of the Poisson approximation. The pool consists of 40 names. The detailed pool description is given in table D, in which the parameter k_g denotes the number of names having the same parameters.

To assess the accuracy of the Monte Carlo method, 100 independent runs (with different seeds) of each experiment are made. Each Monte Carlo run consists of 100,000 random samples. The mean of the 100 estimates of each premium is calculated, as well as the 95% non-parametric confidence interval.

Table E presents the approximate premiums and the simulated mean premiums and confidence intervals. The approximate premiums are calculated using the compound Poisson approximation. All approximate premiums belong to the 95% confidence intervals. The results of this numerical experiment show that even in the case of a more realistic pool composition (that is, heterogeneous, correlated names), the Poisson approximation continues to be a robust approach for the valuation of CDO tranches.

Performance comparisons. Having shown the Poisson approximation to be reliable in the case of heterogeneous pools, we now turn our attention to the performance benefits of using such an approach for valuing single tranches. We compare the performance, in terms of the number of tranches being evaluated, of three methods: our compound Poisson approximation (CPA); an industry-standard analytic (ISA) method that relies on a determination of the complete pool-loss distribution; and the Monte Carlo method (MC). Twenty-eight numerical experiments were carried out. They had the following characteristics in common: 100 names in the pool, 100 single-step market scenarios, and, for Monte Carlo pricing, 100,000 credit-driver scenarios. The distinguishing features among the experiments were: the number of homogeneous groups in the pool, which varied from one (completely homogeneous) to seven; the degree of dispersion between those groups; and the number of tranches being valued (that is, one, the senior tranche with a thickness of 6% of total recovery-adjusted notional, or four, with a combined thickness of 12.125% of total recovery-adjusted notional). The groups were (roughly) the same size for each experiment. The precise correspondence between the experiments and the degree of dispersion is summarised in table F, which displays the notionals used in each experiment.

The calculations were done with a single 1.5 GHz Pentium M processor and 512MB RAM. The results of the experiments are shown in table G. It compares the relative performance of CPA and ISA, in terms of the ratio of computation times: CPA over ISA. (Thus, for an entry r, if $r < 1$, then CPA is faster than ISA, by a factor of $1/r$; if $r > 1$, then ISA is faster than CPA, by a factor of r.) The time for Monte Carlo pricing was the same for all 28 experiments and, compared with the most efficient analytical methods, was about 115 times slower than the fastest case (single tranche, homogeneous pool, using CPA), and about 13 times slower than the slowest case (four tranches, four or five groups, high dispersion, using CPA).

The conclusion from these experiments is unequivocal. For a single tranche, the compound Poisson approximation is the most efficient method in the presence of low notional dispersion and in most of the high-dispersion cases. The relative performance advantage for the low-dispersion cases is around 2.5, except in the homogeneous case (1.8). In the presence of high dispersion, when CPA has the advantage, it ranges from about one to 1.5.

For the valuation of several tranches: the two methods are about equivalent in the low-dispersion setting, except in the homogeneous case where the industry standard method has a 50% performance advantage over the compound Poisson approximation. In the high-dispersion setting, the industry standard method is the most efficient; the performance advantage ranges between about 1.5 to three times that of the compound Poisson approximation.

It should be noted that table G is based on the generic result in (17).

D. Pool of correlated instruments

Notional	Initial CR	$\beta^{(k)}$	k_g
10	Baa2	0.5	2
10	Baa3	0.5	1
10	Baa2	0.6	2
10	Baa3	0.6	2
10	Baa3	0.7	2
10	Baa3	0.8	1
20	Baa3	0.5	3
20	Baa2	0.6	4
20	Baa3	0.6	3
30	Baa2	0.5	7
30	Baa3	0.5	3
70	Baa1	0.4	4
70	Baa2	0.4	3
70	Baa3	0.5	3

