Using Expected Shortfall for Credit Risk Regulation

Kjartan Kloster Osmundsen*

University of Stavanger

February 26, 2017

Abstract

The Basel Committee's minimum capital requirement function for banks' credit risk is based on value at risk. This paper performs a statistical and economic analysis of the consequences of instead basing it on expected shortfall, a switch that has already been set in motion for market risk. The empirical analysis is carried out by means of both theoretical simulations and real data from a Norwegian savings bank group's corporate portfolio. Expected shortfall has some well known conceptual advantages compared to value at risk, primarily a better ability to capture tail risk. It is also sub-additive in general, thus always reflecting the positive effect of diversification. These two aspects are examined in detail, in addition to comparing parameter sensitivity, estimation stability and backtesting methods for the two risk measures. All comparisons are conducted within the Basel Committee's minimum capital requirement framework. The findings support a switch from value at risk to expected shortfall for credit risk modelling.

Keywords: Expected shortfall, credit risk, bank regulation, Basel III, tail risk

1 Introduction

This paper addresses the effects of shifting from value at risk (VaR) to expected shortfall (ES) as the underlying risk measure when computing regulatory capital requirements for banks'

*I would like to direct great thanks to Jacob Laading for stimulating discussions and constructive feedback. Thanks are due to a Norwegian savings bank group that provided the corporate portfolio data material used in the paper. My appreciation also goes to Roy Endré Dahl, Tore Selland Kleppe, Sindre Lorentzen and Atle Øglend for contributing with valuable inputs and feedback.

credit risk. These effects will be measured using both simulated data and real data from a Norwegian savings bank group's corporate portfolio.

The Basel Committee on Banking Supervision aims to enhance financial stability worldwide, partly by setting minimum standards for the regulation and supervision of banks [4]. In 2004, the introduction of the Committee's second international regulatory accord, Basel II [1], opened the possibility for banks to calculate their minimum capital requirements using risk parameters estimated by internal models, instead of using given standard rates (the standardised approach). The Basel Committee's minimum capital requirement is calculated by a function that takes estimates of the following risk parameters as input parameters: probability of default (PD), loss given default (LGD) and exposure at default (EAD). This function is derived using VaR as the underlying risk measure, and we will examine how the function's properties change when it is derived using ES.

Banks have been allowed to use internal models as a basis for calculating their *market risk* capital requirements since 1997 [4], i.e. seven years before the same applied for credit risk. Internal credit risk models were not allowed at an earlier stage due to the fact that they are not a simple extension of their market risk counterparts. Data limitations is a key impediment to the design and implementation of credit risk models [5]. Most credit instruments are not listed with a market value, implying that there are no historical prices to base future projections on. As there is no market values to compare with the book values, there is no impairment loss. Loss occurs only at default events, and the infrequent nature of these events makes it difficult to collect enough relevant data. The long time horizons also make the validation of credit risk models fundamentally more difficult than the backtesting of market risk models.

In January 2016, the Basel Committee published revised standards for calculation of minimum capital requirements for market risk [2], which include a shift from VaR to ES as the underlying risk measure. The Committee stated that the former reliance on VaR largely stems from historical precedent and common industry practice. This has been reinforced over time by the requirement to use VaR for regulatory capital purposes. However, the Committee recognized that a number of weaknesses have been identified with VaR, including its inability to capture tail risk [3]. There has currently not been considered a transition from VaR to ES for measuring credit risk. However, as the development of credit risk models lies a few years behind the market risk models, there is reason to believe that this might be considered in a not so distant future.

The paper is structured as follows. Section 2 introduces the risk parameters used for credit risk modelling, and shows how the Basel Committee's capital requirement function is derived with VaR as the underlying risk measure. Section 3 introduces ES and shows how the capital requirement function changes when applying this risk measure. An extensive comparison of the two different risk measures is presented in Section 4. Aside from a general comparison of the risk measures' properties, the section's main focus is a comparison of the two resulting versions of the capital requirement function. This includes confidence level calibration, backtesting methods and a simulation-based comparison of parameter sensitivity. We also examine how VaR and ES values are affected by the tail properties of the loss distributions, by simulating losses from distribution functions with different tail weights, using real estimates of risk parameters from a Norwegian savings bank group's corporate portfolio. Lastly, the final conclusions are given in Section 5.

2 Credit Risk Modelling

The introduction of Basel II in 2004 opened the possibility for banks to calculate the assets' risk weights using parameter estimates from internal models. To be able to use this *internal ratings based* (IRB) approach, the bank's risk models have to be approved by the national supervisory authorities.

In this section we introduce the risk parameters involved, and describe the model choices made by the Basel Committee when deriving the mathematical function for calculating regulatory capital under the IRB approach.

2.1 Risk Parameters

Probability of default (PD) is the probability that a borrower will be unable to meet the debt obligations. This probability is defined for a particular time horizon, typically one year.

Exposure at default (EAD) is the lender's outstanding exposure to the borrower in case of default.

Loss given default (LGD) is the lender's likely loss in case of default. Usually stated as a percentage of EAD.

The *expected loss* (EL) is the average credit loss a bank can expect on its credit portfolio over the chosen time horizon. The expected loss is calculated as the mean of the loss distribution, and is typically covered by provisioning and pricing policies [6]. The expected loss of a single loan can be calculated as follows:

$$EL = PD \cdot LGD \cdot EAD. \tag{1}$$

Banks typically express the risk of a portfolio with the *unexpected loss* (UL), which is the amount by which the actual credit loss exceeds the expected loss. The economic capital held to support a bank's credit risk exposure is usually determined so that the estimated probability of unexpected loss exceeding economic capital is less than a target insolvency rate. The potential unexpected loss which is judged too expensive to hold capital against, is called *stress loss*, and leads to insolvency. The probability density function of future credit losses is the basis for calculating the unexpected loss, and the target insolvency rate is chosen so that the resulting economic capital will cover all but the most extreme events.

2.2 The Basel Committee's Capital Requirement Function

Basel II made it possible for banks to use internal risk models to estimate PD, EAD and LGD for each individual loan [1]. These estimates are used as input parameters for a mathematical function that calculates the regulatory capital requirement for each loan. This capital requirement function is based on Gordy's *Asymptotic Single Risk Factor* (ASRF) model [7], which models risk using only one systematic risk factor, which may be interpreted as reflecting the state of the global economy. The model is constructed to be *portfolio-invariant*, so that the marginal capital requirement for a loan does not depend on the properties of the portfolio in which it is held.

The probability of default conditional on a systematic risk factor is calculated by Vasicek's adaptation of the Merton model [8]:

$$PD(X) = \Phi\left(\frac{\Phi^{-1}(PD) - X\sqrt{R}}{\sqrt{1-R}}\right), \tag{2}$$

where Φ is the cumulative distribution function of the standard normal distribution and R is the loan's correlation with the systematic risk factor X, i.e. the degree of the lender's exposure to the systematic risk factor. The unconditional PD on the right hand side reflects expected default rates under normal business conditions, and is estimated by the banks.

The ASRF model uses value at risk as the underlying risk measure, meaning that the required capital is calculated so that the loss probability does not exceed a set target q. This is achieved by holding capital that covers up to the q^{th} quantile of the loss distribution. By choosing a realization of the systematic risk factor equal to the q^{th} quantile $\alpha_q(X)$, we obtain the following expression as X is assumed to be normally distributed:

$$PD(\alpha_q(X)) = PD(\Phi^{-1}(1-q)) = PD(-\Phi^{-1}(q)) = \Phi\left(\frac{\Phi^{-1}(PD)+\Phi^{-1}(q)\sqrt{R}}{\sqrt{1-R}}\right). \tag{3}$$

The expected loss for each loan can be calculated using (1) without the EAD-factor, thus being expressed as a percentage of the exposure at default. Inserting (3) for PD, we get the q^{th} quantile of the expected loss conditional on the systematic risk factor X, i.e. the value at risk [7]:

$$\alpha_q(E[L|X]) = E[L|\alpha_q(X)] = PD(\alpha_q(X)) \cdot LGD. \tag{4}$$

The LGD value in (4) must reflect economic downturn conditions in circumstances where loss severities are expected to be higher during cyclical downturns than during typical business conditions [6]. This so-called *"downturn"* LGD value is not computed with a mapping function similar to (3). Instead, the Basel Committee has decided to let the banks provide downturn LGD values based on their internal assessments. The reason for this is the evolving nature of bank practices in the area of LGD quantification.

The Basel Committee's capital requirement only considers the unexpected loss. As the ASRF model delivers the entire value at risk, the expected loss $PD \cdot LGD$ has to be subtracted from (4). This results in the Basel Committee's capital requirement function:

$$K = LGD \cdot \Phi\left(\frac{\Phi^{-1}(PD)+\Phi^{-1}(0{,}999)\cdot\sqrt{R}}{\sqrt{1-R}}\right) - PD \cdot LGD, \tag{5}$$

where the Committee has chosen the confidence level $q = 0.999$. This means that losses on a loan should exceed the capital requirement only once in a thousand years. The reason why the confidence level is set so high is partly to protect against inevitable estimation error in the banks' internal models [6].

Under the Basel III regulation, banks must multiply (5) by a factor of 1.06, based on an impact study of Basel II conducted by the Basel Committee [9]. The capital requirement function is also multiplied by an adjustment factor for the maturity of the loan, as long-term credits have higher risk than short-term credits. The maturity adjustment MA is given by

$$MA = \frac{1 + (M - 2.5) \cdot b(PD)}{1 - 1.5 \cdot b(PD)},$$

where M is years to maturity and $b(PD) = (0.11852 - 0.05478 \cdot \ln(PD))^2$.

