Analytic improvement of the saddle-point approximation and spread risk attribution in a portfolio of tranches[+]

Damian Taras, Christopher Cloke-Browne and Evan Kalimtgis

The saddle-point approximation has rapidly become an established technique for the evaluation of portfolio loss distributions and risk measures such as Value-at-Risk. In this paper, Damian Taras, Chistopher Cloke-Brown and Evan Kalimtgis describe an improved version of the saddle-point approximation which, in contrast to the original version, retains accuracy over the whole distribution. The key applications of the approach are the evaluation of portfolio risk measures and the pricing of structured credit instruments. The next step in the risk management of a portfolio of tranches is to consider movements of the loss distributions, due to spread risk. It is shown how spread risk measures may be evaluated for such a portfolio and a consistent procedure for the attribution of spread risk to names inside tranches is described for the first time. The existence of this procedure has wide-ranging implications for the management and hedging of structured credit portfolios.

1 Introduction

In this paper we describe a new approach to the evaluation of the loss distribution of a portfolio, and discuss simulation of the movements of the loss distribution. The analytic approach has two key applications: in the calculation of Value-at-Risk (VaR) statistics of the portfolio, and in calculation of prices of structured instruments on the portfolio, such as Collateralized Debt Obligation (CDO) tranches or k-to-default baskets. To demonstrate the analytic approach, we show how the saddle-point approximation may be systematically improved by the inclusion of higher-order terms. The resulting approximation is accurate across the whole loss distribution, not just the tail. This leads to a valuable alternative to Monte Carlo simulation, whose sensitivity to noise is problematic.

Turning to movements of the loss distribution curve, we show how these may be generated to derive a spread VaR for a tranche. We describe for the first time a consistent procedure for the attribution of VaR contributions to the constituent names of the tranche. These VaR contributions are analogous to the VaR contributions of names in an unstructured portfolio, and indeed reduce to these values when the thresholds of the tranche are taken to 0% and 100% respectively. The formalism has also been extended to CDO's of CDO's ("CDO-squared" products).

We also describe how the total VaR of a portfolio containing several tranches may be attributed consistently amongst the tranches and all names inside all tranches. Correlation between tranches arises as a natural consequence of the correlation between all the underlying names. This leads to the concept of optimization of a tranche through name selection based on its risk/return characteristics. More generally, these developments have wide implications for the risk management and hedging of structured credit portfolios.

[+] Original version 11th January 2005; updated 4th August 2005

2 Construction of a portfolio loss distribution: analytic approaches

In this section we describe a new analytic approach to the loss distribution of a portfolio, based on improvements to the saddle-point approximation. The approximation appears in the context of using a moment-generating function (mgf) approach, whereby the loss distribution is expressed as an Laplace inversion integral with respect to the mgf. Correlation may be incorporated by use of a discrete latent variable.

Other closed-form solutions to the computation of a default loss distribution have appeared recently, based on recursive approaches, for example by Hull & White, 2003, and Andersen *et al.*, 2003. A drawback of the latter approaches is that, while suited to equally weighted portfolio, they become more complicated for non-equally weighted portfolios. The analytic approximation described here is in contrast as easily applicable to both cases. Of course, while a recursive approach is limited to the aggregation of underlying discrete (binary) distributions, as is required for default risk evaluation or tranche pricing, the improved saddle point approximation is also applicable to the aggregation of underlying continuous distributions, as is required for spread risk (VaR) evaluation.

