Robust and Stable Capital Allocation

Bernd Appasamy, Anne Kleppe, Christian Oehler
d-fine GmbH, Opernplatz 2, 60313 Frankfurt, Germany,
e-mail: christian.oehler@d-fine.de

Abstract

Capital allocation of credit risk is one of the key concepts of modern value based bank management. In order to conduct business decisions according to some risk-reward relationship the reliability of capital allocation figures on transaction level is most critical. Conventional allocation techniques like volatility contributions or credit value at risk (CVaR) are either non-coherent or fail to achieve the numerical stability needed for bank management decisions. Within the financial industry expected shortfall as allocation scheme becomes more and more popular since it fulfils all properties of a risk sensitive and coherent risk measure. However, expected shortfall allocation is impossible to be implemented within crude Monte Carlo portfolio modelling since risk contributions on transaction level turn out to be numerically totally unstable. In this article the impact of variance reduction techniques on the numerical stability of expected shortfall risk contributions is investigated. The authors prove that by applying two different numerical variance reduction techniques in combination it is finally possible to obtain robust and stable expected shortfall risk contributions at transaction level within any real portfolio.

Keywords

Economic capital, capital allocation within real world portfolios, expected shortfall, variance reduction, importance sampling

1 Introduction

Increasing competition within the global financial industry has stimulated banks to find new ways of linking front office, risk controlling and treasury to manage capital on an optimal risk-reward relationship. The success of a financial institution greatly depends on its ability to achieve sustainable high returns above the cost of capital on every single transaction. As such profitability management must be regarded as one of the key ingredients to modern value-based bank management.

Taking a closer look at risk return related profitability measures, it turns out that these are highly sensitive to the accuracy with which contributory economic capital figures can be determined. Even small uncertainties in the estimation of risk contributions can lead to completely unstable profitability figures and can by far outweigh uncertainties in the estimation of risk adjusted return. Thus profitability figures, if based on inaccurate economic capital contributions, are rendered meaningless and may initiate misleading business decisions.

The main focus of this article will be on how to achieve meaningful contributory economic capital figures in the context of credit risk portfolio modelling. Here, Monte Carlo techniques based on asset return processes have been established as the most popular modelling approach. However, crude Monte Carlo alone is insufficient for obtaining stable contributory risk figures at transaction level. Costly scale-up of computer hardware and the development of sophisticated multithreading or parallel processing techniques are the consequence. But even having the best hardware and the most sophisticated parallel processing mechanism in place the goal to calculate stable and robust risk capital figures cannot be achieved without appropriate variance reduction techniques.

In this article two very efficient methods of variance reduction are analysed. These will be proven to finally resolve the issue of robust and stable Monte Carlo based economic capital risk contributions. The level of variance reduction that can be attained for real bank portfolios is analyzed with particular focus on high portfolio loss quantiles. With these methods at hand meaningful risk contribution figures can be assigned at any level of portfolio aggregation and finally down to transaction level.

2 Capital Allocation Schemes

Within the financial industry three different approaches to capital allocation are adopted. On the one hand side, these are volatility contributions and marginal or incremental CVaR, both well established, and on the other side expected shortfall contributions which become more and more popular.

Calculation of the capital charge of an individual transaction based on the volatility contribution is analytically tractable by a so-called variance-covariance technique [2]. Although intrinsically stable, the distribution of volatility of the loss distribution amongst contributory transactions is not suitable as risk measure: Volatility contributions are neither sub-additive nor fully sensitive to risk concentration at high loss quantiles. Volatility contributions have a missing diversification property making the bank blind against real default correlation effects and last but not least may yield contributory economic capital figures which exceed the total exposure of a single transaction [2].

Marginal or incremental CVaR is the most widespread capital allocation scheme. It measures the sensitivity of economic capital at a certain loss quantile with respect to some transaction under consideration. Within a Monte Carlo approach this risk figure is easily obtained by calculating the delta of group economic capital with and without the transaction under consideration. Although sensitive to high default correlations marginal CVaR, just like the volatility contribution scheme, is not sub-additive and hence not suitable as risk measure. Since the marginal CVaR allocation scheme can be easily implemented in any portfolio modelling system it is in widespread usage among financial institutions. This is surprising since marginal CVaR is the most unstable and least robust allocation scheme among the three methods discussed here.

In contrast to volatility contributions and marginal CVaR, expected shortfall (ESF) leads to truly risk sensitive capital allocation figures. Expected shortfall can intuitively be regarded as the average portfolio loss above a pre-defined quantile (ESF-threshold) of the loss distribution minus the expected loss (EL).

$$ESF = E\left[Portfolioloss \;\middle|\; Portfolioloss \geq ESF\ threshold\right] - EL$$

The expected shortfall contribution of a single transaction to the portfolio ESF follows in a straightforward way: It is given by the transaction's average contribution to the portfolio loss, given that the total portfolio loss is above the ESF-threshold.

$$ESF_i = E[Loss_i \mid Portfolioloss \geq ESF\ threshold] - EL_i$$

Expected shortfall is a coherent risk measure [8] and besides its sub-additivity feature it exhibits a long list of good properties which make it the best-practice choice for capital allocation, cf. Figure 1. As an integrated measure it is intrinsically more stable than incremental CVaR. In contrast to the variance-covariance approach risk contributions based on expected shortfall never exceed a transaction's total exposure. Expected shortfall is highly sensitive to risk concentration and therefore enables the clear identification of capital destroyers and value-creators. Consequently, it supports investment decisions most efficiently and effectively. Due to its many beneficial properties expected shortfall allocation currently is regarded as the best practice choice among all allocation schemes.

