DELFT UNIVERSITY OF TECHNOLOGY

REPORT 07-06

Computation of VaR and VaR Contribution in the Vasicek portfolio credit loss model: a comparative study

X. Huang, C.W. Oosterlee, and M.A.M Mesters

ISSN 1389-6520

Reports of the Department of Applied Mathematical Analysis

Delft 2007

Computation of VaR and VaR Contribution in the Vasicek portfolio credit loss model: a comparative study

Xinzheng Huang[a,b,*], Cornelis. W. Oosterlee[a,c] and Mâcé Mesters[b]

[a] *Delft Institute of Applied Mathematics, Delft University of Technology, Mekelweg 4, 2628CD, Delft, the Netherlands*

[b] *Group Risk Management, Rabobank, Croeselaan 18, 3521CB, Utrecht, the Netherlands*

[c] *CWI - National Research Institute for Mathematics and Computer Science, Kruislaan 413, 1098 SJ, Amsterdam, the Netherlands*

Abstract

We compare various numerical methods for the estimation of the VaR and the marginal VaR Contribution (VaRC) in the Vasicek one-factor portfolio credit loss model. The methods we investigate are the normal approximation, the saddlepoint approximation, a simplified saddlepoint approximation and importance sampling. We investigate each method in terms of speed, accuracy and robustness and in particular explore their abilities of dealing with exposure concentration.

Key words: portfolio credit risk, Value at Risk, VaR contribution, normal approximation, saddlepoint approximation, importance sampling

1 Introduction

Credit risk is the risk of loss resulting from an obligor's inability to meet its obligations. More generally it can also include losses due to credit quality changes. For financial institutions it is essential to quantify the credit risk at a portfolio level. The key issue in the portfolio credit loss modeling is the specification of the default dependence among obligors. A common practice is utilizing a factor model, such that the obligors are independent conditional on some common factors, e.g., state of the economy, different industries and geographical regions.

We quantify portfolio credit risk in the Vasicek model, which is the basis of the Basel II (Basel Committee on Bank Supervision, 2005) internal rating based (IRB) approach. It is a Gaussian one factor model such that the default events are driven by a latent common factor that is assumed to follow the Gaussian distribution. It is a one-period model that only considers default risk, i.e., loss only occurs when an obligor defaults in a fixed time horizon. The model is able to reproduce the qualitative behavior of empirical credit loss distributions, namely fat tail and high skewness. Under certain homogeneity conditions, the Vasicek one-factor model leads to very simple analytic asymptotic approximations of the loss distribution, Value at Risk (VaR) and VaR Contribution (VaRC). This asymptotic approximation works extremely well if the portfolio consists of a large

*Corresponding author; E-mail: X.Huang@ewi.tudelft.nl

number of small exposures. The model may be extended to portfolios that are not homogeneous in terms of default probability and pairwise correlation. However, the analytic approximation of the Vasicek model can significantly underestimate risks in the presence of exposure concentrations, i.e., when the portfolio is dominated by a few obligors.

A variety of alternative methods to estimate the portfolio credit risk and the risk contribution have been proposed for more general portfolios. Glasserman & Ruiz-Mata (2006) provide an interesting comparison of methods for computing credit loss distributions. The methods considered there are plain Monte Carlo simulation, a recursive method due to Andersen et al. (2003), the saddlepoint approximation, and the numerical transform inversion as in Abate et al. (2000). They conclude that the plain Monte Carlo method is the best method in a multi-factor setting in terms of speed and accuracy, followed by the saddlepoint approximation. They find that the recursive method performs well when the number of obligors is small but becomes slow as the number of obligors increases, particularly for high loss levels. This is because the recursive method computes the entire loss distribution and when the number of obligors increases, the maximum total loss increases in the mean time. They also find that numerical transform inversion methods gives acceptable estimates for small loss levels but the approximation worsens for higher loss levels. This is not surprising. This method numerically inverts the Bromwich integral, whose integrand becomes highly oscillatory and extremely difficult to handle for high loss levels.

The perspective of our comparison in this paper is quite different from Glasserman & Ruiz-Mata (2006). First we restrict ourselves to the one-factor model. Secondly we are mainly interested in VaR_α when α is close to 1, i.e., high loss levels. Thirdly we are also interested in the estimation of marginal VaR Contribution. Finally we would like to investigate how well the problem of exposure concentration can be handled.

We point out that the conclusions of Glasserman & Ruiz-Mata (2006) are based on portfolios with less than 1000 obligors. But in practice it will not be surprising that a bank's credit portfolio has more than tens of thousands of obligors. The plain Monte Carlo simulation will certainly become more demanding in computation time as the portfolio size increases. After all, a true problem with plain simulation is the estimation of the marginal VaR Contribution, which is based on the scenarios that portfolio loss equals VaR. These are extremely rare events. We should for this reason consider importance sampling as in Glasserman & Li (2005); Glasserman (2006) instead of plain simulation. We will drop the recursive method and the numerical transform inversion method for obvious reasons given above. Note that Debuysscher & Szegö (2003) suggest that the numerical inversion can be expedited by fast Fourier transform. However a straightforward implementation of FFT also suffers from the same problem as the numerical transform inversion. We should instead include the normal approximation method as in Martin (2004), which is a direct application of the central limit theorem. In addition we consider a Simplified saddlepoint approximation for the estimation of VaRC.

The rest of the article is organized as follows. In section 2 we introduce the Vasicek one-factor model. Section 3 reviews various alternative methods we want to investigate, among others are the normal approximation, the saddlepoint approximation, the simplified saddlepoint approximation and importance sampling. A stylized portfolio is considered in section 4. Section 5 discusses the robustness of each method. The 6th section concludes.

2 The Vasicek Model

Consider a credit portfolio consisting of n obligors. Any obligor i can be characterized by three parameters: the exposure at default EAD_i, the loss given default LGD_i and the probability of default PD_i. Obligor i is subject to default after a fixed time horizon and the default can be modeled as a Bernoulli random variable D_i such that

$$D_i = \begin{cases} 1 & \text{with probability } \mathrm{PD}_i, \\ 0 & \text{with probability } 1\text{-}\mathrm{PD}_i. \end{cases}$$

Define the effective exposure of obligor i by $w_i = \mathrm{EAD}_i \times \mathrm{LGD}_i$, then the loss incurred due to the default of obligor i is given by

$$L_i = \mathrm{EAD}_i \times \mathrm{LGD}_i \times D_i = w_i D_i.$$

It follows that the portfolio loss is given by

$$L = \sum_{i=1}^{n} L_i = \sum_{i=1}^{n} w_i D_i.$$

Value at Risk (VaR) is the risk measure chosen in the Basel II Accord for the evaluation of capital requirements. Let α be some given confidence level, the VaR is simply the α-quantile of the loss distribution of L. Thus,

$$\mathrm{VaR}_\alpha = \inf\{x : \mathbb{P}(L \leq x) \geq \alpha\}. \tag{1}$$

The *VaR Contribution* (VaRC) measures how much each obligor contributes to the total VaR of a portfolio. Under some continuity conditions, the VaRC coincides with the conditional expectation of L_i given that the portfolio loss L is equal to $\mathrm{VaR}_\alpha(L)$, i.e.,

$$\mathrm{VaRC}_{i,\alpha} = w_i \frac{\partial \mathrm{VaR}_\alpha}{\partial w_i}(L) = w_i E[D_i | L = \mathrm{VaR}_\alpha(L)], \tag{2}$$

For more details see Tasche (2000) and Gourieroux et al. (2000).