E. Tranche premiums (bp): Poisson approximation versus Monte Carlo

Tranche	Poisson	Monte Carlo	95% CI
Super-senior	3	3	[3, 4]
Senior	110	109	[106, 112]
Mezzanine	285	286	[279, 294]
Mezzanine jr	392	394	[385, 406]
Equity	736	742	[733, 756]

F. Recovery-adjusted notional selection for experiments

No. of groups	Low dispersion	High dispersion
1	70	na
2	60, 70	10, 70
3	50, 60, 70	10, 60, 70
4	40, 50, 60, 70	10, 20, 60, 70
5	30, 40, 50, 60, 70	10, 20, 30, 60, 70
6	20, 30, 40, 50, 60, 70	10, 20, 30, 50, 60, 70
7	na	10, 20, 30, 40, 50, 60, 70

G. Performance ratios: CPA time over ISA time

No. of homo-geneous groups	One tranche (6%)		Four tranches (12.125%)	
	Low dispersion of notionals	High dispersion of notionals	Low dispersion of notionals	High dispersion of notionals
1	0.54	na	1.57	na
2	0.40	1.22	1.05	3.10
3	0.43	1.11	1.01	2.70
4	0.36	0.93	0.88	2.23
5	0.43	0.80	0.99	1.91
6	0.48	0.70	1.08	1.63
7	na	0.65	na	1.49

Since the first row in table G represents a homogeneous pool, we can replace (17) with (14) when assessing relative performance. Doing so improves the performance dramatically. In this case, the figures 0.54 and 1.57 drop to 0.37 and 0.4 respectively.

Comparison with other analytical methods

The most basic analytical approach calculates the conditional pool-loss distribution as the convolution of all the reference names' individual loss distributions, the latter being very simple – zero or recovery-adjusted notional, with complementary probabilities. It was proposed by Laurent & Gregory (2003) that the convolutions be calculated using the fast Fourier transform (FFT). Numerical integration of the conditional pool-loss distribution yields the unconditional one. Finally, the expected tranche loss is calculated from the latter distribution via the transformation (2): $\mathbb{E}[L] = \Sigma_l\ \psi(l)\mathbb{P}(\mathcal{L} = l)$.

Alternatively, Hull & White (2003) propose an analytical approach to the valuation of CDOs with the homogeneous restriction of common recovery-adjusted notionals in the pool (although possibly distinct default probabilities). The calculations reduce to a recursion relation to determine the conditional distribution of the number of defaults in the pool. Numerical integration of the conditional pool-loss distribution yields the unconditional one. In the heterogeneous case, using a bucketing and convolution technique, the results from homogeneous subgroups can be combined to yield the full pool-loss distribution. Finally, the expected tranche loss is obtained as in the previous approach.

Merino & Nyfeler (2002) describe a semi-analytic method in the broader setting of a very large credit portfolio with a very general default correlation model. Recovery rates are all zero. They first determine the conditional loss distribution on homogeneous sub-portfolios (buckets) via a Poisson approximation, in the same spirit as our rigorous approximation in the proposition. The results on buckets are combined via convolution, using the FFT. The unconditional distribution is then calculated via quasi-Monte Carlo integration. (Actually, they propose an overlap of the last two steps: the forward discrete Fourier transforms are multiplied together, then quasi-Monte Carlo integrated before the inverse discrete Fourier transform is taken.)

As stated previously, the common theme among these analytical approaches is the determination of the complete pool-loss distribution prior to handling the tranches. We proposed an alternative approach that focuses attention on only that part of the distribution that is pertinent to a particular tranche (see figure 1).

The approach taken in Andersen, Basu & Sidenius (2003) has some overlap with ours, in that the entire pool-loss distribution is not derived for pricing a single tranche (see (17) in Andersen, Basu & Sidenius). A bucketing and convolution method is used, one name at a time, to calculate the pool loss distribution in the required range (see (10) in Andersen, Basu & Sidenius). As stated by those authors, the computational complexity of their method is generally quadratic in the number of names in the pool. This will be the case with any method that focuses on the pool-loss distribution by a direct (convolution) method. In contrast, both the method in Hull & White (2003) and our method focus on the distribution of the number of defaults in the pool. The computational complexity of this approach is only linear in the number of names in the pool, and therefore our method is faster, especially for pricing a single tranche.