Expected Shortfall Benefits

Conventional allocation techniques (volatility contribution, marginal CVaR)	Capital allocation via Expected Shortfall
✗ Breakdown of volatility inconsistent w.r.t. rare tail events of the loss distribution	✓ Consistent and truly risk-sensitive capital figures
✗ Risk contributions of loans may exceed their exposure	✓ Risk contributions of loans always smaller than their exposure
✗ Sub-additivity violation leading to non-sensitivity against diversification benefits and default risk concentrations	✓ Sub-additivity allowing for detection of concentration risk and diversification benefits
✗ Marginal CVaR as non-integrated measure extremely unstable	✓ Transactions with higher correlations/concentration risk always consume more capital
✓ Stable risk contribution figures when derived from volatility contribution	✗ Without powerful variance reduction techniques numerically unstable due to volatility implied by statistical uncertainty

Figure 1: Pros and Cons of conventional allocation techniques (volatility contribution, marginal CVaR) in comparison to expected shortfall allocation.

3 Why Variance Reduction?

As outlined in the previous section expected shortfall is the best-practice choice for an economic capital allocation scheme. However, while expected shortfall yields risk figures way more stable than marginal/incremental CVaR, the achievable level of stability for ESF-contributions at transaction level is still totally unacceptable within crude Monte Carlo simulations.

The problem with expected shortfall contributions at high loss quantiles is due to the fact that out of the total number of Monte Carlo scenarios generated within a Monte Carlo run only very few scenarios contribute to the expected shortfall estimator. For example, consider a Monte Carlo run consisting of 100,000 scenarios and a relatively high loss quantile of, say, 99.9%. On average only 100 scenarios out of the 100,000 scenarios will contribute to the portfolio-level ESF-estimator. These 100 scenarios above the ESF threshold will usually be mainly driven by a few transactions of counterparties with high exposure, high PD's and high default correlations. Some of the transactions in the portfolio might in fact not contribute to any of the losses of these 100 scenarios. Therefore, on transaction level most risk contribution figures will be completely arbitrary and thus meaningless. As such more or less all variance reduction techniques aim at exploiting the total number of generated Monte Carlo scenarios more effectively.

4 Variance Reduction Techniques

Since stable contributory capital risk figures cannot be achieved by means of crude Monte Carlo simulations, sophisticated statistical variance reduction techniques must be implemented. The variance reduction techniques discussed and analyzed in this article are antithetic variables denoted by method M1, the method of importance sampling (IS) denoted by M2 and a method based on conditional independence referred to as method M3.

Among the different approaches to importance sampling found in the literature [2-4] the results outlined here are based on the ideas developed in [2]. Within this approach the underlying optimization problem is solved by mapping all risk factors of the original bank portfolio to a completely homogeneous portfolio of infinite granularity. The method abbreviated by M3 is based on the idea of conditional independence of default/ migration behaviour between counterparties once the systematic risk in the asset return process is known. The idea of conditional independence is originally described in [5]. Within this method the idea of conditional independence is generalized to the estimation of ESF risk contributions at high economic capital quantiles.

Finally the usage of antithetic variables (method M1) will briefly be discussed. This very simple technique uses mirror scenarios of each Monte Carlo scenario thus generating one extra scenario without additional calculation time. It can be easily implemented resulting in roughly a factor of two in variance reduction compared to crude Monte Carlo. Since more a statistical trick than an elaborate method, antithetic variables are not analyzed further within this article.

5 Definition of Test Portfolios

In order to study the effectiveness of the different variance reduction techniques the achievable level of variance reduction shall be analyzed within structurally very different portfolios. As such we shall consider three different test portfolios. The first portfolio is a generic and, with respect to its risk parameters, completely homogeneous benchmark portfolio. The second test portfolio is an extremely inhomogeneous real portfolio of bank counterparties. The third test portfolio is a random sample of a real credit portfolio of some universal bank.

The importance sampling approach described in [2] which is adopted in this article relies on the approximation of the given portfolio by an infinitely granular and homogeneous one. As such the achievable level of variance reduction naturally depends on the homogeneity of the portfolio under consideration: The less homogeneous the portfolio, the less effective is importance sampling. However, while some portfolios might be relatively close to this idealized homogeneous portfolio, realistic credit portfolios in general are quite far from being homogeneous.

In the following the three test portfolios shall be described in some more detail.

5.1. Benchmark portfolio

This portfolio is a generic, totally homogeneous test portfolio comprising 1000 single transactions. The homogeneity is reflected in the choice of risk parameters being the same for all transactions and in the fact that all transactions are mapped on to the same systematic factor. This kind of portfolio is conventionally referred to as infinitely granular and homogeneous and is frequently discussed in the literature. It leads to a single factor model with the resulting loss distribution being the Vasicek distribution [1]. The capital formula underlying the Basel II capital accord is one example of such a homogenous portfolio model approach [9]. The level of variance reduction that can be achieved via importance sampling is maximal within this portfolio which can hence be regarded as benchmark portfolio. The probability of default is set to 9bp for all transactions and the specific risk factor R^2 is assumed to be 39%.