The Vasicek model is named after a series of Vasicek's papers (1987; 1991; 2002). The modeling of the dependence structure among counterparties in the portfolio is simplified by the introduction of a common factor that affects all counterparties. It is assumed that the standardized asset log-return X_i of obligor i can be decomposed into a systematic part Y and an idiosyncratic part Z_i such that

$$X_i = \sqrt{\rho_i} Y + \sqrt{1-\rho_i} Z_i, \tag{3}$$

where Y and all Z_i are independent standard normal random variables. In case $\rho_i = \rho$ for all i the parameter ρ is called the common asset correlation.

One can derive that the VaR and the VaRC at the $\alpha-$percentile for an infinitely large portfolio without exposure concentration are as follows:

$$\mathrm{VaR}_\alpha = \sum_i w_i \Phi\Big(\frac{\Phi^{-1}(PD_i) + \sqrt{\rho_i}\Phi^{-1}(\alpha)}{\sqrt{1-\rho_i}}\Big), \tag{4}$$

$$\mathrm{VaRC}_{i,\alpha} = w_i \Phi\Big(\frac{\Phi^{-1}(PD_i) + \sqrt{\rho_i}\Phi^{-1}(\alpha)}{\sqrt{1-\rho_i}}\Big), \tag{5}$$

where Φ denotes the CDF of the standard normal distribution.

The asymptotic Vasicek approximations (4) and (5) work well for portfolios consisting of an infinite number of small obligors. These formulas are less suitable for and tend to underestimate risks for portfolios with few obligors or portfolios dominated by a few large exposures. In the next section we compare alternative numerical methods to deal with the problem of exposure concentration.

3 Numerical Methods

Here we give a brief introduction to the numerical methods we want to compare for the estimation of VaR and VaRC, among which the saddlepoint approximation is a method proposed by ourselves (Huang et al., 2006).

3.1 Normal Approximation

The normal approximation (NA) is a direct application of the central limit theorem (CLT) and can be found in Martin (2004). When the portfolio is not sufficiently large for the law of large numbers to hold or not very homogeneous, unsystematic risk arises. We then need to take into account the variability of portfolio loss L conditional on the common factor Y. This can easily be approximated due to the CLT. Conditional on the common factor Y, the portfolio loss L is normally distributed with mean $\mu(Y)$ and variance $\sigma^2(Y)$ such that

$$\mu(Y) = \sum_{i=1}^{n} w_i p_i(Y), \tag{6}$$

$$\sigma^2(Y) = \sum_{i=1}^{n} w_i^2 p_i(Y)(1 - p_i(Y)), \tag{7}$$

where $p_i(Y) = P(D_i = 1|Y)$. It follows that the conditional tail probability reads

$$P(L > x|Y) = \Phi\left(\frac{\mu(Y) - x}{\sigma(Y)}\right).$$

The unconditional tail probability[1] can then be obtained by integrating over Y, i.e.,

$$P(L > x) = E\left[\Phi\left(\frac{\mu(Y) - x}{\sigma(Y)}\right)\right] = \int \Phi\left(\frac{\mu(y) - x}{\sigma(y)}\right)\phi(y)dy. \tag{8}$$

To obtain the VaR contribution in the current setting, we first differentiate $P(L > x)$ with respect to the effective exposure:

$$\frac{\partial}{\partial w_i} P(L > x) = E_Y\left[\left(\frac{1}{\sigma}\frac{\partial \mu}{\partial w_i} - \frac{1}{\sigma}\frac{\partial x}{\partial w_i} - \frac{\mu - x}{\sigma^2}\frac{\partial \sigma}{\partial w_i}\right)\phi\left(\frac{\mu - x}{\sigma}\right)\right], \tag{9}$$

with

$$\frac{\partial \mu(Y)}{\partial w_i} = p_i(Y), \tag{10}$$

$$\frac{\partial \sigma(Y)}{\partial w_i} = w_i p_i(Y)(1 - p_i(Y))/\sigma(Y), \tag{11}$$

and ϕ the pdf of the standard normal distribution. Now replace x by VaR_α in formula (9). Since the tail probability $P(L > \mathrm{VaR}_\alpha)$ is fixed at $1 - \alpha$, the left-hand side should vanish and the VaR contribution is given by

$$w_i \frac{\partial \mathrm{VaR}_\alpha}{\partial w_i} = w_i \frac{E_Y\left[\left(\frac{1}{\sigma}\frac{\partial \mu}{\partial w_i} - \frac{\mu - \mathrm{VaR}_\alpha}{\sigma^2}\frac{\partial \sigma}{\partial w_i}\right)\phi\left(\frac{\mu - \mathrm{VaR}_\alpha}{\sigma}\right)\right]}{E_Y\left[\frac{1}{\sigma}\phi\left(\frac{\mu - \mathrm{VaR}_\alpha}{\sigma}\right)\right]}. \tag{12}$$

[1] There are also attempts to find an analytic approximation to

$$E[P(L(Y) > x) - P(\mu(Y) > x)],$$

which is known as the granularity adjustment. For more details, see Gordy (2003), Wilde (2001).

The normal approximation is also applied in Shelton (2004) for CDO/CDO-squared pricing. Zheng (2006) employs higher order approximations as an improvement to the central limit theorem to compute CDS/CDO-squared transactions. In this article we will restrict ourselves to standard normal approximation as in Martin (2004).

3.2 Saddlepoint Approximation

It is well known that the saddlepoint approximation (SA) provides accurate estimates to very small tail probabilities. This makes it a very suitable technique in the context of portfolio credit loss. The saddlepoint approximation to a random variable of finite sum $X = \sum_{i=1}^{n} X_i$ relies on the existence of the moment generating function (MGF) $M_X(t) = E(e^{tX})$. For X_i with known analytic MGF's M_{X_i}, the MGF of the sum X is the product of MGF of X_i, i.e.,

$$M_X(t) = \prod_{i=1}^{n} M_{X_i}(t).$$

Let $K_X(t) = \log M_X(t)$ be the Cumulant Generating Function(CGF) of X. The inverse MGF of X, known as the Bromwich integral, can then be written as

$$f_X(x) = \frac{1}{2\pi j} \int_{-j\infty}^{+j\infty} \exp(K_X(t) - tx) dt, \tag{13}$$

with $j = \sqrt{-1}$.