As the capital requirement (5) is expressed as a percentage of total exposure, one must multiply by EAD to get the capital requirement stated as a money amount. The total money amount shall constitute at least 8 % of the risk-weighted assets:

$$\sum_{i=1}^{n} K_i \cdot EAD_i \geq 0.08 \cdot \sum_{i=1}^{n} RW_i \cdot EAD_i,$$

where K_i is the calculated minimum capital requirement for asset i, RW_i is the risk weight assigned to asset i and EAD_i is the credit risk exposure of asset i.

Thus, the marginal risk-weight of a single asset is calculated as:

$$RW_i = \frac{K_i}{0.08} = 12.5 \cdot K_i.$$

3 Expected Shortfall

Value at risk is not a *coherent* risk measure, as it has been shown [10] that it is not sub-additive in general. Thus, a merger of two portfolios may have a greater VaR than the sum of the VaR of the individual portfolios. This contradicts basic diversification theory, and is considered as one of the biggest flaws of VaR. Another property of VaR that is often pointed out as a weakness is that it does not give any information about the size of the losses that occurs with a probability less than $1 - q$. This can be particularly problematic if the loss distribution is heavy-tailed, commonly referred to as tail risk. Assets with higher potential for large losses may appear less risky than assets with lower potential for large losses.

However, VaR is sub-additive if the loss distribution belongs to the elliptical distribution family and has finite variance, making it a coherent risk measure in these cases [11]. This includes the normal distribution, Student's t distribution (for $\nu > 2$) and Pareto distribution (for $\alpha > 2$). For these distributions, VaR becomes a scalar multiple of the distribution's standard deviation, which satisfies sub-additivity.

Even though value at risk is not sub-additive in general, it still remains the most widely used risk measure. The reason seems to be that its practical advantages are perceived to outweigh its theoretical shortcomings. Value at risk is considered to have smaller data requirements, easier backtesting and in some cases easier calculation than alternative risk measures [12]. Value at risk is also popular because of its conceptual simplicity. The economic capital calculated by VaR at a confidence level q corresponds to the capital needed to keep the firm's default probability below $100 \cdot (1-q)$ %.

As an alternative to value at risk, Artzner et al. [10] proposed a coherent risk measure called *tailed conditional expectation*(TCE). Acerbi and Tasche [13] proposed an extended version that is also coherent for non-continuous probability distributions:

Definition 1 (Expected Shortfall)**.** *Given a confidence level $q \in (0,1)$, expected shortfall is defined as*

$$ES_q(L) = E\left[L | L \geq VaR_q(L)\right] + \left(E\left[L | L \geq VaR_q(L)\right] - VaR_q(L)\right)\left(\frac{P\left[L \geq VaR_q(L)\right]}{1-q} - 1\right),$$

where $VaR_q(L)$ is the value at risk at the same confidence level.

When $P\left[L \geq \mathrm{VaR}_q(L)\right] = 1-q$, as is the case for continuous distributions, the last term from Definition 1 vanishes, and the expected shortfall equals the TCE.

By using the definition of conditional probability and a change of variables, ES can also be written as an integral over the VaR values for all confidence levels $u \geq q$:

$$\mathrm{ES}_q(L) = \frac{1}{1-q}\int_{u=q}^{1} \mathrm{VaR}_u(L)\, du. \qquad (6)$$

From Definition 1 and (6) it is clear that expected shortfall does not have the same degree of tail risk as value at risk. Unlike VaR, ES can distinguish between two distributions of future net worth that have the same quantile but differ otherwise.

A critique of ES is the fact that tail behaviour is taken into account through an averaging procedure. Medina and Munari [14] claim that averages are poor indicators of risk, thus making ES a potentially deceiving measure of risk.

The ASRF model is also applicable for expected shortfall, as ES-based capital charges are portfolio invariant under the same assumptions as VaR-based capital charges [7]. It is thus possible to derive a version of the Basel Committee's capital requirement function (5) that is based on expected shortfall [15]:

$$K = \frac{LGD}{1-q}\Phi_2\left(\Phi^{-1}(PD), -\Phi^{-1}(q); \sqrt{R}\right) - PD \cdot LGD, \tag{7}$$

where $\Phi_2(\cdot)$ is the bivariate cumulative normal distribution function.

4 Value at Risk Versus Expected Shortfall

In this section we will compare VaR and ES as credit risk measures, and examine how the Basel Committee's capital requirement function is affected by the choice of its underlying risk measure. We will discuss confidence level calibration, backtesting methods, parameter sensitivity and the shape of the loss distribution.

4.1 Sub-additivity and Tail Risk

Value at risk only satisfies sub-additivity when the loss distribution belongs to the elliptical distribution family and has finite variance [11]. Yamai and Yoshiba provide a simple example[1] of how the tail risk of VaR may result in serious practical problems in credit portfolios. A modified version of this example follows: first, suppose a bank holds a credit portfolio consisting of 100 corporate loans to different firms, each with a one year default probability of 1 percent, and a recovery rate of zero (LGD=100 %). The exposure at default is $1 million for each loan. For simplicity, it is assumed that the occurrences of defaults are mutually independent. From (1) we have that the expected loss for each loan is $ 10000. Assuming a 1 % net lending margin ($10000), each loan is thus priced at $20000. This means that the bank earns $20000 for each firm not defaulting, while it loses $1 million for each defaulting firm. Thus, the bank loses money if more than one firm defaults in one year, making the probability of loss approximately 26 % ($1 - 0.99^{100} - 100 \cdot 0.99^{99} \cdot 0.01$). As the probability of loss exceeds 5 %, the 95 % VaR for this diversified investment will have a positive value.

Second, we consider the bank investing the same total amount of $100 million in a large loan to only one of the firms. For this concentrated investment the probability of loss is only 1 % and the 95 % VaR is thus -$2 million: the loan price. As the probability of default is below 5 %, the potential of default is disregarded at the 95 % confidence level. We also observe that value at risk is not sub-additive in this case as the VaR of the diversified portfolio is larger than the VaR for the concentrated portfolio. Table 1 shows the value at risk and expected

[1] Example 2 in [11].

shortfall for both the diversified and the concentrated investment. We see that ES is able to detect the tail risk, resulting in correctly pointing out the concentrated investment as the most risky investment.

	95 % VaR	95 % ES
100 loans	\$1.06 million	\$1.52 million
1 loan	-\$2.0 million	\$18.4 million

Table 1: 95 % value at risk and expected shortfall for a diversified investment and a concentrated investment. Positive numbers correspond to loss, negative numbers indicate profit.

This example shows how value at risk can disregard the increase of potential loss due to credit concentration. One should therefore always ensure that credit concentration is limited by complementary measures when using VaR for risk management. In the Basel Committee's regulatory framework, this issue is addressed in Pillar 2 [1].

4.2 Confidence Level

The expected shortfall version of the Basel Committee's capital requirement function (7) was derived using the same assumptions as for the VaR version (5). Namely, the assumption of a normal distribution for the systematic risk factor, which leads to the loss distribution also being normal. However, a change to the more tail risk sensitive ES would most likely be motivated by real loss distributions being found more heavy-tailed than the normal distribution. A change to ES would thus probably also result in a model that assumes a more heavy-tailed loss distribution. In that case, it could be justifiable to apply a confidence level resulting in a slightly smaller capital requirement, as one could argue the increased tail risk sensitivity reduces the model risk.

Although the derived ES version of the capital requirement function (7) is based on the same loss distribution assumptions as the VaR version, the difference between the two risk measures is significant enough that the two functions behave quite differently. We now try to determine if it is possible to choose a confidence level for the ES version that makes it behave like the 99.9 % VaR version (5). Given the definition of ES, this confidence level must be lower than 99.9 %.

Conducting a least squares fit over the interval $PD \in (0,1)$, we find that the confidence level 99.742 % makes the ES version most similar to the 99.9 % VaR version. There is however considerable differences for the smallest PD values, as shown in Figure 1. The ES ver-

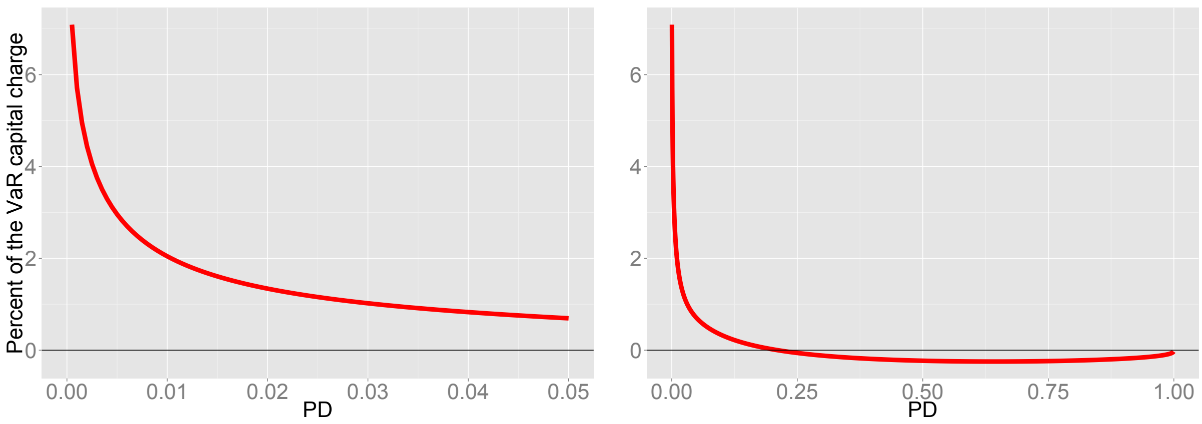

Figure 1: The difference between the calculated capital requirement from the ES version with confidence level 99.742 % and the standard 99.9 % VaR version. Positive values mean that the ES version results in a higher capital charge. The left graph gives a detailed view for small PD values, while the right graph shows the whole (0,1) interval.

sion slightly increase capital charges for loans with low probability of default, and slightly decrease capital charges for loans with probability of default exceeding 21 %. As the Basel Committee has proposed to apply floors for the PD estimates [16], this may be considered a good thing.