5.2. Test portfolio No.1

This test portfolio is a characteristic real bank counterparty portfolio. It comprises 250 aggregated exposures to 250 international banks as counterparties. The portfolio is special in the sense that apart from a few counterparties with low ratings most bank counterparties have very small PD's in the range of 2 to 10 bp according to their rating classes AAA, AA and A as given by the bank's internal Basel II master scale. LGD is set to 100% for all loans as it has no further implication on the test calculations. Most of the counterparties default in very few scenarios only. The population of the loss distribution at high quantiles is therefore mainly dominated by defaults of the few banks with low creditworthiness. Both with respect to the dimensionality of the factor space and with respect to the distribution of the risk parameters this portfolio must be considered as extremely inhomogeneous. Importance sampling is therefore expected to have only a limited effect on the achievable level of variance reduction. The distributions of the risk parameters are illustrated in Figure 2.

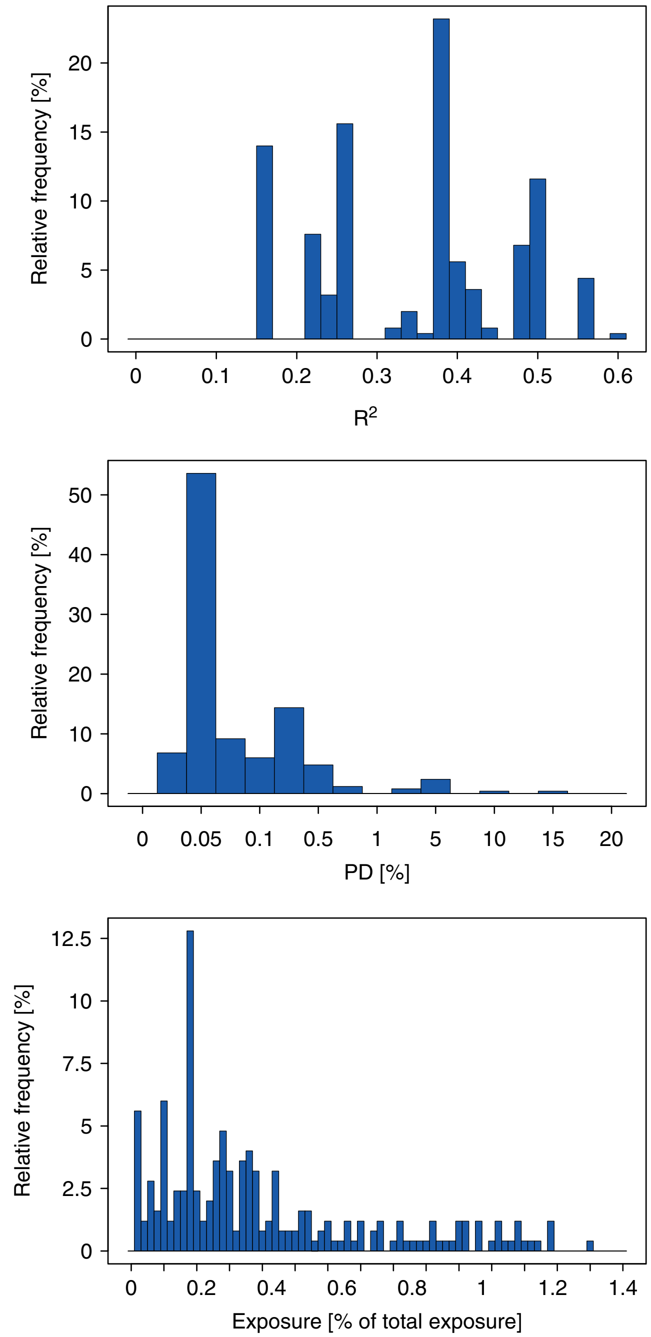

Figure 2: Distributions of risk parameters for test portfolio No. 1, a characteristic portfolio of loans to international banks. The systematic risk represented by R^2 is relatively high, since the risk of financial institutions can well be described by macroeconomic variables. The PD distribution is clustered with most banks exhibiting very small PDs between 1 and 10 bp but some counterparties having PD's up to 15%, e.g. some banks in Russia and South East Europe. Resulting from this is an extremely inhomogeneous EL distribution.

This international bank portfolio is mapped to 31 country specific economic factors, for which the R^2 parameter is determined by regression between the counterparty equity index and the country factor itself [10]. Since all counterparties belong to the financial industry the mapping to industry risk related factors has been omitted: that is, all counterparties are mapped to the same industry risk factor denoted by financial institutions.

5.3. Test portfolio No.2

This test portfolio is a sample comprising 1000 aggregated exposures to clients within different business areas. The sample was drawn randomly from a larger portfolio of a universal bank. The loan portfolio is an inhomogeneous mixture of retail clients, banks, non-bank financial institutions, SME's, as well as medium and large size corporations. The portfolio is totally inhomogeneous by its nature as reflected in the distributions of its risk parameters, cf. Figure 3.

Altogether the counterparties of this portfolio are mapped to 63 individual country and industry specific economic factors. The according systematic risk contribution R^2 is determined by linear regression techniques.

6 Numerical Results

Before presenting numerical results for the three test portfolios outlined in the previous section some definitions and calculation rules regarding the measurement of variance reduction shall be firstly discussed.

6.1. Measurement of variance reduction

The achievable level of variance reduction for ESF estimates on transaction level is assessed within a given test portfolio by performing a series of Monte Carlo runs. The single Monte Carlo runs are based on exactly the same set of parameters apart from the choice of the random seed which is different for each run. This series of Monte Carlo runs is performed twice, with and without variance reduction turned on. For both series the same sequence of random seeds is employed. Monte Carlo simulations performed without any acceleration method will be denoted by crude MC.

