

Valuation and optimization of credit risk using a portfolio model

Per Sidén

2012-09-30

Abstract

In this paper I study a model for credit risk in a portfolio of sovereign bonds, based on (van der Hoorn, 2009). The model is based on historical credit rating changes and the joint distribution of the losses for different bonds is modeled with an assumption of an underlying multivariate Gaussian variable. Different risk measures for the portfolio are calculated using Monte Carlo simulations and the performance is improved by the use of importance sampling. I investigate different methods on how to improve the model and the estimation of the parameters of the model. I also develop methods to value the certainty of the risk measures based on a statistical view on the input data, which give clear indications that the small size of input data gives low accuracy for the risk measures. An attempt to write an algorithm to find the optimal portfolio with respect to the risk measures is also performed and the results are discussed.

Acknowledgements

I would like to thank my supervisor at Sveriges Riksbank, Tomas Edlund, for all his help and guidance throughout this thesis. I would also like to thank Jimmy Olsson at Lunds University and Kristofer Björnström at Sveriges Riksbank.

Contents

1	Introduction	5
1.1	Background	5
1.2	Purpose	7
1.3	Results	8
1.4	Outline	8
2	Credit risk	9
2.1	Credit ratings	9
2.2	Risk measures	12
2.2.1	Expected Loss	12
2.2.2	Unexpected Loss	13
2.2.3	Value-at-Risk	13
2.2.4	Expected Shortfall	15
3	Monte Carlo methods	15
3.1	Importance sampling	18
4	Modeling	21
4.1	One-dimensional model	21
4.2	Multi-dimensional model	23
4.3	Calculating risk measures	26
4.4	Estimating parameters	28
4.4.1	Migration matrix	28
4.4.2	Conditional credit losses	32
4.4.3	Correlations	34
5	Results and Validation	38
6	Optimization	46
6.1	Genetic algorithm	48
7	Discussion	51
7.1	Summary	51
7.2	Conclusions	51
7.3	Further Development	52
	References	53

1 Introduction

1.1 Background

This master thesis is written as a project for the central bank of Sweden, Sveriges Riksbank. One of the Riksbank's main tasks is to maintain financial stability in the Swedish economical system. Therefore, the Riksbank has a large reserve of foreign currency, which could be used for example for emergency loans in case a large Swedish bank cannot repay a loan in foreign currency, to save that bank from bankruptcy. Nowadays, the foreign exchange reserve consists mainly of sovereign bonds from different countries, which has the benefit compared to pure exchange of giving some return while still having their value closely linked to the value of the currency. However, owning bonds always come with some amount of credit risk and that is the topic of this thesis.

Credit risk is the risk that the issuer of a bond do not repay his debt or that the market value of the bond drops because investors believe that the probability of this has increased. When such events happen it usually means large losses for the owner of the bond and therefore, managing credit risk has become really important. For a commercial bank, large credit losses could be devastating, leading the bank into bankruptcy if its capital reserve is not large enough. The Riksbank does not face this kind of threat, but managing credit risk properly is still in the Riksbank's interest, since large credit losses might hurt its reputation and since it is important that the foreign currency reserve remains big enough to handle, for example, a financial crisis.

The use of risk measures has become very popular in credit risk management. Risk measures aim to quantify the risk in a credit portfolio in different ways and some widely-used risk measures will be introduced in section 2.2. These risk measures are the main output of the model I am studying in this thesis. The reason why risk measures have become so popular is because they make it easy to evaluate and compare different credit portfolios with respect to credit risk. For example, the measure Value-at-Risk (VaR) tries to indicate how large the credit loss could be in a "worst-case-scenario" (VaR is essentially a large percentile of the credit loss distribution). For a commercial bank, VaR could be matched with its capital reserve to make sure that losses can be covered. Risk measures are used by investors, but also by policy-makers and regulators to make sure that the banks do not take on too much risk. For example, they play a central role in the Basel Accords which are recommendations on banking laws and regulations by the Bank of International Settlement.