4.3 Loss Distributions

In this section we will examine how value at risk and expected shortfall depend on the tail of the loss distribution. Using parameter values from a real data set, we will simulate loss realizations by assuming different loss distributions. The simulated loss values are used to create VaR and ES estimates for different confidence levels. Both the level and the uncertainty of these estimates are compared. It is also tested how the results depend on the number of simulations.

4.3.1 Data Set

The data set contains information about corporate loans issued by a Norwegian savings bank group from March 2015 to January 2016. It contains about a fifth of the group's total corporate portfolio from this period, picked randomly. This amounts to a total of 109045 loans. For each loan, the data set contains numbers for EAD, LGD and PD. The correlations to the systematic risk factor are also included.

Figure 2 provides some insight about the data set, by displaying density plots of the LGD

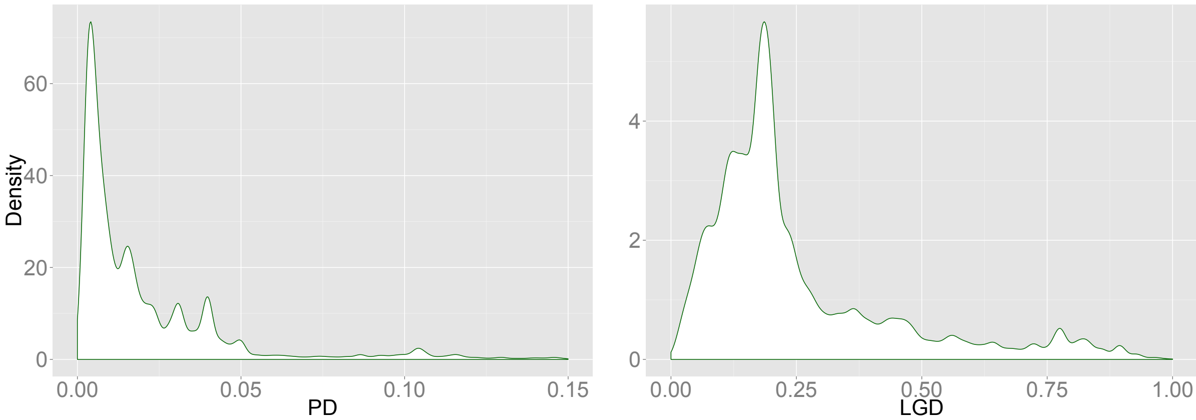

Figure 2: The distribution of LGD and PD values in the data set.

and PD values. We see that most of the loans have low risk. In fact, 90.8 % of the loans have been assigned a probability of default of 0.05 or less. Only 2.1 % of the loans have a PD value greater than 0.15 (not included in Figure 2). The majority of the values for loss given default is also in the low end of the scale, with 68 % of the loans having a LGD value of 0.25 or less. There are however also a substantial number of loans that have high LGD values, unlike what is the case for the PD values.

4.3.2 Simulation

We simulate conditional PD values from (2), by using simulated values for X. The PD and R values used in this calculation are obtained from the data set. For each simulated X value, the conditional PD is calculated for the data set's 109045 loans. We want to simulate loss distributions with different tail weights. This is achieved by simulating the X values by drawing from different probability distributions. Figure 3 shows the probability density function of the distributions that will be used to simulate the X values. The standard normal distribution is used as a baseline. The Cauchy distribution is chosen as it provides different tails weights by changing the scale parameter. The scale parameters 0.5, 1, 1.5, 2 and 2.5 are used.

The conditional probabilities of default for the Cauchy distribution is calculated using a modified version of (2):

$$PD(X) = F_C\left(\frac{F_C^{-1}(PD) - X\sqrt{R}}{\sqrt{1-R}}\right),$$

where F_C is the cumulative distribution function of the Cauchy(0,1) distribution.

To simulate losses, loan j is considered defaulted if a uniformly distributed random num-

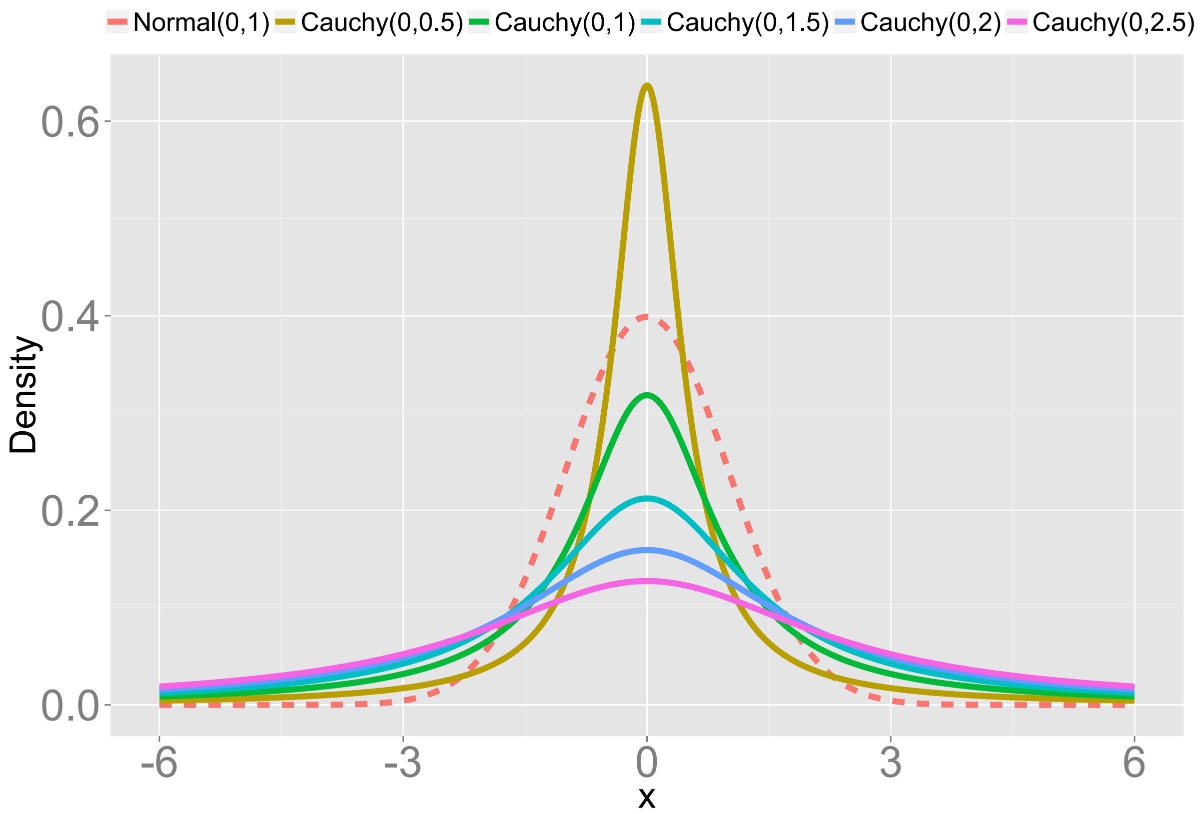

Figure 3: Probability density function for the Cauchy distribution with five different scale parameters ranging from 0.5 to 2.5. The probability density function of the standard normal distribution is also included, with dotted lines.

ber $U_j \in [0, 1]$ is smaller than or equal to the simulated conditional PD value. For the defaulted loans, the conditional PDs are multiplied with the associated LGD and EAD values from the data set to obtain the money amount lost:

$$L(X) = \sum_{j=1}^{J} \mathbf{1}_{\{PD_j(X) \geq U_j\}} \cdot LGD_j \cdot EAD_j, \tag{8}$$

where $\mathbf{1}_{\{\}}$ is the indicator function and $J = 109045$.

The simulation of loss as shown in (8) is then repeated N times, so that the N different loss values constitute a representation of the assumed loss distribution. Following Yamai and Yoshiba [24] we take the VaR estimator at confidence level α as the $(N \cdot (\alpha - 1) + 1)$th largest loss value, and the ES estimator as the mean of the $(N \cdot (\alpha - 1) + 1)$ largest loss values.

We simulate M sets of N different loss values, to obtain better estimates for VaR and ES, namely the means of the two sets of M different estimators. This way we can also study the standard deviations of the two final estimates. The whole procedure is carried out for five different confidence levels.

The size of the data set makes this process quite time consuming for big values of N and M. The simulation code is written in R [25]. Multicore computer processing was enabled to speed up the process, using the packages `foreach`, `parallel` and `doParallel`.