Within such a series of Monte Carlo runs ESF risk contribution estimates at transaction level will show statistical fluctuation from simulation to simulation. As a measure of this fluctuation the variance of ESF risk contributions over all Monte Carlo runs is calculated for both, the simulations with and without variance reduction method employed.

Following the steps above the variance reduction factor for the ESF estimate of some transaction i is calculated in the following way (presented here for the example of importance sampling only)

$$\text{Var\_Reduction}_i^{IS} = \text{Var}_i(\text{crude MC})/\text{Var}_i(\text{IS}).$$

Here, in contrast to crude MC, IS denotes simulations with importance sampling turned on. Similar notation will be used in the case of method M3.

The variance reduction factor can be directly translated into the Monte Carlo acceleration factor: the level of stability achieved within accelerated Monte Carlo can be achieved within crude MC when the number of scenarios

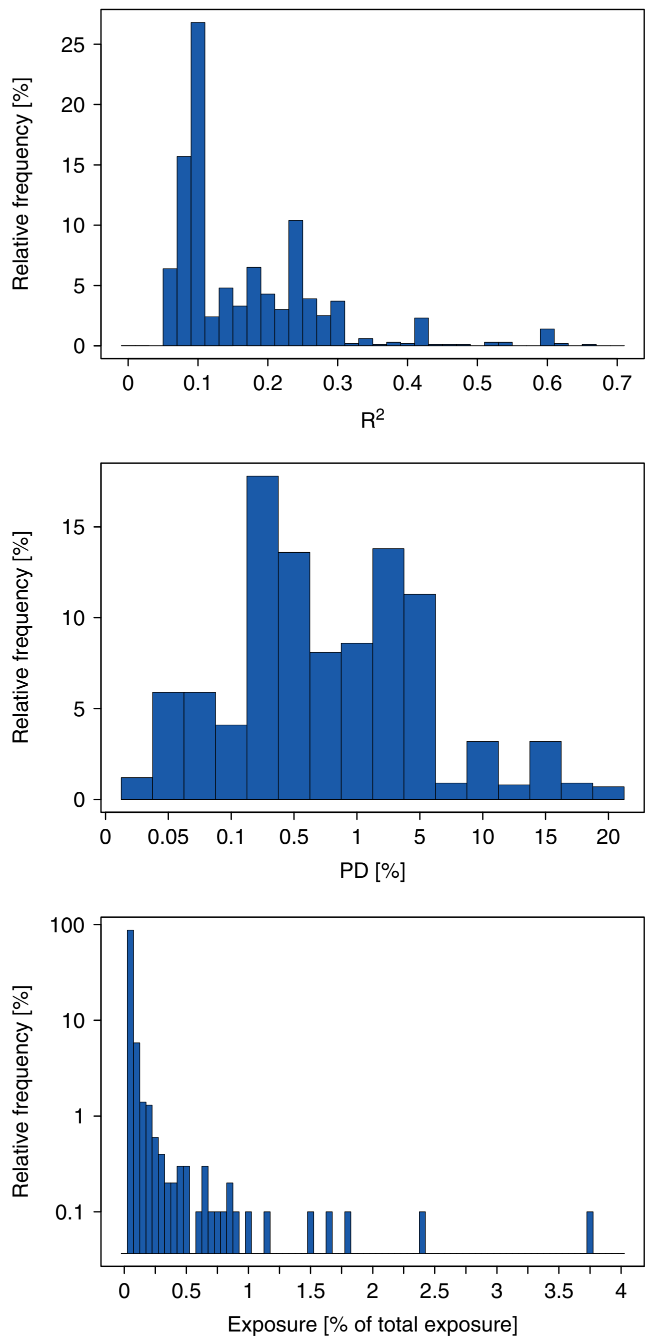

Figure 3: Distributions of risk parameters for test portfolio No. 2, a realistic credit portfolio. It is an inhomogeneous mixture of 1000 aggregated exposures to clients in different business areas.

within each run is multiplied by the variance reduction factor. As an example consider a run with 200.000 Monte Carlo scenarios and assume that for some given transaction the employed variance reduction technique has been shown to result in a variance reduction factor of 500. Therefore, in order to get the same level of accuracy for this transaction as can be achieved within accelerated Monte Carlo, 200.000 X 500 = 100 million scenarios would have to be performed within crude MC.

The arithmetic mean of the variance reduction factors on transaction level is calculated and defined as the portfolio specific average variance reduction factor (here again for the example of importance sampling only)

$$\langle \mathrm{Var\_Reduction}^{IS} \rangle_{portfolio} = \frac{1}{\#\mathrm{transactions}} \sum_{i=1}^{\#\mathrm{transactions}} \mathrm{Var}_i(\mathrm{crude\ MC})/\mathrm{Var}_i(\mathrm{IS})$$

It should be noted, however, that the average factor of variance reduction determined this way can only be considered as a rough estimate for the effectiveness of the method employed. In particular, it does not convey any information about the achievable level of accuracy of ESF risk contributions at transaction level. Most importantly the effectiveness of some employed variance reduction techniques greatly depends on whether the achievable level of accuracy is the same for all transactions. As such the variance reduction should be the greater the greater the relative variance of the ESF risk contributions within crude MC is.