The saddlepoint, i.e., the point at which $K_X(t) - tx$ is stationary, is a $t = \tilde{t}$ such that

$$K'_X(\tilde{t}) = x. \tag{14}$$

The density $f_X(x)$ and the tail probability $\mathbb{P}(X > x)$ can be approximated by $K_X(t)$ and its derivatives at $\tilde{t}$.

There are several variants of saddlepoint approximations available and we take the Daniels (Daniels, 1987) formula for the density

$$f_X(x) = \frac{\phi(z_l)}{\sqrt{K''(\tilde{t})}} \left\{ 1 + \left[-\frac{5K'''(\tilde{t})^2}{24K''(\tilde{t})^3} + \frac{K^{(4)}(\tilde{t})}{8K''(\tilde{t})^2} \right] \right\}, \tag{15}$$

and the Lugannani-Rice (Lugannani & Rice, 1980) formula for the tail probability

$$\mathbb{P}(X > x) = 1 - \Phi(z_l) + \phi(z_l) \left(\frac{1}{z_w} - \frac{1}{z_l} \right), \tag{16}$$

where $z_w = \tilde{t}\sqrt{K''(\tilde{t})}$ and $z_l = sgn(\tilde{t})\sqrt{2[x\tilde{t} - K(\tilde{t})]}$.

The saddlepoint approximation is usually highly accurate in the tail of a distribution. The use of saddlepoint approximation in portfolio credit loss is pioneered by a series of articles by Martin et al. (2001*a*; *b*). Gordy (2002) showed that saddlepoint approximation is fast and robust when applied to CreditRisk$^+$. All of them apply saddlepoint approximation to the *unconditional* MGF of loss L, despite the fact that the L_i are not independent. Annaert et al. (2006) show that the procedure described in Gordy (2002) may give inaccurate results in case of portfolios with high skewness and kurtosis in exposure size.

Huang et al. (2006) apply the saddlepoint approximation to the *conditional* MGF of L given the common factor Y, so that the L_i are independent. This is the situation for which the saddlepoint approximation will work well. In this way a uniform accuracy of density and tail probability for different levels of portfolio loss L is achieved at the expense of some extra computational cost: Eq. (14) needs to be solved once for each realization of Y. It is also shown by a numerical example that the accuracy of the saddlepoint approximation is not impaired by a high skewness and/or kurtosis in the exposure size.

Martin & Ordovás (2006) compare the application of saddlepoint approximation to the *unconditional* MGF and to the *conditional* MGF. They confirm that the latter (called indirect approach therein) is more accurate and more generally applicable. The use of the saddlepoint approximation is also recommended by Yang et al. (2006) and Antonov et al. (2005) in the context of CDO pricing, both adopting the indirect approach.

As for the computation of VaRC, Huang et al. (2006) give the following formula,

$$w_i \frac{\partial \mathrm{VaR}_\alpha}{\partial w_i} = w_i \frac{E_Y\left[\int_{-j\infty}^{+j\infty} \frac{p_i(Y)e^{w_i t}}{1-p_i(Y)+p_i(Y)e^{w_i t}} \exp\left(K(t,Y) - t\mathrm{VaR}_\alpha\right) dt\right]}{E_Y\left[\int_{-j\infty}^{+j\infty} \exp(K(t,Y) - t\mathrm{VaR}_\alpha) dt\right]}, \tag{17}$$

and propose a double saddlepoint approximation for both integrals in the numerator and denominator. This requires finding for each obligor i and each Y a saddlepoint $\tilde{t}_i$ in addition to Eq. (14) that solves

$$\sum_{k \neq i} \frac{w_k p_k(Y) e^{w_k t}}{1 - p_k(Y) + p_k(Y)e^{w_k t}} = \mathrm{VaR}_\alpha - w_i. \tag{18}$$

3.3 Simplified Saddlepoint Approximation

For the calculation of VaRC, Martin et al. (2001*b*) propose the following estimate, also under the name of a saddlepoint approximation,

$$\mathrm{VaRC}_{i,\alpha} = \frac{w_i}{\tilde{t}} \frac{\partial K_L(t)}{\partial w_i}\Big|_{t=\tilde{t}} = \frac{w_i p_i e^{w_i \tilde{t}}}{1 - p_i + p_i e^{w_i \tilde{t}}} \tag{19}$$

in the case of independent obligors. Here $K_L(t) = \log E\left[e^{tL}\right]$ is the cumulant generating function of L, $\tilde{t}$ is the solution of $K_L'(t) = \mathrm{VaR}_\alpha$ and p_i is the default probability of obligor i. This estimate is also derived by Thompson & Ordovás (2003) based on the idea of an ensemble and Glasserman (2006) as a result of an asymptotic approximation.

It is straightforward to extend the independent case to the conditionally independent case as in the Vasicek model, which reads

$$\mathrm{VaRC}_{i,\alpha} \approx \frac{E_Y\left[f_L(\mathrm{VaR}_\alpha|Y) \frac{w_i p_i(Y) e^{w_i \tilde{t}}}{1 - p_i(Y) + p_i(Y) e^{w_i \tilde{t}}}\right]}{E_Y\left[f_L(\mathrm{VaR}_\alpha|Y)\right]}, \tag{20}$$

where $f_L(\mathrm{VaR}_\alpha|Y)$ can be computed efficiently by the saddlepoint approximations. This formula can also be found in Antonov et al. (2005).

We call the estimate given by (20) a simplified saddlepoint approximation (SSA), in the sense that it is a simplified version of the double saddlepoint approximation to (17). For a portfolio with n distinct obligors, the double saddlepoint approximation requires solving one time (14) and n times (18) for each realization of the common factor Y, whereas the SSA only needs the solution $\tilde{t}$ to (14). It then assumes that $\tilde{t}$ and $\tilde{t}_i$, the solutions to (14) and (18), are more or less the same for each obligor i and simply replace all $\tilde{t}_i$ by the saddlepoint $\tilde{t}$. Consequently the SSA is generally faster than the SA, but it may give less accurate results if the above assumption is violated.

3.4 Importance Sampling

Monte Carlo (MC) simulation is an all-around method which is very easy to implement. However, Monte Carlo simulation can be extremely time-consuming. The typical error convergence rate of plain Monte Carlo simulation is $O(1/\sqrt{N})$, where N is the number of simulations, requiring a large number of simulations to obtain precise results. See Boyle et al. (1997) for a review in the finance context.

Two main variance reduction techniques for Monte Carlo methods applied to portfolio credit loss can be found in the literature. Control variates are employed by Tchistiakov et al. (2004) where the Vasicek distribution is considered as a control variable. Importance sampling (IS) is adopted by Kalkbrener et al. (2004); Merino & Nyfeler (2005) for the calculation of Expected Shortfall Contribution and by Glasserman & Li (2005); Glasserman (2006) for the calculation of VaR and VaRC. We note that the difficulty with Monte Carlo simulation mainly concerns the determination of VaRC since the estimate expressed in formula (2) is based on the very rare event that portfolio loss $L = \text{VaR}$. In this respect control variates do not provide any improvement. IS as suggested in Glasserman & Li (2005); Glasserman (2006) seems a more appropriate choice and will be adopted here.