4.3.3 Results

Figure 4 shows three different simulated loss distributions, each consisting of N loss values simulated using (8). The systematic risk factor X has been drawn from, respectively, the Normal(0,1), the Cauchy(0,1) and the Cauchy(0,2.5) probability distributions. We see that the Cauchy distributed risk factors result in distributions with much heavier tails than for the normal distributed risk factor. For the Cauchy distributed risk factors, the scale parameter does not seem to have a big impact on the tail length. The bigger scale parameter does however result in a noticeably heavier tail. Table 2 shows the means and relative standard de-

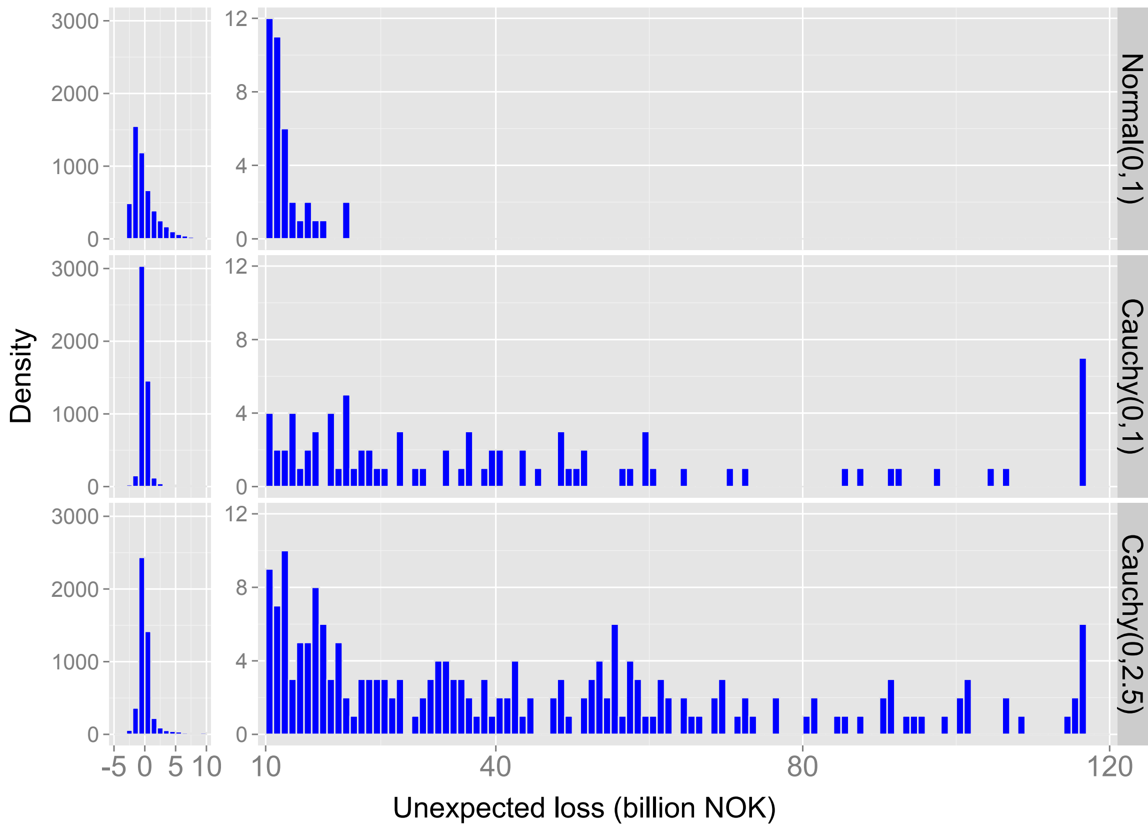

Figure 4: Simulated loss distributions given different probability distributions for the systematic risk factor. From the top: Standard normal distribution, Cauchy distribution with scale parameter 1 and scale parameter 2.5. All three distributions consist of $N = 5000$ simulated loss values. Two different y-axis are used to make the unlikely tail events more visible.

viations of $M = 100$ VaR and ES estimates produced using the simulation method described above, in addition to the percentage difference between the corresponding results for ES and VaR. Each estimate is calculated using $N = 5000$ simulated loss values, and is calculated at five different confidence levels for each of the six loss distributions.

As one would expect, we see from Table 2 that the VaR and ES estimates are most dependent on the confidence level for the most heavy-tailed loss distributions. This applies especially to the ES estimates, as they are affected by the whole tail regardless of confidence level. The gap between the VaR and ES values is decreasing for higher confidence levels, as this causes VaR to take into account a greater part of the distribution function. Since an increase in scale parameter for the Cauchy distribution results in a noticeably heavier tail, but not a longer tail, the difference between the two risk measures is actually decreasing when increasing the scale parameter. Note that the largest scale parameters used do result in loss distributions that are probably more heavy-tailed than what one would realistically expect.

When it comes to the estimates' relative standard deviation, the results are more varying. Only the standard normal loss distribution leads to increasing relative SD for higher confidence levels. This is also the only loss distribution that results in the ES estimates having the highest relative SD for all five confidence levels. This implies that the losses beyond the VaR quantile level varies more than the quantile level itself, meaning that the ES estimates require a larger sample size to ensure the same level of accuracy as the VaR estimates. This is not the case where we have negative values in Table 2c, as this means the VaR estimate has the highest relative SD. This is the case for all the Cauchy loss distributions when the confidence level is 99 % or higher. The reason being that these loss distributions are so long-tailed that the largest simulated losses do not vary much between each simulation set.

The most notable result from Table 2 is that the difference between value at risk and expected shortfall is highly dependent on the loss distribution. As mentioned in Section 4.3.1, the closest equivalent to the 99.9 % VaR is a 99.742 % ES. For the 99 % VaR, the closest equivalent is 97.465 %. Thus, it does not make sense to compare the mean and relative standard deviation of the 99.9 % VaR with the corresponding numbers for the 99.9 % ES, and so on. Considering this, ES has the lowest relative SD also for the standard normal loss distribution.

Appendix A shows how the number of simulations, N, affects the relative standard deviations of the VaR and ES estimates. A graphic representation of a selection of these results are shown in Appendix B. As one would expect, the relative SD decreases when you increase the number of simulations. The size of this reduction appears to be about the same for both the

VaR	Mean (billion NOK)					Relative SD				
Dist \CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
Normal(0,1)	7.150	8.844	11.266	14.528	17.367	0.020	0.025	0.038	0.046	0.083
Cauchy(0,0.5)	3.407	4.232	9.969	49.370	87.391	0.010	0.047	0.161	0.203	0.168
Cauchy(0,1)	4.216	7.916	28.724	78.021	111.764	0.033	0.094	0.192	0.138	0.071
Cauchy(0,1.5)	5.703	12.699	47.093	95.243	116.444	0.061	0.123	0.150	0.100	0.029
Cauchy(0,2)	7.627	20.165	59.256	104.599	118.500	0.069	0.162	0.108	0.079	0.010
Cauchy(0,2.5)	9.804	28.666	68.016	111.612	119.038	0.067	0.132	0.077	0.051	0.004

(a) Value at Risk

ES	Mean (billion NOK)					Relative SD				
Dist \CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
Normal(0,1)	9.764	11.619	14.163	17.504	20.256	0.024	0.030	0.041	0.066	0.092
Cauchy(0,0.5)	12.017	20.228	40.833	80.496	105.006	0.103	0.120	0.136	0.115	0.093
Cauchy(0,1)	20.506	35.335	66.161	102.409	116.822	0.092	0.100	0.096	0.070	0.026
Cauchy(0,1.5)	28.172	47.888	80.862	110.684	118.442	0.081	0.085	0.070	0.040	0.009
Cauchy(0,2)	34.977	57.992	89.703	115.064	119.077	0.084	0.083	0.061	0.025	0.003
Cauchy(0,2.5)	41.637	66.734	96.889	117.241	119.248	0.068	0.065	0.050	0.017	0.001

(b) Expected Shortfall

100 · (ES-VaR)/VaR	Mean					Relative SD				
Dist \CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
Normal(0,1)	36.6	31.4	25.7	20.5	16.6	19.5	21.7	9.8	43.4	11.4
Cauchy(0,0.5)	252.7	378.0	309.6	63.0	20.2	987.2	154.9	-15.9	-43.2	-44.8
Cauchy(0,1)	386.3	346.4	130.3	31.3	4.5	176.8	6.9	-50.1	-49.7	-63.7
Cauchy(0,1.5)	394.0	277.1	71.7	16.2	1.7	34.1	-30.8	-53.7	-60.3	-68.7
Cauchy(0,2)	358.6	187.6	51.4	10.0	0.5	21.9	-48.5	-43.2	-68.7	-67.1
Cauchy(0,2.5)	324.7	132.8	42.4	5.0	0.2	1.1	-51.1	-34.5	-67.9	-69.5

(c) Percentage difference

Table 2: The mean and relative standard deviation of $M = 100$ VaR and ES estimates, for different confidence levels (CL). Each estimate is calculated using $N = 5000$ simulated loss values, simulated using the probability distribution indicated in the leftmost column for the systematic risk factor. The percentage difference between the corresponding ES and VaR results is also shown.

VaR and ES estimates, even if the levels differ. The relative reduction thus being largest for the estimates with the smallest relative SD. However, do still keep in mind that ES must have a smaller confidence level than VaR if the two risk measures are to result in the same capital charge.

To summarize, this section shows that the difference between value at risk and expected shortfall is highly dependent on the loss distribution.

4.4 Backtesting

Backtesting is a method used for model validation, where statistical procedures are used to compare actual losses to former risk measure forecasts. We will compare backtesting for value at risk and expected shortfall, with respect to both theoretical properties and practical implementation.

4.4.1 Elicitability

Gneiting [17] proved in 2010 that expected shortfall is not elicitable, as opposed to value at risk. This discovery led many to erroneously conclude that ES would not be backtestable, see for instance [18]. Elicitability is defined as follows [19]:

Definition 2 (Elicitability)**.** *A statistic $\phi(Y)$ of a random variable Y is said to be elicitable if it minimizes the expected value of a scoring function S:*

$$\phi(Y) = \arg\min_x E[S(x, Y)].$$

If you want to compare different forecasting procedures, this is typically done by using a scoring function (error measure), such as the absolute error or the squared error, which is averaged over forecast cases. Thus, the performance criterion takes the form [17]

$$\bar{S} = \frac{1}{n} \sum_{i=1}^{n} S(x_i, y_i), \tag{9}$$

where x_i are point forecasts, the y_i are the corresponding realizations and S is the scoring function. Most scoring functions are negatively oriented, that is, the smaller, the better. Thus, we favour the forecasting procedure that minimizes (9).

In simple terms, Definition 2 says that a statistic is elicitable if there exists a scoring function S that makes this statistic the best forecasting procedure according to (9). The mean and the median represent popular examples, minimizing the mean square and absolute error, respectively. The q^{th} quantile, hence VaR, is elicitable with the scoring function $S(x, y) = (\mathbf{1}_{\{x \geq y\}} - q)(x - y)$, where $\mathbf{1}_{\{\}}$ is the indicator function [19].