For obvious reasons, risk contributions for which the relative standard deviation (standard deviation/average risk contribution) is already small within crude MC, say, of the order of a few percent, the variance reduction factor does not need to be as high as for risk contributions for which the standard deviation/average is large within crude MC.

Therefore, the effectiveness of an acceleration method can be analyzed and represented best by means of a scatter plot (e.g. Figure 4 B. to Figure 7 B.). Here the single transactions' relative standard deviations within the accelerated Monte Carlo runs are displayed as a function of the relative standard deviation within the crude MC simulations. Each point in such a scatter plot denotes the level of accuracy of some transaction's ESF risk contribution before and after acceleration. Points on the bisecting line correspond to no acceleration, i.e. for such transactions the acceleration method has no impact on the level of stability at all. Points above the bisector correspond to transactions for which stability after employing variance reduction is worse compared to crude MC, i.e. on these transactions the acceleration method has even the contrary effect. The ideal case is a constant line at the level of a few percent of ESF uncertainty, such that each ESF estimate shows more or less the same level of accuracy regardless how unstable it was within crude MC.

6.2. Variance reduction results

In this section the achievable level of variance reduction will be discussed in detail for the three test portfolios as specified in section 5.

The numerical results for these test portfolios are based on three series of Monte Carlo runs each. The run parameters are summarized in Table 1.

The EC-quantile for the calculation of the expected shortfall risk contributions is set to 99.9% for all simulations. The number of Monte Carlo simulations per run is chosen appropriately to reflect the risk structure of the respective test portfolio. They correspond to the minimal number of simulations needed to obtain un-biased variance reduction factors.

The results of the variance reduction analysis are illustrated in Figure 4 to Figure 7 separately for each test portfolio. Variance, standard deviation and average of the transaction specific ESF risk contributions are determined over each series of Monte Carlo runs. In all plots Figures A.) show the

Table 1: Run parameters for the three different sets of Monte Carlo simulations for the benchmark portfolio and test portfolios No. 1 and No. 2.

Portfolio	Number of runs for each set (crude MC, IS, IS & M3)	Number of simulations per run
Benchmark	20	50,000
Test portfolio 1	20	300,000
Test portfolio 2	20	200,000

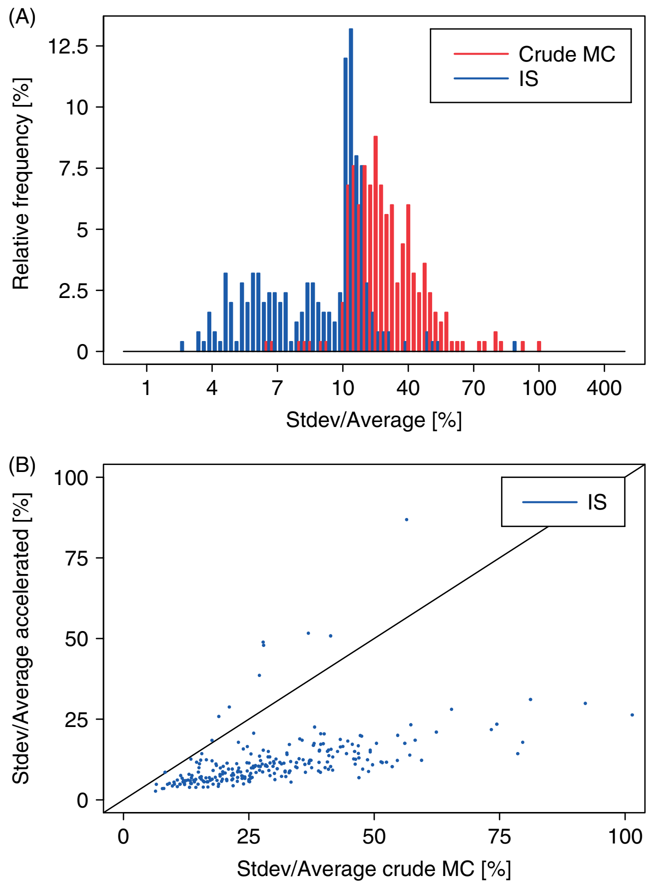

Figure 4: Test portfolio No. 1: A) Distribution of the relative standard deviation of ESF-risk contributions for single transactions; B) Relative standard deviation for single transactions within crude MC versus relative standard deviation when employing IS.

frequency distributions of the standard deviation/average ratio of the transaction specific ESF estimates. Figures B.) show the scatter plots of standard deviation/average of ESF estimates using accelerated Monte Carlo versus standard deviation/average within crude MC.

The numerical results of the variance reduction at portfolio level are summarized in Table 2. Here the average variance reduction factors at portfolio level, the minimal and maximal variance reduction at transaction level and the average level of accuracy are given for crude MC, importance sampling and importance sampling plus method M3, respectively. For a given portfolio and series of runs the average level of accuracy is defined as the expectation value of the frequency distributions (arithmetic mean of the single transactions' ratios of standard deviation and average), cf. Figure 4 A.) to Figure 7 A.).

Taking a closer look at the results for the benchmark portfolio it becomes obvious that importance sampling alone is already sufficient in order to achieve the desired level of accuracy for the transaction specific risk contributions. The average level of accuracy of 3% in the case of importance sampling, cf. Table 2, must be considered as highly satisfactory given the relatively low number of Monte Carlo scenarios of 50,000. Table 2 also shows the minimal and the maximal variance reduction attained by means of IS in this benchmark portfolio. The minimal variance reduction factor of 11 corresponds to those transactions which where already close to stable within crude MC whereas the maximal value of 503 corresponds to those transactions with the highest uncertainty within crude MC.