2.1 Moment Generating Functions and the Saddle-Point Approximation

The saddle-point approximation is applied in the context of using a moment generating function (mgf) to generate the loss distribution of a portfolio. The mgf($M(s)$) is defined as the expectation of the exponential of total loss (X):

$$M(s) = \mathbf{E}[e^{sX}] = \int e^{st} f(t)\, dt \qquad (1)$$

where $f(t)$ is the distribution function of the loss X . In the absence of correlation, the mgf simplifies to the product of the mgf's of the individual assets:

$$M(s) = \mathbf{E}[e^{s\sum X_i}] = \prod M_i(s) \qquad (2)$$

where X_i is the loss of asset i and $M_i(s) = \mathbf{E}[e^{sX_i}]$. Correlation between the assets may in certain models be incorporated analytically, for example through the use of a latent variable. The distribution function, $f(t)$, of the total loss may then recovered from the mgf by application of an inverse Laplace transform:

$$f(t) = \frac{1}{2\pi \mathrm{i}} \int_{-\mathrm{i}\infty}^{+\mathrm{i}\infty} \exp(K(s) - st)\, ds\,. \qquad (3)$$

where we have defined, $K(s) = \log M(s)$, known as the cumulant generating function.

The inverse Laplace transform integral (3) is not in general analytically tractable. A traditional approach is to use a Fast Fourier Tranform approach, although this is well-known to lead to problems such as ill-conditioning and truncation errors (see Press *et al.*, for further details).

The saddle-point approximation (see e.g. Daniels, 1954; Martin *et al.*, 2001) represents an analytic alternative for this inversion step, with its analytic form leading to significant increases in efficiency. The approximation consists in finding the point at which the term in the exponential is stationary, and then Taylor-expanding as far as a quadratic and doing the resulting Gaussian integral. The saddle-point is $s = \hat{t}$ obeys $K'(\hat{t}) = t$, where the prime denotes differentiation with respect to *s*. We perform a Taylor expansion on $K(s) - st$, as far as the quadratic term:

$$K(s) - st = K(\hat{t}) - t\hat{t} + \tfrac{1}{2}(s-\hat{t})^2 K_Y''(\hat{t}) + \dots \tag{4}$$

We then compute the resulting Gaussian integral:

$$f(t) \approx \frac{\exp(K(\hat{t}) - t\hat{t})}{2\pi \mathrm{i}} \int_{-\mathrm{i}\infty}^{+\mathrm{i}\infty} \exp\left(\tfrac{1}{2}(s-\hat{t})^2 K''(\hat{t})\right) ds = \frac{\exp(K(\hat{t}) - t\hat{t})}{\sqrt{2\pi K''(\hat{t})}}$$

Proceeding similarly for the tail probability $P(t) = \int_t^\infty f(t')dt'$, we come to

$$P(t) \approx \theta(-\hat{t}) + \exp(K(\hat{t}) - t\hat{t})\left\{\mathrm{sgn}(\hat{t})\Phi\left(-\left|\hat{t}\right| K''(\hat{t})^{1/2}\right)\exp\left(K''(\hat{t})\hat{t}^2/2\right)\right\}, \tag{5}$$

where $\theta(t)$ is the Heaviside step function.

A strength of the saddle-point approximation is that it is generally accurate in the tail of the distribution, in fact becoming more accurate the further into the tail. This makes it well suited to a Value-at-Risk calculation which searches for the loss corresponding to a small tail probability (e.g. 5% or 1%). In contrast, more conventional approximations, such as the Edgeworth or Cornish-Fisher expansions, become more inaccurate for smaller tail probabilities.

Another feature that makes the saddle-point approximation suited to VaR calculations is that, due to its analytic form, it becomes possible to search for the loss corresponding to the tail probability in Laplace space, that is, by searching values of $\hat{t}$. As a result it is not necessary to compute a Laplace inversion for each trial value of the VaR, unlike for other numerical methods, leading to a significant increase in speed.

2.3 Higher-order corrections to the saddle-point approximation

So far we have stopped at the quadratic term in the expansion (4). Instead, we can continue up to the fourth order, to include third and fourth derivatives of the kgf. We can expand the exponential in these terms and the integration over s remains analytic. For the density, we find

$$f(t) \approx \frac{\exp(K(\hat{t}) - t\hat{t})}{\sqrt{2\pi K''(\hat{t})}} \left(1 + \frac{K^{IV}(\hat{t})}{8K''(\hat{t})^2} - \frac{5K'''(\hat{t})^2}{24K''(\hat{t})^3} + \ldots \right)$$