It turns out that even if ES is not elicitable, it is still '2nd order' elicitable in the following sense [20]:

Definition 3 (Conditional Elicitability)**.** *A statistic $\phi(Y)$ of a random variable Y is called conditionally elicitable if there exist two statistics $\tilde{\pi}(Y)$ and $\pi(Y)$ such that*

$$\phi(Y) = \pi(Y, \tilde{\pi}(Y)),$$

where $\tilde{\pi}(Y)$ *is elicitable and* $\pi(Y)$ *is such that* $\pi(Y,c)$ *is elicitable for all* $c \in \mathbb{R}$.

Conditional elicitability is a helpful concept for the forecasting of risk measures which are not elicitable. Due to the elicitability of $\tilde{\pi}(Y)$ we can first forecast $\tilde{\pi}(Y)$ and then, in a second step, regard this result as fixed and forecast $\pi(Y,c)$ due to the elicitability of $\pi(Y)$. With regard to backtesting and forecast comparison, conditional elicitability offers a way of splitting up a forecast method into two component methods and separately backtesting and comparing their forecast performances [20]. This applies to ES, as it is simply a mean of quantiles, and both the quantiles and the mean are elicitable.

4.4.2 Practical implementation

A popular backtesting method for value at risk is based on the following *violation process*:

$$I_t(q) = \mathbf{1}_{\{L(t) > \mathrm{VaR}_q(L(t))\}},$$

where $\mathbf{1}_{\{\}}$ is the indicator function and t denotes the time period.

Christoffersen [21] shows that VaR forecasts are valid if and only if the violation process $I_t(q)$ satisfies the unconditional coverage hypothesis: $E[I_t(q)] = 1-q$ in addition to $I_t(q)$ and $I_s(q)$ being independent for $s \neq q$. Under these two conditions, the violations are independent and identically distributed Bernoulli random variables with success probability $1-q$. Hence, the number of violations has a Binomial distribution. The unconditional coverage hypothesis can be tested by comparing the fraction of violations to the VaR confidence level, using a standard likelihood ratio test.

Backtesting ES does not have to be more complicated than backtesting VaR. Tasche et al. [20] proposes a backtesting method for ES that is as simple as the VaR violation method, based on the following approximation:

$$\begin{aligned}\mathrm{ES}_q(L) &= \frac{1}{1-q}\int_{u=q}^{1} \mathrm{VaR}_u(L)\,du \\ &\approx \frac{1}{4}\left[\mathrm{VaR}_q(L) + \mathrm{VaR}_{0.75q+0.25}(L) + \mathrm{VaR}_{0.5q+0.5}(L) + \mathrm{VaR}_{0.25q+0.75}(L)\right]. \qquad (10)\end{aligned}$$

If the four different VaR values in (10) are successfully backtested, then also the estimate of

$\text{ES}_q(L)$ can be considered reliable subject to careful manual inspection of the observations exceeding $\text{VaR}_{0.25q+0.75}(L)$. These tail observations must at any rate be manually inspected in order to separate data outliers from genuine fair tail observations.

In practise, backtesting of credit risk models can be quite problematic. The infrequent nature of default events makes it difficult to collect enough relevant data, especially for the tail of the loss distribution. The long time horizons further complicate the data collection. The backtesting approach based on (10) is attractive not only for its simplicity but also because it illustrates the fact that ES backtesting requires more data than the VaR backtesting, since loss beyond the VaR level is infrequent, thus the average of them is hard to estimate accurately [11]. For market risk, the Basel Committee uses a similar backtesting approach for a 97.5 % ES, which is based on testing VaR violations for the 97.5 % and 99 % confidence levels [2].

Acerbi and Szekely [19] have recently argued that elicitability has to do with model *selection* and not with model *testing*, and is therefore irrelevant for the choice of a regulatory risk standard. They show that expected shortfall is directly backtestable, by introducing three model-free, nonparametric backtesting methods for ES. These tests generally require more storage of information than typical VaR tests, but introduce no conceptual limitations or computational difficulties of any sort. Compared to these test procedures, the simple backtesting method based on (10) has the advantage of not relying on Monte Carlo simulation for the statistical test [20].

4.5 Parameter Sensitivity

We will examine how the uncertainty of the banks' parameter estimates affects the output from the Basel Committee's capital requirement function, using both value at risk and expected shortfall. This is carried out by simulating LGD and PD values. The estimation uncertainty of these two parameters is represented by the relative standard deviation of the probability distributions they are sampled from.

4.5.1 Simulating LGD Values

The simulation method used for the LGD values is based on a model for recovery rates ($1 - LGD$) developed by Jon Frye [22]. Frye's model adapts some of Michael Gordy's work, and lets the recovery rate r depend on the systematic risk factor X:

$$r = \mu + \sigma h X + \sigma\sqrt{1-h^2}Z, \tag{11}$$

where h is the correlation between the recovery rate r and the systematic risk factor X, and Z is a standard normal variable independent of X.

Following a proposal from Schönbucher [2], we apply a logistic transformation $F(Y) = \frac{\exp(Y)}{1+\exp(Y)}$ on (11), to limit the r values (and thus also the LGD values) to the interval $[0,1]$.

As for the capital requirement function (5) we use $x = -\Phi^{-1}(0,999) = \Phi^{-1}(0,001)$. This gives the following distribution for the LGD values:

$$\begin{aligned} r &\sim N\left(\mu + \sigma h \Phi^{-1}(0,001), \sigma^2\left(1-h^2\right)\right), \\ \widehat{LGD} &= 1 - F(r) = 1 - \frac{\exp(r)}{1+\exp(r)}. \end{aligned} \tag{12}$$

4.5.2 Simulating PD Values

Over a time period, a firm either meets the loan terms or defaults. This makes it natural to model the number of default occurrences m by applying a binomial distribution:

$$m \sim Bin(n, PD),$$

where PD is chosen as the bank's estimated value for the unconditional probability of default, and n is the number of simulations.

After completing the n simulations, the probability of default is estimated as

$$\widehat{PD} = \frac{m}{n}.$$

The simulated $\widehat{PD}$ values will thus be centered around the bank's estimate of PD, and the deviations from this value represents the uncertainty in the bank's PD estimate. The variance of $\widehat{PD}$ is inversely proportional to the number of simulations:

$$\text{Var}\left(\widehat{PD}\right) = \text{Var}\left(\frac{m}{n}\right) = \frac{1}{n^2}\text{Var}(m) = \frac{PD(1-PD)}{n}.$$

[2] See [23], pages 147-150.

4.5.3 Calculation of Parameter Sensitivity

To simulate $\widehat{LGD}$ values representing different degrees of estimation uncertainty, five different values are used for σ in (12). The chosen values range from 0.05 to 0.45, with increments of 0.1. h is kept constant at 0.2. The $\widehat{PD}$ values are simulated in a similar way, where five different values of n represent varying degrees of estimation uncertainty. For each of the five σ values there are simulated N different $\widehat{LGD}$ values, and N different $\widehat{PD}$ values are simulated for each of the five n values. These five values for n are chosen so that the five different series of $\widehat{PD}$ values have the same relative standard deviations as the corresponding five different series of $\widehat{LGD}$ values . This is achieved by selecting n values that satisfy the following equation:

$$\frac{\sqrt{\text{Var}\left(\widehat{PD}\right)}}{PD} = \sqrt{\frac{1-PD}{nPD}} = \sigma_{\widehat{LGD}\text{rel}} \Longrightarrow n = \frac{1-PD}{\left(\sigma_{\widehat{LGD}\text{rel}}\right)^2 PD},$$

where $\sigma_{\widehat{LGD}\text{rel}}$ is the relative standard deviation of a $\widehat{LGD}$ value series, and PD is the expected value of $\widehat{PD}$. Figure 5 shows how the distributions of the simulated values depends on σ and n.

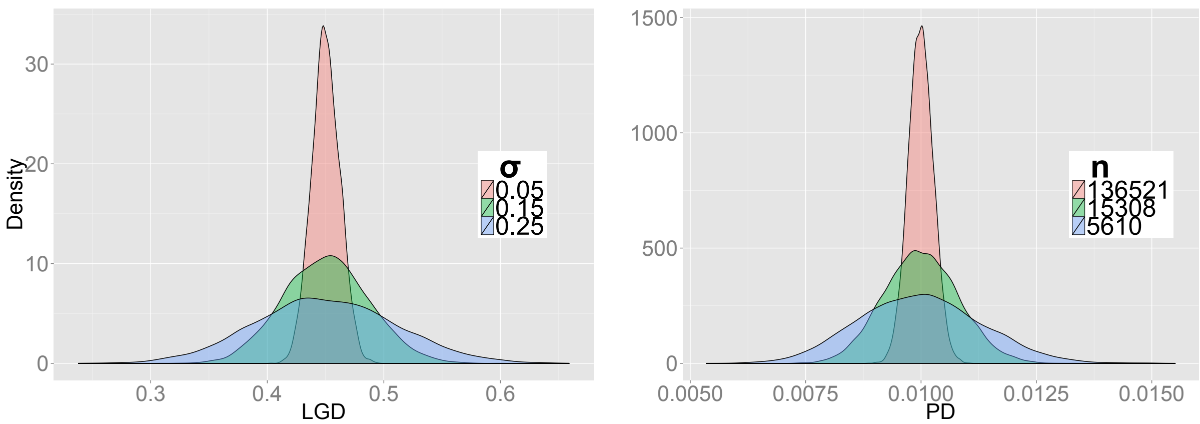

Figure 5: The distributions of 10000 simulated LGD and PD values, with expected values of respectively 0.45 and 0.01. The LGD values shown are simulated with three lowest σ values, and the PD values are simulated with the three n values that result in the same relative standard deviations.