These highly satisfactory results come to no surprise as the benchmark portfolio is completely homogeneous and as such by its very construction ideal for the application of importance sampling as developed in [2].

As such the level of variance reduction achievable with importance sampling within this benchmark portfolio must be regarded as an upper limit. For portfolios with risk parameters deviating from the homogeneity assumption the achievable level of variance reduction by means of importance sampling will be in general way less than this upper limit.

In addition to importance sampling, method M3 is applied to the benchmark portfolio. In this case the average variance reduction over all transactions is 1663. The maximal variance reduction factor at transaction level is 5291. The overall level of accuracy is at a level of 0.7% which must be regarded as highly stable given only 50,000 Monte Carlo scenarios.

The results for test portfolio No. 1 are displayed in Figure 4 (IS) and Figure 5 (IS + method M3). The variance reduction factors resulting from the application of importance sampling as the only acceleration method are far from optimal. Large overlapping regions and missing separation between the frequency distributions generated with crude MC and importance sampling indicate the poor variance reduction for many transactions. The scatter plot, however, looks a bit more promising. Here a certain trend between the size of original variance within crude MC and variance when importance sampling is turned on can be recognized: ESF estimates having a larger uncertainty within crude MC show a higher variance reduction on average. However there are some transactions with corresponding data entries above the bisector resulting in lower stability after importance sampling is applied. Within this real bank portfolio importance sampling cuts down the variance on average but is far from sufficient in order to obtain stable risk contributions on transaction level.

In a second step the variance reduction analysis is performed for IS in conjunction with method M3. The corresponding results are displayed in Figure 5.

The frequency distributions are now completely separated. The scatter plot shows that variance is drastically cut down and that the same high level of stability can be achieved for each transaction within the portfolio. Those transactions with the highest variance in crude MC exhibit the highest acceleration factor. The maximum value of variance reduction observed at transaction level is 3272 and the average level of stability is 2% (see Table 2). Taking into account the extreme inhomogeneity of this test portfolio the results must be regarded as excellent. Even though this portfolio of bank counterparties can be regarded as a worst-case portfolio with respect to the estimation of stable ESF risk contributions at transaction level the employed variance reduction techniques lead to sufficiently stable ESF estimates. As such it can be concluded that employing both variance reduction techniques in conjunction, importance sampling and method M3, will lead to stable expected shortfall estimates on transaction level within any realistic portfolio.

Test portfolio No. 2 is a portfolio of a universal bank comprising clients within different business areas. Due to this mixture of business areas the distribution of exposure and PD is more uniformly populated compared to the bank counterparty portfolio discussed above. However, this portfolio still heavily deviates from the homogeneity assumption

Table 2: Numerical variance reduction results for the benchmark portfolio and test portfolios No. 1 and No. 2.

		Benchmark portfolio	Portfolio No. 1	Portfolio No. 2
Average variance reduction on transaction level	IS	103	10	11
	IS& M3	1663	254	403
Minimal variance reduction on transaction level	IS	11	0.3	0.3
	IS & M3	312	11	2
Maximal variance reduction on transaction level	IS	503	62	294
	IS & M3	5291	3272	36306
Level of accuracy of ESF contributions (Stdev/Average)	Crude MC	30%	30%	150%
	IS	3%	10%	30%
	IS & M3	0.7%	2%	5%

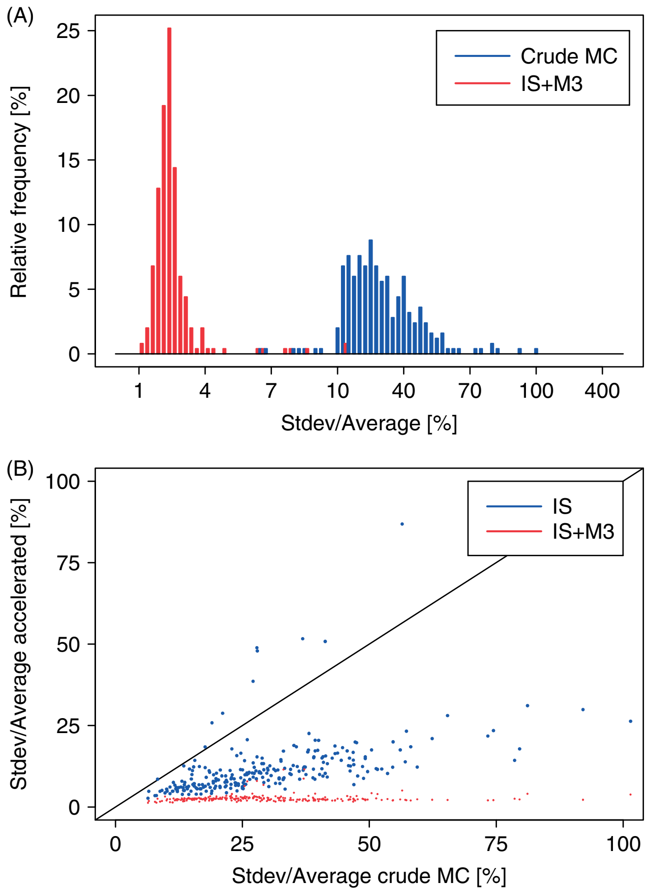

Figure 5: Test portfolio No. 1: A) Distribution of the relative standard deviation of ESF-risk contributions for single transactions; B) Relative standard deviation for single transactions within crude MC versus relative standard deviation when employing IS in connection with method M3.