Similarly, for the tail probability, we come to

$$\begin{aligned} P(t) \approx \theta(-\hat{t}) + \exp(K(\hat{t}) - t\hat{t}) \Big[& \operatorname{sgn}(\hat{t}) \Phi\left(-|\hat{t}| K''(\hat{t})^{1/2} \right) \exp\left(K''(\hat{t}) \hat{t}^2 / 2 \right) \\ & \times \left(1 - \frac{\hat{t}^3 K'''(\hat{t})}{6} + \frac{\hat{t}^4 K^{IV}(\hat{t})}{24} + \frac{\hat{t}^6 K'''(\hat{t})^2}{72} \right) \\ & + \frac{1}{72\sqrt{2\pi} K''(\hat{t})^{5/2}} \Big\{ 3K''(\hat{t})(1 - K''(\hat{t})\hat{t}^2)(\hat{t} K^{IV}(\hat{t}) - 4K'''(\hat{t})) \\ & - \hat{t} K'''^2(\hat{t})(3 - \hat{t}^2 K''(\hat{t}) + \hat{t}^4 K''(\hat{t})^2) \Big\} \Big] \end{aligned} \tag{6}$$

These types of corrections to the saddle-point approximation are well-known in many branches of physics, one example being the corrections to the classical (or mean-field) properties due to quantum interference processes (see e.g. Feynman & Hibbs 1954).

The saddle-point is generally accurate within the tail of the probability distribution, although it loses accuracy towards the centre. In contrast, we have found that incorporating corrections to the saddle-point approximation provides an analytic approximation that retains its validity over the *whole* distribution.

As a simple example, consider a portfolio whose loss distribution is given by the exponential distribution, given by $f(t) = a\exp(-at)$. The mgf is $M(s) = a/(a-s)$, defined for $s < a$. Applying the saddle-point approximation and its corrections to the probability density using this mgf, we find

$$f(t) \approx c.a\exp(-at),$$

where c is a constant, with $c = e/\sqrt{2\pi}$ =0.92.. in the saddle-point approximation, and $c = 11e/(12\sqrt{2\pi})$ = 0.994…when the first higher-order corrections are included. Thus we see an error of around 8% for the saddle-point approximation is drastically reduced to only 0.6% when higher-order corrections are included.

As a more realistic example, we also examine the performance of the approximation for the default loss distribution of a more typical portfolio, of around 200 assets with an average correlation of around 20%. Fig. 1 shows the relative performance of the various approximations to the tail probability: Monte Carlo, saddle-point (Eq. (5)) and corrections to saddle-point (Eq. (6)). This figure could be viewed equivalently as the VaR for each tail probability. Correlation was modelled by use of a discrete latent variable, calibrated by discretization of an underlying Gaussian risk factor (as described in Martin *et al.*, 2001). We chose a sufficiently large number of points (250,000) for the Monte Carlo that this approximation can be viewed as a benchmark for the ‘correct’ distribution. The key observation from this figure is that

- the corrections to the saddle-point approximation lead to a dramatic improvement in accuracy of the saddle-point approximation across the *whole* distribution.

Of course, since the improved saddle-point approach is analytic, it avoids the noise problems of Monte Carlo, which makes it particularly useful for calculating sensitivities of this distribution, e.g. for VaR contributions. Notice that the 'steps' of the true distribution are smoothed within the saddle-point approximation and its improved version.

Besides VaR calculations, the above analysis is particularly significant for pricing structured credit instruments, such as CDO tranches or k-to-default baskets. For example, the fair spread of a tranche or basket is expressed through an expectation with respect to the loss distribution of its underlying pool. While a senior tranche (or high-k basket) will depend on the upper tail of the distribution, a mezzanine tranche (or low-k baskets) will depend on the centre of the distribution, where the saddle-point approximation becomes unreliable. Again, the analytic nature of the approximation becomes especially useful for calculation of sensitivities of quantities such as the fair spread.