The simulated $\widehat{LGD}$ and $\widehat{PD}$ values are used pairwise to calculate the corresponding capital requirement values. The loan maturity is chosen to one year, so that the adjustment factor MA equals one. The correlation factor is chosen to be calculated as for loans to firms with annual revenue above 50 million euros. These choices result in the capital requirement function (5) taking the following form

$$\widehat{K} = \widehat{LGD} \cdot \Phi\left(\frac{\Phi^{-1}(\widehat{PD}) + \Phi^{-1}(0,999) \cdot \sqrt{0.24 - 0.12\left(\frac{1-e^{-50\cdot\widehat{PD}}}{1-e^{-50}}\right)}}{\sqrt{0.76 - 0.12\left(\frac{1-e^{-50\cdot\widehat{PD}}}{1-e^{-50}}\right)}} \right) - \widehat{PD} \cdot \widehat{LGD}. \tag{13}$$

The corresponding version of the expected shortfall capital requirement function (7) takes the form

$$\widehat{K} = \frac{\widehat{LGD}}{1-q} \cdot \Phi_2\left(\Phi^{-1}(\widehat{PD}), -\Phi^{-1}(q); \sqrt{0.24 - 0.12\left(\frac{1-e^{-50\cdot\widehat{PD}}}{1-e^{-50}}\right)} \right) - \widehat{PD} \cdot \widehat{LGD}. \tag{14}$$

The capital requirements (13) and (14) are calculated for all 25 possible combinations of σ and n. At last, the relative standard deviation of the calculated capital requirements for each of these combinations are computed:

$$\sigma_{K\text{rel}} = \sqrt{\frac{\sum_{i=1}^{N}\left(K_i - \overline{K}\right)^2}{N-1}} \Big/ \overline{K}.$$

Since (13) and (14) are proportional to $\widehat{LGD}$, it is not interesting to vary the expected value of $\widehat{LGD}$, as it will only result in a linear scaling of the capital requirement's variation. The expected value of the simulated $\widehat{LGD}$ values is set to 0.45 for all the different σ values. To achieve this, for each different σ value in (12), we must calculate a μ which gives this desired expectation value. However, several different expectation values will be used for the $\widehat{PD}$ simulations, to see how this impacts the resulting standard deviations of the calculated capital requirements.

4.5.4 Results

Figure 6 shows the relative standard deviation of the simulated capital requirement for 99.9 % value at risk, for different expectation values for PD. For each of the different expectation values, there are 25 different values for the standard deviation of the capital requirement, which corresponds to different combinations of the five levels of uncertainty for both the PD and LGD simulations. As one would expect, the capital requirement's uncertainty increases with increasing parameter uncertainty. What we are interested in, is which of the two parameters' uncertainty that affects the capital requirement the most.

We see from Figure 6 that the parameter uncertainties have different impact depending

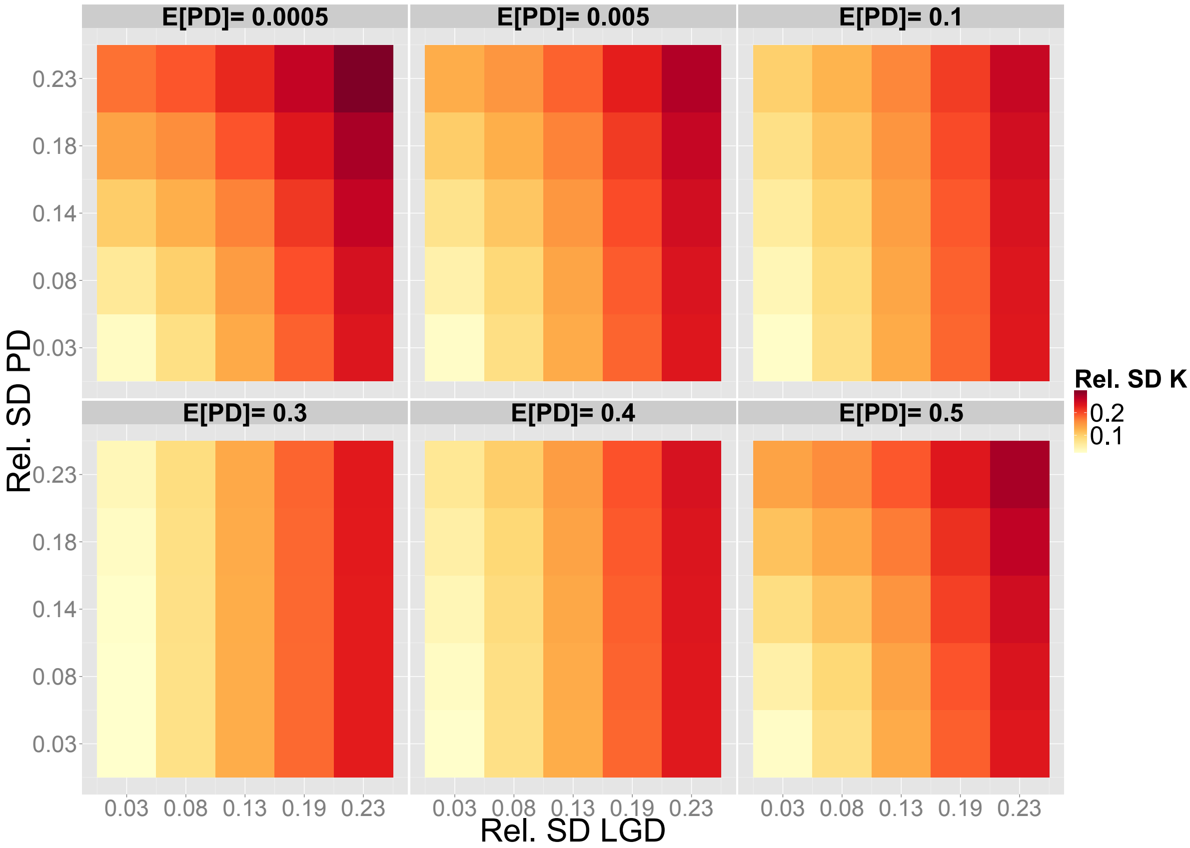

Figure 6: Relative standard deviation of the 99.9 % VaR capital requirement (13), given the relative standard deviation for PD and LGD. Calculated for six different expectation values of PD, with the expected value of LGD equal to 0.45. Using $N = 10000$ simulations for each calculation.

on the the expected value of the simulated PD values. When the expected probability of default is close to 0.3, the capital requirement uncertainty is almost only influenced by the uncertainty of the simulated LGD values. The uncertainty of the PD values plays a greater role when the expectation of PD is either small or above 0.4. But even for $E[PD] = 0.0005$ the LGD uncertainty is most influential, as we see that the rightmost column is a slightly darker red than the upper row. However, when the expected value of the probability of default exceeds 0.7, the PD uncertainty is extremely influential, and the standard deviation of the capital requirement increases significantly. This can be seen in Figure 7, where the expected value of the simulated PD values are 0.8. Note that the colors in this figure correspond to larger relative standard deviations than in Figure 6.

Looking at (13) it is clear that $e^{-50\cdot\widehat{PD}}$ is the part of the capital requirement function that explains the influence of the PD values' uncertainty when the expectation of PD is small. For expectation values of PD close to zero, this part of the function is sensitive to very small

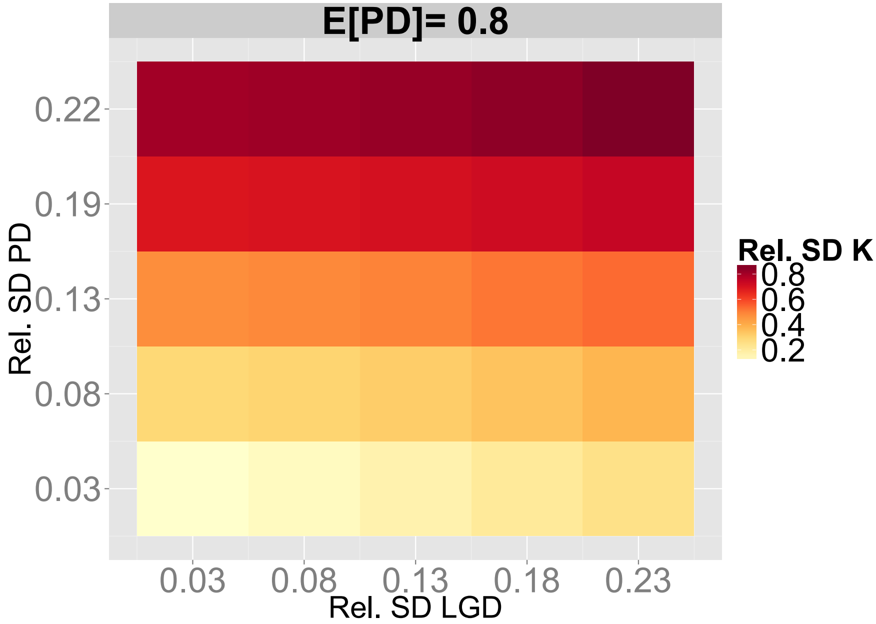

Figure 7: Relative standard deviation of 99.9 % VaR capital requirement (13), given the relative standard deviation for PD and LGD. Calculated for $E[PD] = 0.8$ and $E[LGD] = 0.45$, using $N = 10000$ simulations.

changes in the PD value. This sensitivity gradually becomes smaller for larger PD values, and for PD values larger than 0.1 this part of the function will remain approximately constant. The function part $\Phi^{-1}(\widehat{PD})$ is particularly sensitive when the expected value of PD is close to zero or one.

The major impact on the capital requirement's uncertainty for large expectation values of PD is due to the last term in (13). For large PD values, the value of the last term becomes large enough so that its variation affects the variation of the whole function.