The importance sampling procedure consists of two steps:

- Mean shifting - shift of the mean of common factors,
- Exponential twisting - change of distribution to the (conditional) default probabilities.

With mean shifting the common factor Y is sampled under probability measure $\mathbb{R}$ which is equivalent to the original measure $\mathbb{P}$ such that under $\mathbb{R}$, Y is normally distributed with mean mean $\mu \neq 0$ and variance 1. The tail probability is then given by

$$P(L > x) = E^{\mathbb{R}}\left[1_{\{L>x\}} e^{-\mu Y + \mu^2/2}\right]. \tag{21}$$

This step will increase the probability $L > x$, making a rare event less rare.

The idea of exponential twisting is to choose

$$q_{i,\theta(Y)}(Y) = \frac{p_i(Y) e^{\theta(Y) w_i}}{1 + p_i(Y)(e^{\theta(Y) w_i} - 1)}, \tag{22}$$

which increases the default probability if $\theta > 0$. This step will cluster the losses around x, which is particularly useful for the estimation of VaRC. With these two techniques the tail probability can be formulated as

$$\begin{aligned} P(L > x) &= E\left\{E^{\mathbb{Q}}\left[\mathbf{1}_{\{L>x\}} \prod_i \left(\frac{p_i(Y)}{q_i(Y)}\right)^{D_i} \left(\frac{1-p_i(Y)}{1-q_i(Y)}\right)^{1-D_i} \Big| Y\right]\right\} \\ &= E\left\{E^{\mathbb{Q}}\left[\mathbf{1}_{\{L>x\}} e^{-\theta(Y)L + K(\theta(Y),Y)} \Big| Y\right]\right\} \\ &= E^{\mathbb{R}}\left\{e^{-\mu Y + \mu^2/2} E^{\mathbb{Q}}\left[\mathbf{1}_{\{L>x\}} e^{-\theta(Y)L + K(\theta(Y),Y)} \Big| Y\right]\right\}. \end{aligned} \tag{23}$$

To find suitable parameters for the procedures of exponential twisting and mean shifting, Glasserman & Li (2005) and Glasserman (2006) propose to choose $\hat{\theta}(y)$ that solves

$$K'(\theta(y), y) = x. \tag{24}$$

and μ as the solution to

$$\max_y K(\hat{\theta}(y), y) - \hat{\theta}(y)x - \frac{1}{2} y^\top y,$$

where $\hat{\theta}(y)$ is given by (24) and the exponential of which to be maximized is the upper bound of $P(L > x|Y = y)\exp(-y^\top y)$. Note that Eq. (24) is identical to Eq. (14) in the saddlepoint approximation since both methods employ the idea of an Esscher transform.

The estimation of the VaR contribution is trivial. It is given by

$$\mathrm{VaRC}_i = w_i \frac{\sum_k D_i l^k \mathbf{1}_{\{L^k=\mathrm{VaR}\}}}{\sum_k l^k \mathbf{1}_{\{L^k=\mathrm{VaR}\}}},$$

where the superscript k denotes the k-th simulated scenario and l is the likelihood ratio $e^{-\mu Y+\mu^2/2-\theta(Y)L+K(\theta(Y),Y)}$.

4 A Stylized Portfolio

We consider a stylized portfolio A consisting of 11,325 obligors which only differ in exposure size. They are categorized in 6 buckets and the exposure per obligor and the number of obligors in each bucket are the following,

bucket	1	2	3	4	5	6
Exposure	1	10	50	100	500	800
# of obligors	10000	1000	200	100	20	5

Other parameters are

$$\rho = 20\%, \quad PD = 0.33\%. \tag{25}$$

The portfolio has a total exposure of 54000. It is a portfolio of so-called lower granularity since the largest obligor has an exposure 800 times larger than the smallest obligor. Exposure concentration is not really significant as the weight of the largest obligor is less than 1.5%.

Both the normal approximation and saddlepoint approximation calculate the tail probability instead of the VaR directly. The VaR can then be obtained by inverting the loss distribution. A not very sophisticated iterative solver, the bisection method, is used for this purpose. We search the VaR in the interval with as a lower bound the portfolio expected loss $E(L)$ and as an upper bound the total portfolio exposure. The two approaches also require the discretization of the common factor Y. In a one-factor setting, numerical integration methods rather than simulation can be used for efficient and accurate calculation of unconditional loss density and tail probability. We employ the Gaussian quadrature method and truncate the domain of Y to the interval $[-5, 5]$. The probability of Y falling out of this interval is merely 5.7×10^{-7}. The speed of saddlepoint methods strongly depends on the number of abscissas N in the discretization of Y. Most of the CPU time is spent in finding the saddlepoints. This also applies to IS with exponential twisting. We find generally that $N = 100$ abscissas is sufficient in terms of accuracy for the saddlepoint methods, while for IS many more points are necessary to obtain an estimate with small variance. For the normal approximation we also adopt $N = 100$.

In the tables that follow "Vasicek" denotes the asymptotic approximation of the Vasicek model and "NA" denotes the normal approximation. The results given by the saddlepoint approximations are labeled by "SA". "IS-10K" stands for importance sampling with ten thousand scenarios. Its VaR estimate and the sample standard deviations are computed by dividing the ten thousand scenarios into 10 equally-sized subsamples.

Table 1 presents both the $\mathrm{VaR}_{99.9\%}$ and $\mathrm{VaR}_{99.99\%}$ of the portfolio given by various methods. CPU times are in seconds. Monte Carlo simulation based on 10 subsamples with $16M$ scenarios each serves as our benchmark. We also report on the standard deviation and the 95% confidence intervals (CI) beside the point estimates. The standard deviations of $\text{VaR}_{99.9\%}$ and $\text{VaR}_{99.99\%}$ are 7.7 (0.1% of the corresponding VaR) and 38.4 (0.5% of the corresponding VaR), respectively.

Even though the portfolio has no serious exposure concentration, the VaR estimates at both confidence levels obtained from the asymptotic Vasicek approximation are far from the benchmark VaR (relative errors around 5%). The normal approximation provides a significant improvement in accuracy with only little additional computational time. The relative errors for both VaR estimates are less than 1%. The saddlepoint approximation is even more accurate while remaining fast. Both VaR estimates, which can be obtained in several seconds, fall within the 95% confidence interval and have relative errors less than 0.2%. The variance reduction of IS compared to plain simulation is more effective in the far tail. With only one thousand scenarios in each subsample, the standard deviations of the VaR estimate are not really small. Although the VaR estimates given by IS are comparable to those given by SA, the former is significantly more computational intensive.