3 VaR and VaR contributions of a CDO tranche or k-to-default basket

Having discussed how to generate a loss distribution of a portfolio, we now turn to consideration of *movements* of the loss distribution. Recall that the fair spread of a tranche derives from expectations of the loss distribution, at payment times up to maturity. The loss distribution at each payment time is generated from the default probabilities of the underlying names, while the default probabilities are stripped out from the current underlying spreads. To calculate a spread VaR for a tranche, we need the distribution of the tranche spread loss, which corresponds to the distribution of the tranche fair spread[1] at a risk horizon. Movements of the tranche fair spread correspond to movements of the portfolio loss distribution: we therefore need to describe movements over time of the entire portfolio loss distribution curve up to the risk horizon. This problem is conceptually similar to models of the interest rate curve, where possible movements of the entire curve are simulated into the future.

The steps required for a spread VaR calculation of a tranche are summarized in Fig.2. The simplest way to visualize the calculation is through a Monte Carlo approach, where we build up the distribution of the spread loss of the tranche by repeatedly sampling its fair spread at the risk horizon. Working backwards from the right in Fig. 2: for each sample of the tranche fair spread, we need a sample realisation of the loss distribution curve. For each sample realisation of this curve, we need a sample default probability, and corresponding spread, for each underlying instrument. Thus the VaR calculation is hence driven by repeatedly sampling correlated configurations of the underlying spreads. Within the context of a structural model linking credit to equity,

[1] The loss of the tranche is given as the change in its fair spread times its duration, times its notional: this represents a linear approximation in the fair spread.

movements of the underlying spreads may in turn be generated from movements of the underlying equity values.

Although the calculation of the VaR of a tranche is conceptually not difficult, there are two further nontrivial steps that are of vital interest:

1. Given a value for the spread VaR of a tranche, can we attribute this VaR to the constituent names of the tranche? In other words, can we generate consistent *spread VaR contributions* for names inside the tranche?
2. Given 1., can we generalise the approach to a portfolio containing more than one tranche, together with unstructured assets?

The answer to both these questions is affirmative. We address questions 1. and 2. in sections 3.1 and 3.2 respectively.

3.1 VaR contributions of names inside a standalone tranche

For an unstructured portfolio, the existence of a consistent definition of VaR contributions is ensured by the property of 1-homogeneity (see e.g. Martin *et al.*, 2001). This property states that, if all the weights of underlying names (a_i) are scaled by a constant factor, then the statistic in question (here, VaR) scales by the same factor. As a consequence, the VaR (T_P) may be expressed as an exact sum of VaR contributions:

$$T_P = \sum_i a_i \frac{\partial T_p}{\partial a_i}.$$

.

To find the VaR contributions of the tranche, we would like to apply a similar observation. To do so, however, we need to scale the thresholds of the tranche in the same proportion as we scale the notionals of the underlying assets: this ensures that the spread of the tranche is invariant under scaling. Let us define $T = s\sum a_i$, where s is the tranche fair spread: the loss of the tranche at the risk horizon is then given[2] by the change in T . Furthermore T satisfies the required homogeneity property and hence may be decomposed into contributions $T = \sum T_i$, where

$$T_i = a_i \frac{d^{mov}}{da_i}\left(s\sum_j a_j \right). \tag{7}$$

Here d^{mov} / da_i represents a 'moving' derivative, whereby the thresholds move with changes in the weight a_i, in such a way that each threshold represents a fixed fraction of the total (pool) notional. Thus, if the thresholds are given by losses L_1 and L_2, they are a function of the set of weights $\{a_j\}$:

$$L_{1,2}\left(\{a_j\}\right) = \ell_{1,2} \sum_j a_j$$

[2] Up to constant factors (the tranche duration and the ratio of the tranche notional to the pool notional)

where $\ell_{1,2}$ are fixed. T is then a function of the weights $\{a_j\}$ and the thresholds $L_{1,2}$, so that

$$\frac{d^{mov}}{da_i} T(\{a_j\}, L_{1,2}\{a_j\}) \equiv \left(\frac{\partial}{\partial a_i} + \frac{\partial L_1}{\partial a_i}\frac{\partial}{\partial L_1} + \frac{\partial L_2}{\partial a_i}\frac{\partial}{\partial L_2} \right) T$$