When the same simulation method was carried out using expected shortfall, the confidence level was chosen to 99.742 %, as it was shown in Section 4.3.1 that this confidence level results in the capital requirement closest to the 99.9 % VaR. For large expected values of PD, there were virtually no difference in the capital requirement's uncertainty between the ES and VaR approach. For smaller expected values of PD the ES approach resulted in reduced uncertainty, as shown in Figure 8. We see that the relative reduction is largest when the LGD uncertainty is low and the PD uncertainty is high. This reduction is however only a few percent, so there is not that much of a difference between the two approaches regarding the parameter sensitivity.

Because the relative reduction of the capital requirement's relative standard deviation is largest for the combination of the lowest LGD uncertainty and the highest PD uncertainty,

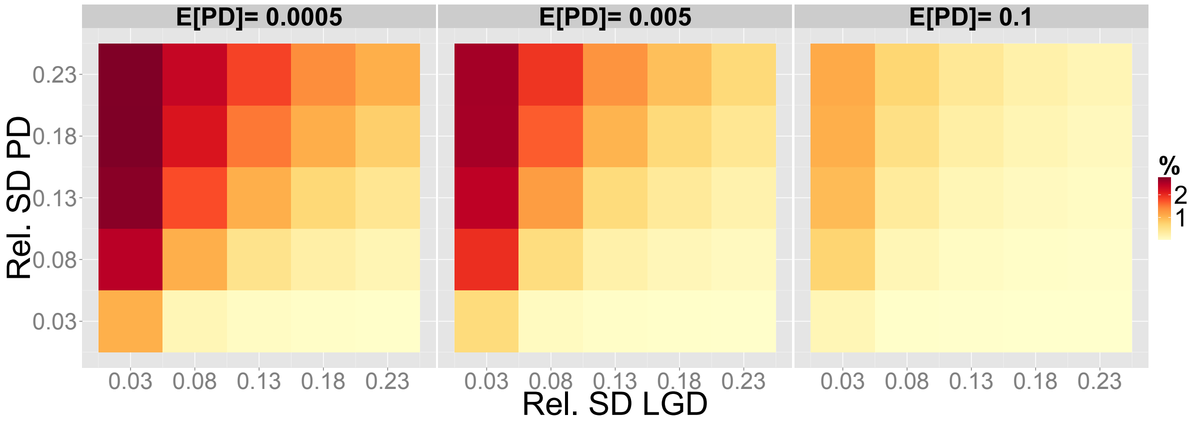

Figure 8: Percentage reduction in the relative standard deviation of the capital requirement by switching from 99.9 % VaR (13) to 99.742 % ES (14), given the relative standard deviation for PD and LGD. Calculated for three different expectation values of PD, with the expected value of LGD equal to 0.45. Using $N = 10000$ simulations for each calculation.

we decide to calculate the reduction for this case using different confidence levels for the expected shortfall. Figure 9 shows the results for confidence levels ranging from 99.4 % to 99.9 %. We see that the relative uncertainty reduction is largest for the combination of high confidence levels and low PDs. The combination of large PDs and smaller confidence levels also stands out. We see that switching to ES also increases the uncertainty in a few cases, especially for large PD values at the 99.9 % confidence level.

As both the VaR version (13) and the ES version (14) of the capital requirement function are based on the same assumptions and models, the only distinction between the two versions is the risk measure. The results in this section thus show how a credit model's parameter sensitivity can depend on the chosen risk measure. In Section 4.2 we saw from Figure 1 that the 99.742 % ES version resulted in a higher capital charge than the 99.9 % VaR version for PD values below 0.21, with considerable increasing relative difference for PD values below 0.01. Figure 8 shows that the lowest PD values also cause the most notable difference between the VaR and ES version when it comes to the relative standard deviation of the capital requirement. This is of course no coincidence. As the difference between the ES and VaR version increases for the small PD values, the ES version will result in a smaller change in capital charge than the VaR version for the same change in these PD values. The ES version thus makes the capital charge depend on these values in a more stable manner, thereby reducing the relative standard deviation.

There could be both advantages and disadvantages with a capital requirement function that is less sensitive to the PD parameter in the lowest end of the scale. On the plus side,

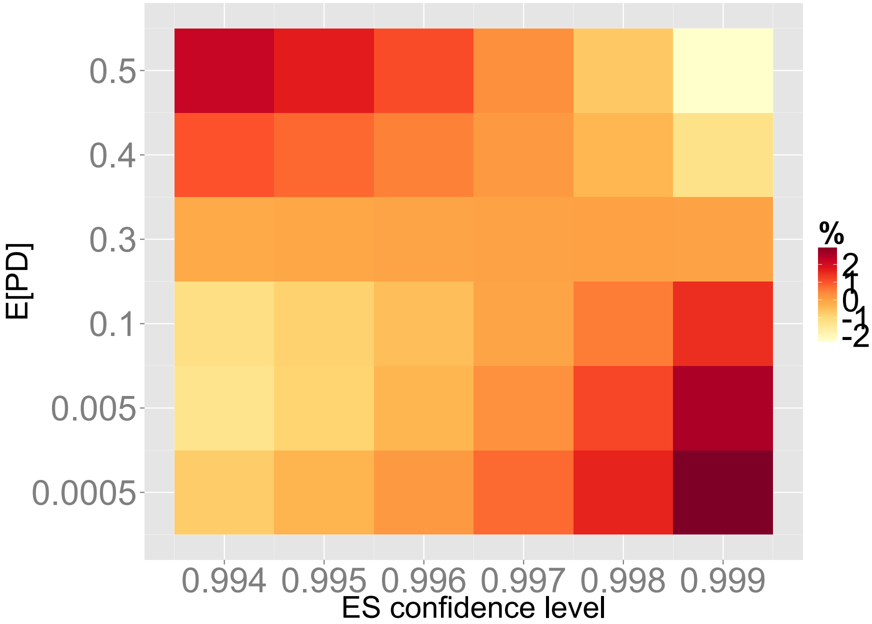

Figure 9: Percentage reduction in the capital requirement's relative standard deviation by switching from 99.9 % VaR to ES with different confidence levels. The results are shown for different expectation values for the probability of default. The relative standard deviation of the simulated PD and LGD values is set constant to 23 % and 3 % respectively. The expected value of LGD is equal to 0.45. Using $N = 10000$ simulations for each calculation.

one could argue that this to some degree reduces the banks' incentive to estimate artificially low PD values. At the same time this might be viewed as counterproductive, since the fundamental idea behind the IRB approach is a more risk sensitive capital charge. In case of a change of risk measure from VaR to ES, one would possibly also make some changes to the model assumptions, particularly the shape of the loss distribution function. A change to the more tail risk sensitive ES is essential for more heavy-tailed loss functions.

5 Conclusion

The Basel Committee's minimum capital requirement function for banks' credit risk is based on the risk measure value at risk (VaR). The paper performs a statistical and economic analysis of the consequences of replacing VaR with Expected Shortfall (ES), a switch that has already been set in motion for market risk. This analysis uses both theoretical simulations and real data from a Norwegian savings bank group's corporate portfolio.

By correctly calibrating the ES confidence level, it will produce approximately the same capital requirement for credit risk as with VaR, where the largest difference occurs for loans with low default probability. A switch from VaR to ES will involve some clear conceptual improvements, primarily a better ability to accurately capture tail risk. ES is also sub-additive in general, unlike VaR, so that it always reflects the positive effect of diversification. There has been some uncertainty regarding the backtesting abilities of ES, but we show that backtesting of ES does not have to be more complicated than backtesting VaR. The parameter sensitivity and estimation stability of ES have also been examined, and appear to be similar as for VaR, if not slightly less sensitive and more stable.

This paper shows that the difference between ES and VaR is highly dependent on the assumed loss distribution. Since ES considers the entire loss distribution, it is more suitable for credit risk models with more heavy-tailed loss distributions than the normal distribution. For such distributions, we show that the estimation stability of ES is clearly better than for VaR.

The advantages of switching to ES must be weighed against costs and challenges associated with a transition to this risk measure, especially concerning practical implementation. However, as this risk measure switch has already been set in motion for market risk, banks are going to have practical experience with ES before this switch potentially also happens for credit risk. In addition, ES is after all based on VaR, so we are talking about adjusting the existing system, not creating a new system from scratch. Taking all this into consideration, the conclusion is that the findings of the paper support a switch from VaR to ES for credit risk modelling.

A

VaR	Normal(0,1)					Cauchy(0,0.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.041	0.056	0.070	0.088	0.117	0.020	0.105	0.397	0.532	0.430
2500	0.030	0.038	0.053	0.067	0.090	0.014	0.060	0.320	0.321	0.255
5000	0.020	0.025	0.038	0.046	0.083	0.010	0.047	0.161	0.203	0.168
10000	0.013	0.019	0.025	0.037	0.043	0.007	0.038	0.120	0.168	0.125
	Cauchy(0,1)					Cauchy(0,1.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.078	0.259	0.396	0.292	0.228	0.133	0.376	0.373	0.243	0.170
2500	0.051	0.164	0.327	0.173	0.149	0.083	0.170	0.192	0.139	0.084
5000	0.033	0.094	0.192	0.138	0.071	0.061	0.123	0.150	0.100	0.029
10000	0.025	0.062	0.139	0.104	0.036	0.039	0.088	0.109	0.077	0.019
	Cauchy(0,2)					Cauchy(0,2.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.147	0.274	0.238	0.193	0.117	0.176	0.288	0.192	0.130	0.089
2500	0.100	0.201	0.144	0.113	0.058	0.098	0.172	0.120	0.077	0.029
5000	0.069	0.162	0.108	0.079	0.010	0.067	0.132	0.077	0.051	0.004
10000	0.046	0.104	0.055	0.057	0.005	0.059	0.100	0.055	0.035	0.001