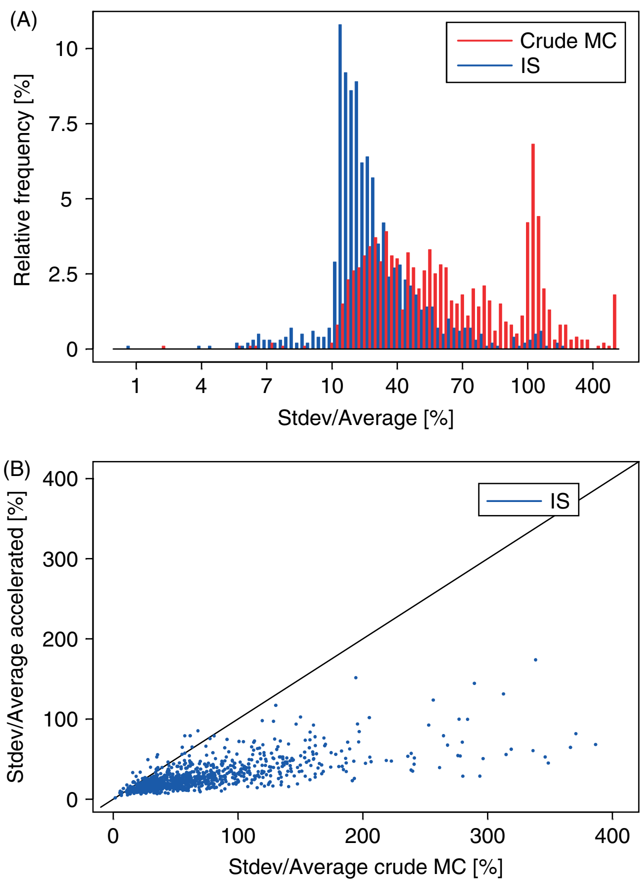

Figure 6: Test portfolio No. 2: A) Distribution of the relative standard deviation of ESF-risk contributions for single transactions; B) Relative standard deviation for single transactions within crude MC versus relative standard deviation when employing IS.

made within the importance sampling approach by [2]. The results for test portfolio No. 2 are displayed in Figure 6 for the case of importance sampling and Figure 7 for importance sampling in conjunction with method M3. The variance reduction effects are very similar to those observed for test portfolio No. 1. Importance sampling alone is insufficient in order to achieve the desired level of accuracy at transaction level. However, method M3 in conjunction with importance sampling is again highly efficient and resolves the stability issue of ESF estimates for each transaction within the test portfolio. The maximal variance reduction at transaction level is 36,306 (cf. Table 2) meaning that the same level of accuracy for this transaction would have been achieved by 36,306 x 200,000 scenarios within crude MC. The average level of accuracy achieved for this portfolio is 5% which is a little bit less compared to the value of 3% achieved for test portfolio No. 1. Nevertheless the level of stability is in good concordance since the number of Monte Carlo scenarios for each run for test portfolio No. 2, which comprises 1,000 transactions, was set to only 200,000 instead of 300,000 for test portfolio No. 1, which comprises only 250 transactions.

As mentioned previously the described results of the variance reduction analysis were obtained at a fixed EC-quantile of 99.9%. However, the level of variance reduction that can be achieved strongly depends on the EC-quantile under consideration. The dependence of the variance reduction factor on the considered EC-quantile shall be thus considered in the following in some more detail. This dependence is of particular interest since concentration risk of banks affects most heavily those loss scenarios within the extreme tail of the portfolio loss distribution. Since ESF as a risk measure is particularly sensitive to such risk concentrations capital allocation is usually conducted at high EC-quantiles. However, though stable risk contribution figures at high quantiles are most urgently needed they normally exhibit the lowest stability. In order to study the effect of variance reduction on risk contributions at high EC-quantiles, the average variance reduction factor was estimated at different EC-quantiles starting at 90% and going up to 99.99%. This analysis was performed both for the case of importance sampling only and for importance sampling in addition to method M3. The test portfolio under consideration was test portfolio No 2 as defined in section 5.

The results of this analysis are displayed in Figure 8. It shows the average variance reduction as a function of the EC-quantile (equivalent to the expected shortfall threshold) in the case of both IS and IS + M3. In both cases a steep increase of the average variance reduction above EC-quantiles