Hence we see the appearance of extra partial derivatives of the quantity T, with respect to the thresholds L_1 and L_2, in the definition of the VaR contributions T_i.
Some strengths of this definition of tranche VaR contributions are that

- In the limit of the lower and upper thresholds of the tranche being taken to 0% and 100% respectively, the tranche VaR contributions converge to the VaR contributions given by the unstructured portfolio, and
- the VaR contributions of the names inside the tranche sum *exactly* to the total VaR of the tranche, as a direct consequence of their definition and without the need for a normalization factor.

Fig. 2 shows typical results of the VaR contributions of the names inside a single (standalone) tranche, and how these contributions change with the subordination of the tranche. The illustration in fact provides a valuable check of the formalism with intuition, in particular with regard to the correlation of the name with the portfolio, as follows:

- Names with a high correlation (such as Delphi here) produce VaR contributions that *increase* as the tranche becomes more senior. Conversely, names with low correlation (such as Cable and Wireless) produce VaR contributions that *decrease* with seniority.

This observation coincides with intuition since it is reasonable that high-correlation, or systemic, names (such as Delphi) are most risky to the holder of a senior tranche: if a single name defaults, others are more likely to follow if that name is systemic. Similarly, the holder of an equity tranche is most sensitive to a small total number of defaults, which is more likely to be generated by idiosyncratic names (such as Cable and Wireless).

3.2 VaR attribution amongst names inside several tranches plus unstructured names

The above reasoning may be applied also to the spread VaR of a portfolio containing more than one tranche, as well as unstructured names. This total VaR may be attributed amongst the portfolio instruments, that is, the tranches and the unstructured names. Furthermore, the VaR contribution of each tranche may be further attributed as VaR sub-contributions of each name inside each tranche. This concept is illustrated in Fig. 4. Notice that Ford appears as both a tranche constituent and as an unstructured name in the portfolio.

Again, it is relatively easy to generate the total VaR of the portfolio: correlated movements in all the underlying spreads, whether inside or outside tranches, are simulated, so as to generate movements of the loss distributions of the tranches, and hence movements of tranche spreads. Correlation between tranches arises as a natural consequence of the correlations between the underlying names.

However attribution of the total spread VaR to tranches and names inside tranches is more difficult. We have found that it is possible to arrive at a consistent attribution procedure using the decomposition of the tranche loss as by Eq. (7) – whether it appears as an unstructured instrument, or inside a tranche, or both (as does the example Ford in Fig. 4). Details of this attribution procedure will be described elsewhere.

A direct benefit of this technology is the ability to optimize a tranche in terms of its risk/reward characteristics. Consider for example an equally weighted synthetic tranche. The reward of each name may be represented by its spread, while its risk is its contribution the tranche spread VaR. This leads to the concept of a "tranche Sharpe ratio", defined as spread of the name divided by its VaR contribution. Removal of names with a low tranche Sharpe ratio leads to a tranche customised on a risk-reward basis.

Moreover:

- Since the spread VaR contribution depends on the subordination, the optimisation procedure is tailored to the tranche subordination. A name with a low Sharpe ratio at a AAA subordination may have a high Sharpe ratio at a BBB subordination.
- The tranche may also be optimised with respect to a reference portfolio: if we include the tranche inside a reference portfolio, the VaR contribution of a name inside the tranche will be altered accordingly.

We remark also that we have extended this technology to CDO's of CDO's, or "CDO-squared" products. The analytics is correspondingly more complicated although the reasoning is conceptually similar. Another remark is that other risk measures, such as expected shortfall, may also be evaluated within the same approximations described in this paper, and also attributed amongst names within tranches in an analogous way.

4 Conclusions

In this paper we have described a new analytic approximation to the loss distribution of a portfolio, based on the systematic improvement of the saddle-point approximation by the inclusion of higher-order terms. We have demonstrated that the resulting approximation is accurate across the whole loss distribution, not just the tail, even in the presence of correlation. This approximation provides a valuable alternative to Monte Carlo simulation, for both the calculation of loss statistics such as VaR for the portfolio, and in the pricing of structured instruments, such as CDO tranches or k-to-default baskets.