Conclusion

CDO valuation continues to be a hot topic among practitioners and academics. Monte Carlo valuation methods, despite being general, are not practical for risk management applications due to their performance constraints. Many of today's analytical methods for valuing CDO tranches are significantly faster than traditional Monte Carlo methods, but ignore the fact that investors often only hold single tranches in a given CDO pool. In these cases, applying a compound Poisson approach to derive tranche loss distributions directly, without the need for deriving the entire pool's loss distribution, can result in performance speed-ups of the order of two to three times above the speed of competing analytical valuation methods. ■

Ben De Prisco is vice-president, research and financial engineering, Ian Iscoe is principal mathematician and Alex Kreinin is director of quantitative research at Algorithmics in Toronto. Email: ben@algorithmics.com, iscoe@algorithmics.com, alex@algorithmics.com

REFERENCES

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

Bielecki T and M Rutkowski, 2002
Credit risk: modeling, valuation and hedging
Springer, Berlin

Gordy M and D Jones, 2003
Random tranches
Risk March, pages 78–83

Hull J and A White, 2003
Valuation of a CDO and an nth to default CDS without Monte Carlo simulation
Working paper, Joseph Rotman School of Management, University of Toronto

Iscoe I and A Kreinin, 2004
Pricing and valuation of synthetic CDOs
Technical paper, Algorithmics Incorporated

Iscoe I, A Kreinin and D Rosen, 1999
An integrated market and credit risk portfolio model
Algo Research Quarterly 2(1), pages 21–37

Lando D, 2004
Credit risk modelling: theory and applications
Princeton University Press, New Jersey

Laurent J-P and J Gregory, 2003
Basket default swaps, CDO's and factor copulas
Presentation at the Conference Quant '03, London, September, available at *www.defaultrisk.com*

Merino S and M Nyfeler, 2002
Calculating portfolio loss
Risk August, pages 82–86

Panjer H and G Willmot, 1992
Insurance risk models
Society of Actuaries, Shaumburg

Schönbucher P, 2003
Credit derivatives pricing models
John Wiley & Sons

Vasicek O, 1987
Limiting loan loss probability distribution
Technical document, KMV Corporation

Guidelines for the submission of technical articles

Risk welcomes the submission of technical articles on topics relevant to our readership. Core areas include market and credit risk measurement and management, the pricing and hedging of derivatives and/or structured securities, and the theoretical modelling and empirical observation of markets and portfolios. This list is not an exhaustive one.

The most important publication criteria are originality, exclusivity and relevance – we attempt to strike a balance between these. Given that *Risk* technical articles are shorter than those in dedicated academic journals, clarity of exposition is another yardstick for publication. Once received by the technical editor and his team, submissions are logged, and checked against the criteria above. Articles that fail to meet the criteria are rejected at this stage.

Articles are then sent to one or more anonymous referees for peer review. Our referees are drawn from the research groups, risk management departments and trading desks of major financial institutions, in addition to academia. Many have already published articles in *Risk*. Depending on the feedback from referees, the technical editor makes a decision to reject or accept the submitted article. His decision is final.

We also welcome the submission of brief communications. These are also peer-reviewed contributions to *Risk* but the process is less formal than for full-length technical articles. Typically, brief communications address an extension or implementation issue arising from a full-length article that, while satisfying our originality, exclusivity and relevance requirements, does not deserve full-length treatment.

Submissions should be sent to the technical team at *technical@incisivemedia.com*. The preferred format is MS Word, although Adobe PDFs are acceptable. The maximum recommended length for articles is 3,500 words, and for brief communications 1,000 words, with some allowance for charts and/or formulas. We expect all articles and communications to contain references to previous literature. We reserve the right to cut accepted articles to satisfy production considerations. Authors should allow four to eight weeks for the refereeing process.