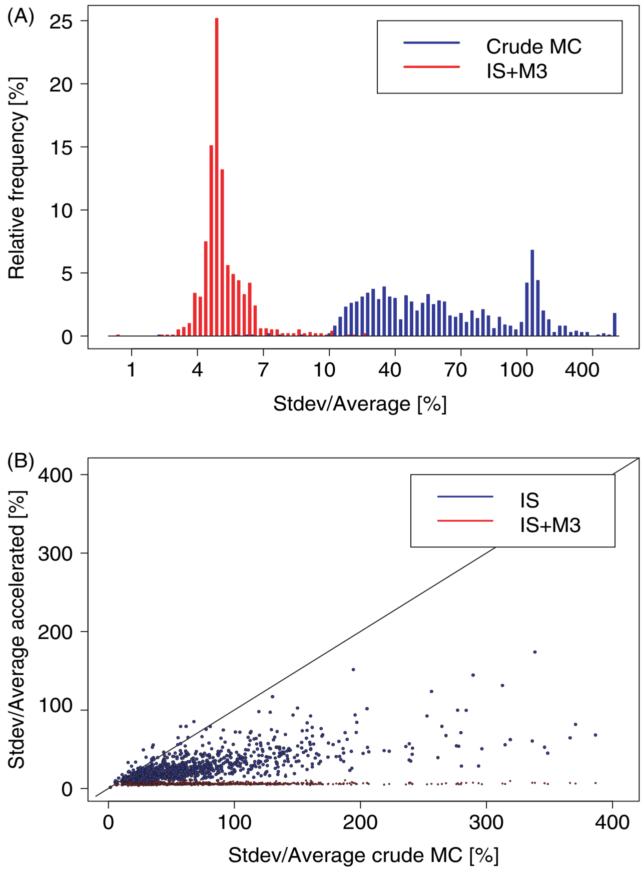

Figure 7: Test portfolio No. 2: Distribution of the relative standard deviation of ESF-risk contributions for single transactions; B) Relative standard deviation for single transactions within crude MC versus relative standard deviation when employing IS in connection with method M3.

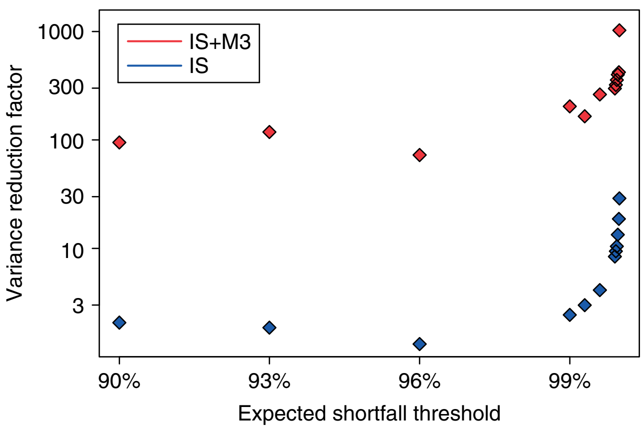

Figure 8: Dependence of the average variance reduction factor of test portfolio No. 2 on the expected shortfall threshold.

of 99% is observed. The impact of the employed variance reduction techniques is thus particularly high within the heavy tail regions of the loss distribution which are in turn of highest interest for risk management and bank steering aspects. Applying importance sampling only, the average variance reduction factor is only 2 at the 90% quantile. It does increase to a level of 10 at the 99.9% quantile and then steeply rises to a factor of 28 at the 99.99% confidence level. However, in the case of applying importance sampling and method M3 in conjunction the variance reduction factors are already significant at lower quantiles. Compared to the case when applying importance sampling only the average variance reduction factors increase by more than an order of magnitude. The lowest value at the 90% EC-quantile is roughly a factor of 150 to be compared to the factor of 2 with IS only. At 99.9% the resulting variance reduction factor is already at a level of 300 and further increases to a value of more than 1000 at the 99.99% level.

7 Conclusion

In this article the impact of variance reduction techniques on the numerical stability of expected shortfall risk contributions at transaction level has been investigated. In order to obtain stable ESF risk contributions two very efficient variance reduction techniques were employed and tested on a homogeneous benchmark portfolio as well as on two real portfolios. The variance reduction achieved by importance sampling turned out to be insufficient in the case of the considered real portfolios. In contrast, applying method M3 in conjunction with importance sampling turned out to be highly efficient and resolved the stability issue of ESF estimates for each transaction within the test portfolios. With 200,000 simulated Monte Carlo scenarios an average level of accuracy of the order of a few percent could be achieved at arbitrary EC-quantiles of up to 99.99%. Therefore the combination only of the two variance reduction methods outlined in this article finally resolves the issue of unstable expected shortfall risk contributions at transaction level.

REFERENCES

[1] Christian Bluhm, Ludger Overbeck and Christoph Wagner. *An Introduction to Credit Risk Modeling*. Financial Mathematics. Chapman & Hall/CRC 2003.

[2] Michael Kalkbrenner, Hans Lotter and Ludger Overbeck. *Sensible and efficient capital allocation for credit portfolios*. RISK, January 2004.

[3] Daniel Egloff, Markus Leipold, Stephan Jöri and Curdin Dalbert. *Optimal Importance Sampling for Credit Portfolios with Stochastic Approximation*. Working Paper, March 2005, University of Zürich (Swiss Banking Institute) and Züricher Kantonalbank.

[4] Paul Glasserman, Jingyi Li. *Importance Sampling for Portfolio Credit Risk*. Management Science, Vol. 51, No. 11, November 2005.

[5] Christopher C. Finger. *Conditional Approaches for CreditMetrics Portfolio Distributions*. CreditMetrics® Monitor. April 1999.

[6] Richard Durret. *Probability: Theory and Examples* - Third Edition. Duxbury Advanced Series. Thomson - Brooks/Cole 2005.

[7] Steven E. Shreve. *Stochastic Calculus for Finance II - Continuous Time Models*. Springer Finance. Springer, 2004.

[8] Artzner P, F Delbaen, J-M Eber and D Heath. *Coherent measures of risk*. Mathematical Finance 9, pages 203-228, 1999.

[9] Basel Committee on Banking Supervision. *International Convergence of Capital Measurement and Capital Standards - A Revised Framework*. Bank for International Settlements, June 2006.

[10] Venables, W.N., and Ripley B.D. 1999/2000. *Modern Applied Statistics with S-Plus*. 3rd ed. (2 volumes). Springer-Verlag, New York.