We have also discussed how the spread VaR of a tranche may be generated from simulation of the movements of the loss distribution, and how this VaR may be attributed consistently to the constituent names of the tranche. We have also generalised the approach to CDO's of CDO's ("CDO-squared" products), and portfolios containing several tranches as well as unstructured names.

These developments have wide implications for the practices of risk management of structured credit portfolios. The systematic and fully consistent risk monitoring of a structured portfolio enables the identification of its leading risks, a valuable feature for example for lightly-managed ("flexi") tranches. It also opens the way to effective hedging, whether by single-name hedging or by shorting further structured products (such as baskets). We have also introduced the concept of optimization of a tranche through name selection based on its risk/return characteristics, potentially with reference to a base portfolio.

Damian Taras, Christopher Cloke-Browne and Evan Kalmitgis are members of the Portfolio Management group in Capital Markets at Dresdner Kleinwort Wasserstein. Damian Taras is a Vice-President managing the quantitative analysis group. Christopher Cloke-Browne is a Director running the Portfolio Advisory team. Evan Kalimtgis is a Managing Director and Head of Correlation and Capital Arbitrage proprietary trading.

References

L. Andersen, J. Sidenius and S. Basu
"All Your Hedges in One Basket"
Risk, November 2003.

H. E. Daniels
"Saddle-point approximations in statistics"
Ann. Math. Stats. **25** 631-650 (1954)

J. Hull and A. White
"Valuation of a CDO and an nth. to default CDS Without Monte-Carlo Simulation"
www.defaultrisk.com/pp_crdrv_14.htm

R. P. Feynman and A. R. Hibbs
Quantum Mechanics and Path Integrals
New York: McGraw-Hill, 1965

R. J. Martin, K. E. Thompson and C. J. Browne
"Taking to the Saddle"; "VaR: Who Contributes and How Much?";"Default Correlation Modelling"
RISK (2001)

W. H. Press, S. A.Teukolsky, W. T. Vetterling and B. P. Flannery
Numerical Recipes in C : The art of scientific computing
CUP, 2nd ed, 1992