(a) Value at Risk

ES	Normal(0,1)					Cauchy(0,0.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.049	0.061	0.081	0.112	0.140	0.266	0.310	0.348	0.325	0.293
2500	0.037	0.046	0.062	0.088	0.116	0.172	0.200	0.219	0.190	0.162
5000	0.024	0.030	0.041	0.066	0.092	0.103	0.120	0.136	0.115	0.093
10000	0.017	0.022	0.029	0.041	0.057	0.089	0.104	0.114	0.087	0.060
	Cauchy(0,1)					Cauchy(0,1.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.204	0.218	0.205	0.166	0.138	0.215	0.223	0.192	0.137	0.108
2500	0.152	0.164	0.147	0.109	0.075	0.104	0.110	0.089	0.069	0.037
5000	0.092	0.100	0.096	0.070	0.026	0.081	0.085	0.070	0.040	0.009
10000	0.066	0.072	0.069	0.047	0.012	0.059	0.064	0.052	0.032	0.005
	Cauchy(0,2)					Cauchy(0,2.5)				
N\CL	0.95	0.975	0.99	0.997	0.999	0.95	0.975	0.99	0.997	0.999
1000	0.157	0.161	0.133	0.093	0.061	0.156	0.145	0.110	0.071	0.049
2500	0.111	0.112	0.088	0.053	0.024	0.092	0.087	0.065	0.032	0.013
5000	0.084	0.083	0.061	0.025	0.003	0.068	0.065	0.050	0.017	0.001
10000	0.055	0.056	0.043	0.019	0.002	0.052	0.047	0.037	0.008	0.000

(b) Expected Shortfall

Table 3: The relative standard deviations of $M = 100$ VaR and ES estimates, for different confidence levels (CL) and different number of simulated loss values, N, used for each estimate. The loss values are simulated using the probability distribution indicated in the table headers for the systematic risk factor.

B

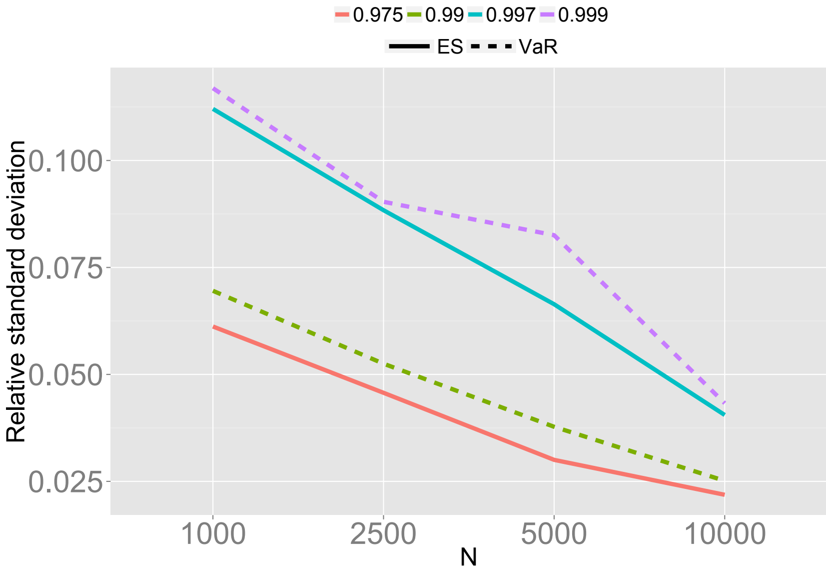

(a) Normal(0,1)

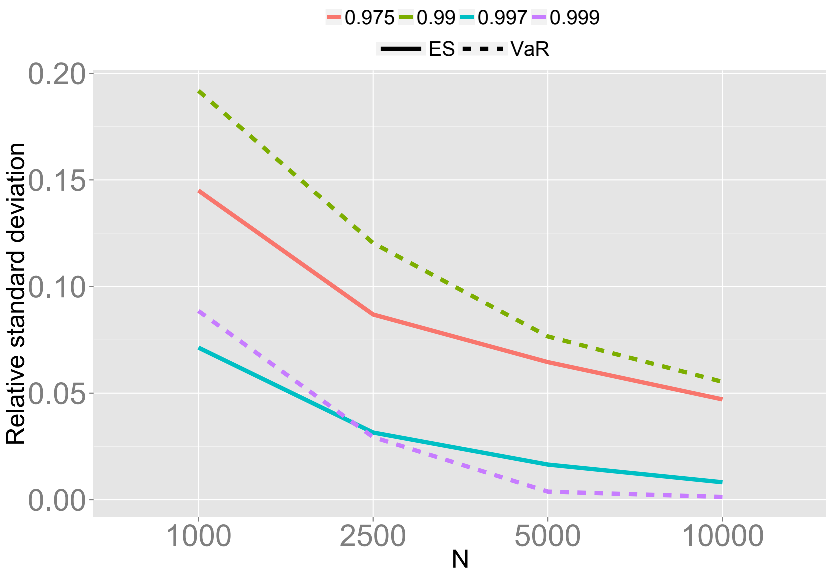

(b) Cauchy(0,2.5)

Figure 10: Relative standard deviation of the VaR (dotted lines) and ES (solid lines) estimates, for N simulations. 99.9 % and 99 % confidence levels are used for VaR, while 99.7 % and 97.5 % confidence levels are used for ES. The estimates are calculated from simulated loss values, from both a standard normal distribution and a Cauchy(0,2.5) distribution.

C References

[1] Bank for International Settlements, "International Convergence of Capital Measurement and Capital Standards." http://www.bis.org/publ/bcbs107.pdf, 2004 [Retrieved 08.23.15].

[2] Bank for International Settlements, "Minimum capital requirements for market risk." http://www.bis.org/bcbs/publ/d352.pdf, 2016 [Retrieved 02.08.16].

[3] Bank for International Settlements, "Fundamental review of the trading book." http://www.bis.org/publ/bcbs219.pdf, 2012 [Retrieved 02.08.16].

[4] Bank for International Settlements, "A brief history of the Basel Committee." http://www.bis.org/bcbs/history.pdf, 2014 [Retrieved 08.23.15].

[5] Bank for International Settlements, "Credit risk modelling: Current practices and applications." http://www.bis.org/publ/bcbs49.pdf, 1999 [Retrieved 04.04.16].

[6] Bank for International Settlements, "An Explanatory Note on the Basel II IRB Risk Weight Functions." http://www.bis.org/bcbs/irbriskweight.pdf, 2005 [Retrieved 10.14.15].

[7] M. B. Gordy, "A risk-factor model foundation for ratings-based bank capital rules," *Journal of financial intermediation*, vol. 12, no. 3, pp. 199–232, 2003.

[8] O. Vasicek, "The distribution of loan portfolio value," *Risk*, vol. 15, no. 12, pp. 160–162, 2002.

[9] Bank for International Settlements, "Instructions for Basel III monitoring." http://www.bis.org/bcbs/qis/biiiimplmoninstr_feb13.pdf, 2013 [Retrieved 08.23.15].

[10] P. Artzner, F. Delbaen, J. Eber, and D. Heath, "Thinking coherently," *RISK*, vol. 10, pp. 68–71, November 1997.

[11] Y. Yamai and T. Yoshiba, "On the Validity of Value-at-Risk: Comparative Analyses with Expected Shortfall," *Monetary and economic studies*, vol. 20, no. 1, pp. 57–85, 2002.

[12] J. Daníelsson, B. N. Jorgensen, G. Samorodnitsky, M. Sarma, and C. G. de Vries, "Fat tails, var and subadditivity," *Journal of econometrics*, vol. 172, no. 2, pp. 283–291, 2013.

[13] C. Acerbi and D. Tasche, "Expected shortfall: a natural coherent alternative to value at risk," *Economic notes*, vol. 31, no. 2, pp. 379–388, 2002.

[14] P. Koch-Medina and C. Munari, "Unexpected shortfalls of Expected Shortfall: Extreme default profiles and regulary arbitrage," *Journal of Banking & Finance*, vol. 62, pp. 141–151, 2016.

[15] M. Hibbeln, *Risk Management in Credit Portfolios: Concentration Risk and Basel II.* Springer Science & Business Media, 2010.

[16] Bank for International Settlements, "Reducing variation in credit risk-weighted assets - constraints on the use of internal model approaches." `http://www.bis.org/bcbs/publ/d362.pdf`, 2016 [Retrieved 03.31.16].

[17] T. Gneiting, "Making and evaluating point forecasts," *Journal of the American Statistical Association*, vol. 106, no. 494, pp. 746–762, 2011.

[18] L. Carver, "Mooted var substitute cannot be back-tested, says top quant," *Risk, March*, vol. 8, 2013.

[19] C. Acerbi and B. Szekely, "Backtesting Expected Shortfall." `https://www.msci.com/resources/research/articles/2014/Research_Insight_Backtesting_Expected_Shortfall_December_2014.pdf`, 2014 [Retrieved 05.05.16].

[20] S. Emmer, M. Kratz, and D. Tasche, "What is the best risk measure in practice? A comparison of standard measures." `https://arxiv.org/pdf/1312.1645.pdf`, 2015 [Retrieved 05.05.16].

[21] P. F. Christoffersen, *Elements of financial risk management.* Academic Press, 2012.

[22] J. Frye, "Depressing recoveries," *RISK-LONDON-RISK MAGAZINE LIMITED-*, vol. 13, no. 11, pp. 108–111, 2000.

[23] P. J. Schönbucher, *Credit derivatives pricing models: models, pricing and implementation.* John Wiley & Sons, 2003.

[24] Y. Yamai and T. Yoshiba, “Comparative analyses of expected shortfall and value-at-risk: their estimation error, decomposition, and optimization,” *Monetary and economic studies*, vol. 20, no. 1, pp. 87–121, 2002.

[25] R Core Team, “R: A Language and Environment for Statistical Computing.” `http://www.R-project.org/`, 2014.