	$\text{VaR}_{99.9\%}$	$\text{VaR}_{99.99\%}$	time
Benchmark	3960.3(7.7)	6851.6(38.4)	
95% CI	[3945.2, 3975.3]	[6776.3, 6926.9]	
Vasicek	3680.5	6477.0	8E-4
NA	3924	6804	2E-2
SA	3965	6841	6E+0
IS-10K	3975.3(56.4)	6836.8(84.9)	2E+3

Table 1: $\text{VaR}_{99.99\%}$ of portfolio A. The Benchmark and IS-10K sample standard deviations (in parentheses) are calculated using 10 simulated subsamples of 16M and 1K scenarios each, respectively.

Regarding the VaR Contribution, we in fact compute the VaRC of an obligor scaled by its effective weight w_i, i.e.,

$$\frac{\partial \text{VaR}_\alpha}{\partial w_i}(L) = P(L_i = 1 | L = \text{VaR}).$$

This represents the VaRC of an obligor as a percentage of its own effective exposure. Expressed as a probability, it always lies in the interval $[0, 1]$.

VaRCs of the obligors in each bucket at loss level $L = 4000$ and $L = 6800$ are given in Table 2. The simulated portfolio loss L is so sparse in the vicinity of the VaR, that we have to replace the event $\{L = \text{VaR}\}$ by

$$\frac{|L - \text{VaR}|}{\text{VaR}} < \gamma \tag{26}$$

to make our VaRC estimates meaningful[2]. We face a dilemma here. A small γ reduces bias but at the expense of having only very few useful scenarios. We here choose $\gamma = 0.5\%$ for $L = 4000$ and $\gamma = 1\%$ for $L = 6800$. The former event has a probability around 0.004% and the latter around 0.001%. Our benchmark VaRC estimates are both based on 12000 such events, resulting from roughly 300M and 1200M scenarios respectively. The benchmark standard deviations (in parenthesis) and confidence intervals are computed by dividing the 12000 scenarios into ten equally-sized subsamples. For IS we simply use the same γ as MC for both loss levels. There are 316 and 772 out of 10K IS scenarios, hence 3.16% and 7.72% respectively, for which L falls in the desired ranges. The effect of clustering losses around the level of interest by IS is truly significant compared to plain Monte Carlo simulation.

It appears that SA is the only method that is able to give all VaRC estimates within the 95% confidence interval. Its maximum absolute error of 0.33% is also the smallest among all methods.

[2]As an alternative Mausser & Rosen (2004) suggest the use of Harrell-Davis estimate: an L-estimator that computes a quantile estimate as a weighted average of multiple order statistics.

The estimates from SSA are similar to those with SA, especially for small exposures. In terms of speed SSA is about seven times faster than SA. At the same time it has 2 estimates outside the 95% confidence interval. The normal approximation and importance sampling have 7 and 5 estimates outside the 95% confidence interval, respectively. The differences to the benchmark for all the three methods are quite small though, with maximums 1.14% (SSA), 1.27% (NA) and 1.24% (IS). NA overestimates the VaRC of small exposures and underestimates the VaRC of large exposures. On the contrary, SSA overestimates the VaRC of large exposures. A problem with IS is that the VaRCs are not monotonically increasing with the effective weight w, which is counterintuitive. From this perspective 10K scenarios do not seem enough.

5 Analysis of Robustness

Both the normal approximation and the saddlepoint approximation are asymptotic approximations that become more accurate when the portfolio size increases. The normal approximation stems from the central limit theorem and uses merely the first two moments of the conditional portfolio loss $L(Y)$. Higher order approximations such as the Edgeworth expansion provide an improvement as they take the higher cumulants of $L(Y)$ into account. As for the saddlepoint approximations, the Daniels formula can be considered as a generalization of the Edgeworth expansion that makes use of explicit knowledge of the moment generating function (see Jensen, 1995). In this respect it is expected that the saddlepoint approximations are generally more accurate than the normal approximation, which is confirmed by our example above. A drawback is that the tail probability given by the Edgeworth expansion is not necessary in the range of $[0, 1]$ and is not always monotone. Similarly the quantile approximations are not always monotone in the probability levels (cf. Wallace, 1958). The Lugannani-Rice formula may also suffer from the same problems. On the contrary, importance sampling/simulation always gives estimates to a probability in $[0, 1]$.

An important concern is whether the conditions of the central limit theorem hold if severe exposure concentration is present in a portfolio. Apparently if the conditions do not hold the normal approximation will fail. Let us now consider a portfolio B consisting of a bucket of 1000 obligors with effective exposure $w_1 = 1$ and one large obligor with effective exposure $w_2 = S$, $S \in \{20, 100\}$, i.e.,

bucket	1	2
Exposure	1	S, $S \in \{20, 100\}$
# of obligors	1000	1

For other parameters ρ and PD we adopt (25). The weight of the large obligor relative to the total exposure is almost 2% when $S = 20$ and 10% when $S = 100$. The latter should be considered as serious exposure concentration. The Binomial Expansion Method (cf. Huang et al., 2006), by which the VaR and VaRC can be computed almost exactly, will be used as the benchmark.

We consider the quantile $\alpha = 99.99\%$. Table 3 gives the VaR of portfolio B obtained by various methods. The approximation error of VaR is measured by the relative error (RE) defined as

$$\text{RE}=(\text{Estimate-Benchmark})/\text{Benchmark}.$$

When $S = 20$, we see that all methods, except Vasicek, have relative errors of less than 2%. When S is increased to 100, both Vasicek and NA become erratic (relative errors $> 10\%$). On the contrary the effect of a large S on the accuracy of SA is marginal. We remark that we have tested for even larger S up to 1000 (50% of the total exposure of correponding portfolio), and SA manages to consistenly give $\text{VaR}_{99.99\%}$ estimates with $|\text{RE}| < 2\%$. IS is also insusceptible to the size of S. It is as accurate as SA, but demands much more CPU time.

(a) VaRC at the loss level $L = 4000$

	VaRC1	VaRC2	VaRC3	VaRC4	VaRC5	VaRC6	time
Benchmark	6.33%(0.04%)	6.38%(0.05%)	6.54%(0.03%)	6.86%(0.08%)	9.36%(0.17%)	11.32%(0.38%)	
95% CI	[6.25%, 6.41%]	[6.28%, 6.48%]	[6.49%, 6.59%]	[6.70%, 7.02%]	[9.02%, 9.70%]	[10.58%, 12.06%]	
Vasicek	7.41%	7.41%	7.41%	7.41%	7.41%	7.41%	3E-3
NA	6.55%	6.59%	6.78%	7.02%	8.92%	10.35%	1E-2
SA	6.35%	6.39%	6.58%	6.82%	9.21%	11.65%	2E+0
SSA	6.35%	6.37%	6.5%	6.68%	9.12%	12.46%	3E-1
IS-10K	6.54%	6.46%	6.77%	6.7%	9.4%	10.33%	1E+3