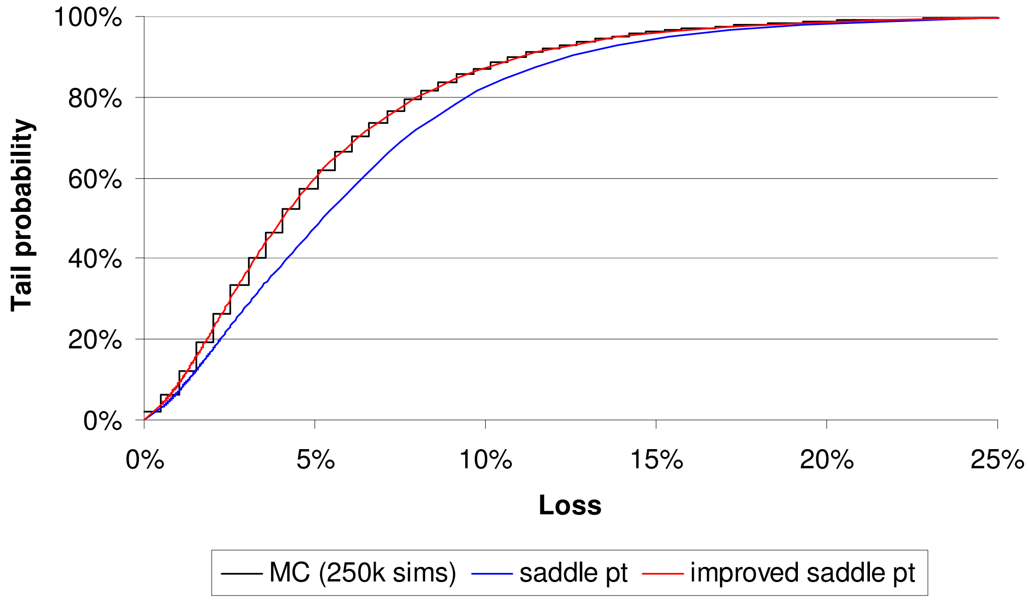

Fig. 1 A comparison of the performance of the saddle-point approximation to the tail probability (Eq. (5)) and its improved version (Eq. (6)) against Monte Carlo (with 250,000 points). A portfolio of around 200 names was used, with average correlation of around 20%, and realistic default probabilities stripped from 5-year market credit spreads.

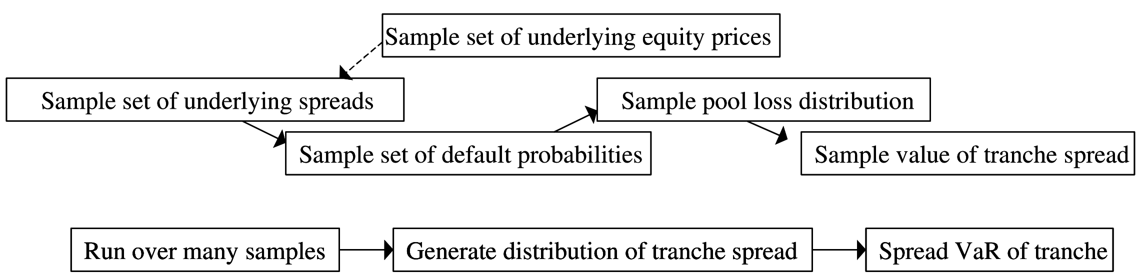

Fig. 2 Required steps to generate a distribution of the fair spread of a tranche, and hence perform a spread VaR calculation for a tranche.

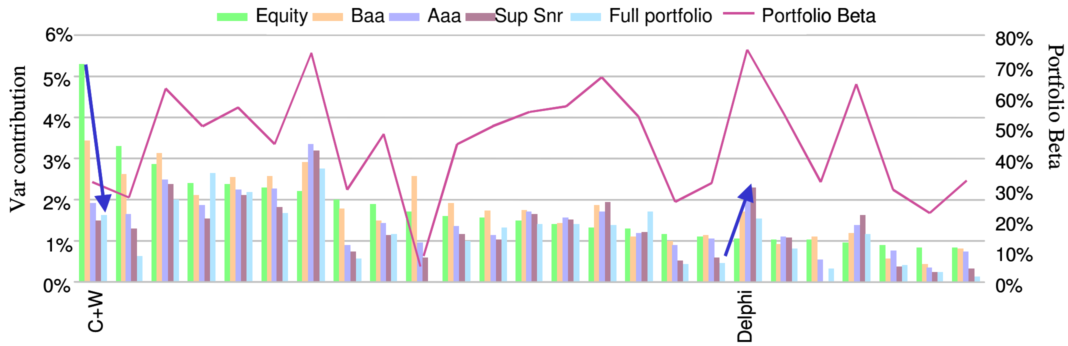

Fig.3: Attribution of the spread VaR of a standalone tranche to its constituent names, and how this depends on the subordination of the tranche.

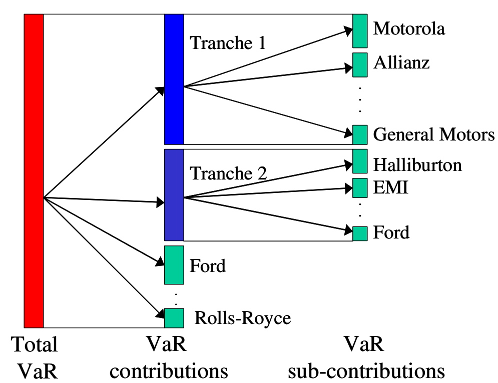

Fig.4: How the total spread VaR of a portfolio may be attributed as to its contents, which may contain several tranches as well as unstructured names. The VaR contribution from each tranche may be further attributed to constituent names inside each tranche.