(b) VaRC at the loss level $L = 6800$

	VaRC1	VaRC2	VaRC3	VaRC4	VaRC5	VaRC6	time
Benchmark	11.23%(0.09%)	11.29%(0.09%)	11.56%(0.11%)	11.87%(0.12%)	14.89%(0.21%)	17.86%(0.59%)	
95% CI	[11.06%, 11.41%]	[11.11%, 11.48%]	[11.35%, 11.77%]	[11.63%, 12.11%]	[14.48%, 15.30%]	[16.70%, 19.03%]	
Vasicek	12.59%	12.59%	12.59%	12.59%	12.59%	12.59%	3E-3
NA	11.42%	11.48%	11.74%	12.06%	14.65%	16.59%	1E-2
SA	11.23%	11.29%	11.55%	11.88%	14.94%	17.78%	2E+0
SSA	11.23%	11.27%	11.48%	11.75%	14.89%	18.44%	3E-1
IS-10K	11.34%	11.52%	11.62%	12.03%	14.83%	16.62%	1E+3

Table 2: VaRC of portfolio A at the loss levels $L = 4000$ and $L = 6800$. The benchmark sample standard deviations (in parentheses) are calculated using 10 simulated subsamples of 1,200 relevant scenarios each.

	$S = 20$				$S = 100$			
	VaR	std	error	time	VaR	std	error	time
Benchmark	125				170			
Vasicek	122.3		-2.13%	6E-4	131.9		-22.39%	1E-3
NA	125		0.00%	1E-2	149		-12.35%	9E-3
SA	126		0.80%	3E+0	168		-1.18%	3E+0
IS-10K	124.1	1.7	-0.72%	2E+2	170.5	3.1	0.29%	2E+2

Table 3: $\text{VaR}_{99.99\%}$ of portfolio B. Errors reported are relative errors.

The reason why the normal approximation does not work for $S = 100$ is not difficult to explain. Conditional on the common factor Y, NA tries to approximate the loss density by a normal distribution (due to the central limit theorem). This works quite well when S is as large as 20. However, when we have $S = 100$, which is almost 10% of the total exposure, the loss density will no longer be unimodal. A normal approximation is not able to capture this pattern and therefore can be problematic. This is illustrated in Figure 1.

It is also worthwhile explaining how the exposure concentration is handled by the saddlepoint approximation. Therefore, instead of computing only a quantile of the portfolio loss, we calculate the whole loss distribution when $S = 100$ using our benchmark and the SA. This is demonstrated in Figure 2(a). We notice that the true loss distribution is not smooth in the vicinity of 100, which is precisely the size of the large exposure S.

Recall that the SA relies on the formulation of the Bromwich integral (13) representing a probability density function. It is thus implicitly assumed that the portfolio loss L, which is discrete when LGD is constant, can be closely approximated by a continuous random variable which has an absolutely continuous cumulative distribution function. The saddlepoint method thus produces a smoothed version of the loss distribution. A more detailed discussion of the saddlepoint approximations as smoothers is in Davison & Wang (2002). We see in Figure 2(a) that the saddlepoint approximation to the tail probabilities is incorrect for almost all quantiles preceding the point of non-smoothness (around the 99.6%-quantile) but is again accurate for higher quantiles. It entails that, with one or a few exceptional exposures in the portfolio, a uniform accuracy of the loss distribution may not be achieved by a straightforward saddlepoint approximation. This can be a problem if the quantile we are interested in precedes the non-smoothness in the loss distribution, which usually occurs at the size of large exposures.

A very easy algorithm can be used to retain the uniform accuracy. Suppose a portfolio has m large exposures $S_i, i = 1, \ldots, m$ with $S_1 \leq S_2 \leq \cdots \leq S_m$. For any loss level $x < S_k$ the tail probability conditional on Y can be written as

$$P(L > x|Y) = 1 - P(L - \sum_{i \geq k} L_i \leq x|Y) \prod_{i \geq k} P(D_k = 0|Y). \tag{27}$$

The above reformulation takes into account the implicit information that when $L < x$ the obligors with exposure larger than x must not default. An application of SA to the probability $P(L - \sum_{i \geq k} L_i > x|Y)$ furthermore removes the exposure concentration $S_i, i \geq k$. It is apparently more accurate than a direct SA to $P(L > x|Y)$. A similar idea is discussed in Beran & Ocker (2005). We call this method the adaptive saddlepoint approximation (ASA) here. As an experiment we apply the ASA to portfolio B with $S = 100$ and plot in Figure 2(b) the loss distribution for loss levels up to but excluding $L = S$ (in the estimation of the tail probabilities the ASA only differs from a direct SA for loss levels $L < S$.). The loss distribution given by the ASA matches the benchmark almost exactly for all $L < S$.

Now we consider the VaR Contribution. Table 4 presents the VaRC of both a small obligor (VaRC1) and the large obligor (VaRC2). For the cases $S = 20$ and $S = 100$, we report four

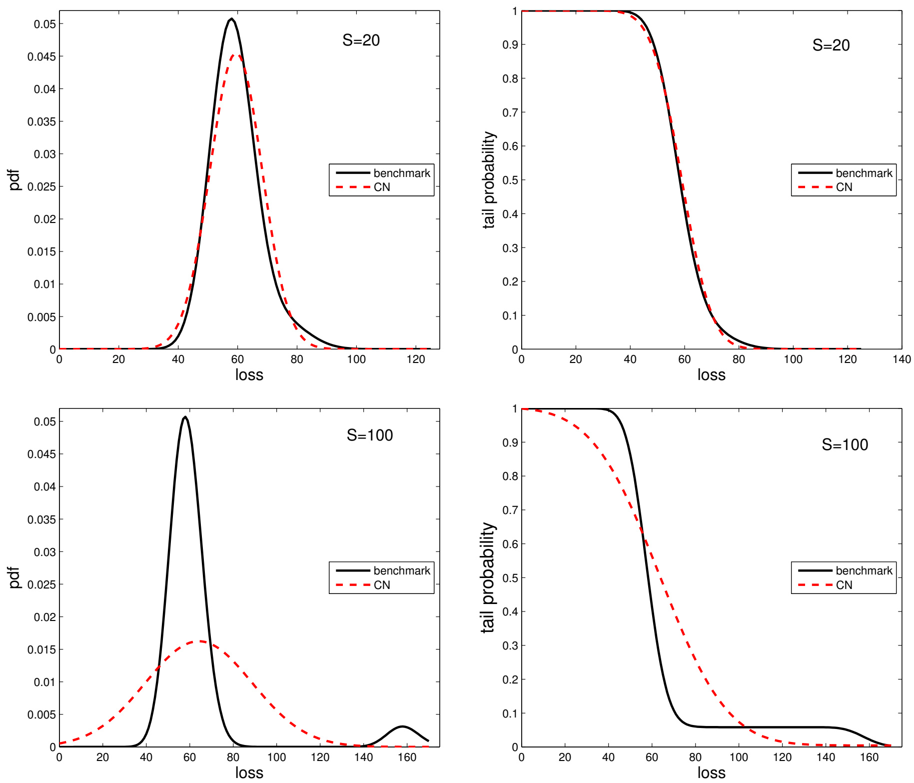

Figure 1: Loss density and tail probability of Portfolio B given by the normal approximation conditional on an arbitrarily chosen common factor Y.

estimates to the VaRC given by each method. The error we report here is absolute error. NA gives fair VaRC estimates for both VaRC1 and VaRC2 when $S = 20$ but deviates dramatically from the benchmark when $S = 100$. This is in line with its performance on the VaR estimation. SA is quite accurate for VaRC1 but becomes less accurate for VaRC2 as S increases. SSA resembles SA in the estimates of VaRC1 but does not give satisfactory estimates to VaRC2 at all: both errors are larger than 5%. This can be understood by the fact that, as mentioned in §3.3, the solutions to (14) and (18) are indeed close for small exposures but can differ substantially for large exposures. Further experiments show that NA, SA and SSA may all give VaRC values that are not in the interval $[0, 1]$ in the presence of more exceptional exposure concentrations (as is pointed out at the beginning of this section). IS appears to be the best method in terms of accuracy and robustness in this case.

In both portfolios A and B importance sampling seems to perform fine for determining VaRC. An important reason for this is that the obligors in a bucket are considered identical and we are able to take the average of all obligors in the same bucket when estimating VaRC. This makes the simulated VaRC estimates much less volatile. We must point out that even though IS is able to cluster the simulated losses around the VaR of interest and thus significantly increases the probability $P(L = \text{VaR})$, a rather large number of replications would still be necessary.

Let us consider a portfolio C consisting of 100 obligors with exposures all different from each other such that

$$w_i = i, \quad i = 1, \ldots, 100. \tag{28}$$

The parameters ρ and PD are again the same as in (25).

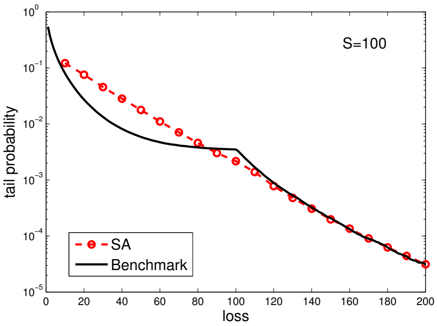

(a) SA

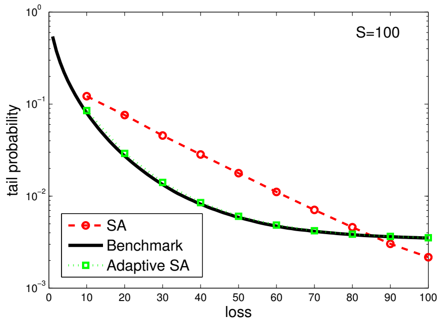

(b) ASA

Figure 2: The loss distribution of portfolio B given by the saddlepoint approximation (SA) and adaptive saddlepoint approximation (ASA) when the loss distribution is not smooth at the vicinity of S. PD= 0.0033, $\rho = 0.2$, $S = 100$.

(a) $S = 20$

$S = 20$	VaRC1	error	VaRC2	error	time
Benchmark	12.06%		21.78%		
Vasicek	12.25%	0.19%	12.25%	-9.53%	3E-3
NA	12.12%	0.06%	18.94%	-2.84%	1E-2
SA	12.05%	-0.01%	21.70%	-0.08%	8E-1
SSA	11.96%	-0.10%	27.06%	5.28%	3E-1
IS-10K	12.04%	-0.02%	22.89%	1.11%	1E+3

(b) $S = 100$

$S = 100$	VaRC1	error	VaRC2	error	time
Benchmark	8.29%		87.07%		
Vasicek	15.45%	7.16%	15.45%	-71.62%	3E-3
NA	12.68%	4.39%	43.18%	-43.89%	2E-1
SA	8.89%	0.60%	90.79%	3.72%	8E-1
SSA	9.15%	0.86%	78.52%	-8.55%	3E-1
IS-10K	8.12%	-0.17%	88.85%	1.78%	1E+3

Table 4: $\text{VaRC}_{99.99\%}$ of portfolio B. Errors reported are absolute errors.

Figure 3 are are scatterplots of the (scaled) VaRC (y-axis) at the loss level $L = 700$, which is around $\text{VaR}_{99.99\%}$, against the EAD (x-axis). In the left-side figure we show the results given by the saddlepoint approximation, the simplified saddlepoint approximation and the normal approximation. All methods clearly show that the VaRC increases as the EAD increases, which is highly desirable for practical purposes. SSA again gives results very close to the saddlepoint approximation. Compared to the SA, the NA overestimates the VaRC of small exposures and underestimates the VaRC of large exposures. This is consistent with the pattern shown in Portfolio A. The estimates given by importance sampling with 10K scenarios and 100K scenarios are presented in the right-side figure along with those given by SA. γ as in (26) is set to be 1%. The relation between the VaRC and EAD is not clear at all with only 10K simulated scenarios. The estimates, resulting from 256 relevant scenarios, disperse all over the area. Improvement in the performance of VaRC estimation is discernable when we increase the number of scenarios of IS by ten times. The VaRC estimates are then based on 2484 relevant scenarios and the upward trend of VaRC with increasing EAD is evident. However due to simulation noises the curve remains highly oscillatory and an even higher number of scenarios seems necessary.

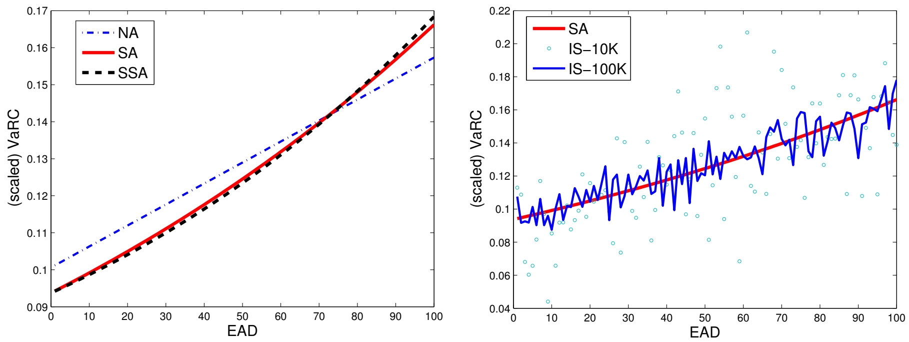

Figure 3: VaR Contribution of Portfolio C as a function to EAD at the loss level $L = 700$.

6 Conclusions

We have examined various numerical methods for the purpose of calculating the credit portfolio VaR and VaRC under the Vasicek one-factor model. Each method provides a viable solution to VaR/VaRC estimation for lower granular portfolios and portfolios with medium exposure concentration. However there is no perfect method that prevails under all circumstances and the choice of preferred method turns out to be a trade-off among speed, accuracy and robustness.

The normal approximation is the fastest method and is able to achieve a fair accuracy. It is however rather vulnerable because it is incapable of handling portfolios dominated by one or a few obligors (or, portfolios with multi-modal loss density). The simplified saddlepoint approximation is second to the normal approximation in speed and may suffer from the same problem when estimating the VaRC.

Importance sampling does not guarantee to be the most accurate method but it always works fine provided a sufficient number of scenarios are drawn. It makes no assumption on the composition of a portfolio and thus is certainly the best choice from the perspective of robustness. Unlike the other methods, it always gives estimates to the scaled VaRC in $[0, 1]$. The downside of IS is that it is rather time-consuming when compared to other methods. Moreover IS is not strong in the estimation of VaRC, which is really demanding in the number of simulated scenarios.

The saddlepoint approximation is generally more accurate than normal approximation and simplified saddlepoint approximation. It is also more reliable in the sense that it can handle more extreme exposure concentration. Consequently it may well serve as a fast alternative to IS with a good balance between accuracy and speed. It must be emphasized that, if the loss distribution is not smooth due to exceptional exposure concentrations and the target quantile precedes the non-smoothness in the loss distribution, a straightforward implementation of SA is likely to be insufficient. The adaptive saddlepoint approximation should be employed in this situation.

Finally we would like to point out again that the normal approximation and the saddlepoint approximation methods are all based on asymptotic approximations. They become more accurate when the portfolio size increases. On the other hand, importance sampling become even more demanding in computation time when the portfolio size increases.

References

Abate, J., Choudhury, G. L. & Whitt, W. (2000), An introduction to numerical transform inversion and its application to probability models, *in* W. K. Grassman, ed., 'Computational Probability', Kluwer, Norwell, MA.

Andersen, L., Sidenius, J. & Basu, S. (2003), 'All your hedges in one basket', *RISK* (November), 67–72.

Annaert, J., Garcia, J. B. C., Lamoot, J. & Lanine, G. (2006), Don't fall from the saddle: the importance of higher moments of credit loss distributions, Technical Report 06/367, Faculty of Economics and Business Administration, Ghent University, Belgium.

Antonov, A., Mechkov, S. & Misirpashaev, T. (2005), Analytical techniques for synthetic CDOs and credit default risk measures, Technical report, Numerix.

Basel Committee on Bank Supervision (2005), 'Basel II: International convergence of capital measurement and capital standards: a revised framework'.

Beran, J. & Ocker, D. (2005), 'Small sample asymptotics for credit risk portfolios', *Journal of Computational & Graphical Statistics* **14**(2), 339–351.

Boyle, P., Broadie, M. & Glasserman, P. (1997), 'Monte carlo methods for security pricing', *Journal of Economic Dynamics and Control* **21**, 1267–1321.

Daniels, H. (1987), 'Tail probability approximations', *International Statistical Review* **55**, 37–48.

Davison, A. & Wang, S. (2002), 'Saddlepoint approximations as smoothers', *Biometrika* **89**(4), 933–938.

Debuysscher, A. & Szegö, M. (2003), The Fourier transform method - technical document, Technical report, Moody's Investors Service.

Glasserman, P. (2006), 'Measuring marginal risk contributions in credit portfolios', *Journal of Computational Finance* **9**(2), 1–41.

Glasserman, P. & Li, J. (2005), 'Importance sampling for credit portfolios', *Management Science* **51**(11), 1643–1656.

Glasserman, P. & Ruiz-Mata, J. (2006), 'Computing the credit loss distribution in the Gaussian copula model: a comparison of methods', *Journal of credit risk* **2**(4), 33–66.

Gordy, M. B. (2002), 'Saddlepoint approximation of CreditRisk$^+$', *Journal of Banking and Finance* **26**, 1335–1353.

Gordy, M. B. (2003), 'A risk-factor model foundation for ratings-based bank capital rules', *Journal of Financial Intermediation* **12**(3), 199–232.

Gourieroux, C., Laurent, J.-P. & Scaillet, O. (2000), 'Sensitivity analysis of values at risk', *Journal of Empirical Finance* **7**, 225–245.

Huang, X., Oosterlee, C. W. & van der Weide, J. A. M. (2006), Higher order saddlepoint approximations in the vasicek portfolio credit loss model, Technical Report 06-08, Department of Applied Mathematical Analysis, Delft University of Technology.

Jensen, J. (1995), *Saddlepoint Approximations*, Oxford University Press.

Kalkbrener, M., Lotter, H. & Overbeck, L. (2004), 'Sensible and efficient capital allocation for credit portfolios', *RISK* (January), S19–S24.

Lugannani, R. & Rice, S. (1980), 'Saddlepoint approximations for the distribution of the sum of independent random variables', *Advances in Applied Probability* **12**, 475–490.

Martin, R. (2004), *Credit portfolio modeling handbook*.

Martin, R. & Ordovás, R. (2006), 'An indirect view from the saddle', *RISK* (October), 94–99.

Martin, R., Thompson, K. & Browne, C. (2001*a*), 'Taking to the saddle', *RISK* (June), 91–94.

Martin, R., Thompson, K. & Browne, C. (2001*b*), 'VAR: who contributes and how much?', *RISK* (August), 99–102.

Mausser, H. & Rosen, D. (2004), Allocating credit capital with VaR contributions, Technical report, Algorithmics Inc.

Merino, S. & Nyfeler, M. A. (2005), 'Applying importance sampling for estimating coherent credit risk contributions', *Quantitative Finance* **4**, 199–207.

Shelton, D. (2004), Back to normal. proxy integration: a fast accurate method for CDO and CDO-squared pricing, Technical report, Citigroup Global Structured Credit Research.

Tasche, D. (2000), Conditional expectation as quantile derivative, Technical report, Technische Universitat Muenchen.

Tchistiakov, V., de Smet, J. & Hoogbruin, P.-P. (2004), 'A credit loss control variable', *RISK* (July), 81–85.

Thompson, K. & Ordovás, R. (2003), 'Credit ensembles', *RISK* (April), 67–72.

Vasicek, O. (1987), 'Probability of loss on loan portfolio'.

Vasicek, O. (1991), 'Limiting loan loss probability distribution'.

Vasicek, O. (2002), 'Loan portfolio value', *RISK* (December), 160–162.

Wallace, D. L. (1958), 'Asymptotic approximations to distributions', *The Annals of Mathematical Statistics* **29**(3), 635–654.

Wilde, T. (2001), 'Probing granularity', *RISK* (August), 103–106.

Yang, J., Hurd, T. & Zhang, X. (2006), 'Saddlepoint approximation method for pricing CDOs', *Journal of Computational Finance* **10**(1), 1–??

Zheng, H. (2006), 'Efficient hybrid methods for portfolio credit derivatives', *Quantitative Finance* **6**, 349–357.