To calculate the risk measures, one needs a model for the credit risk. The challenge in modeling credit risk is the limited amount of data combined with the wish of estimating the tail of the credit loss distribution. The risk of bond issuers defaulting is usually very small, especially for sovereign issuers. If a default happens, however, the credit loss is usually very large. This gives a

loss distribution with a very long and fat tail. At the same time, the small number of observed defaults makes it hard to estimate the tail with high confidence. This is a typical extreme value problem. At portfolio level, another big challenge is how to model the joint loss distribution of different bonds. The correlation is known to be positive in most cases, but how to model and estimate the distribution is problematic and also suffers from the limited amount of data.

The model I will use in this thesis is essentially the same as the one presented in (van der Hoorn, 2009), which is a model that has been used by the European Central Bank (ECB). The differences between my model and the ECB model are minor, but include how future cashflows are discounted and the number of rating classes used. However, when it comes to estimating the model parameters and evaluating the certainty of the model I have investigated some different methods than those in the ECB paper, for example by estimating asset correlations from CDS spreads. The ECB model belongs to a group of models that started becoming popular in late 90s, when the RiskMetrics™ Group launched the benchmark model CreditMetrics™, which is presented in (Gupton et al., 1997). What is special and in common for this kind of models is that

- they use credit rating classes to describe the different states that a bond can end up in and default is the lowest rating class. The probability of migration to another class is based on historical data of migrations. These are collected into a so called migration matrix.
- they use a mark-to-market approach to evaluate credit losses, which means that a decrease in bond value one year from now due to a rating change is seen as a credit loss, even if the bond issuer never defaults later on.
- they model the joint distribution of losses for bonds with asset correlation, by assuming an underlying multivariate Gaussian variable, i.e. they use a Gaussian copula, with a fixed correlation matrix.

The benefit of the first two points, compared to a basic default/no-default model, is that it smoothens the credit loss distribution, better reflecting the actual credit risk that a bond owner faces. Furthermore, one can now make use of credit rating data in the model. As drawbacks on the first two points one could mention that one now has to trust the correctness of the credit rating agencies and that it is still the default state that will give the largest losses, which could mean that the use of rating classes contributes more to making the model more complicated, than to making the estimate of the loss distribution tail better. The reason for the third point, is that in the old default/no-default models it was hard to estimate the default correlation between different bonds with any certainty, since defaults are so rare. By instead assuming some underlying asset in the fashion of (Merton, 1974), there is a whole new range of data that could be used for the correlation estimation. In CreditMetrics, they mainly looked at corporate bonds, which meant that they could use stock movements and industry variables for the correlation estimation. For sovereign bonds, this is not

possible and therefore the estimation becomes harder. In the ECB paper they use a constant asset correlation of 24%, which of course removes the benefit of the asset correlation model entirely. Voices has also been risen, for example in (Wilmott, 2007, pp. 274-275), that models like this with fixed correlation parameters are generally bad, since the correlation does not seem to be constant when looking at asset time series.

In addition to CreditMetrics-type models, there exist a wide range of well known models for credit risk that are used in the industry. These include the KMV-model and the Credit Portfolio View-model. The KMV-model is similar to the CreditMetrics-model in that it uses the Merton-model with an underlying asset. While the CreditMetrics-model uses the underlying asset mainly to model correlation, the KMV-model consider corporations' total asset and equity values to extract the probability of default and KMV's own risk measure, Distance to Default. The KMV-model is however best suited for corporate bonds. The Credit Portfolio View-model includes macroeconomical factors to explain the credit risk. It seems natural that the probability of default for a bond issuer would increase during unstable periods in the world economy and therefore the macroeconomic variables should have some explanatory value. This can be seen as an alternative to model the correlation between bonds explicitly. For a more detailed introduction to these and other common credit risk models, see (Bluhm et al., 2003, chapter 2).

1.2 Purpose

The main purpose of the thesis is to implement, develop and evaluate the ECB model for the Riksbank's portfolio. This includes investigating different methods for estimating the parameters of the model. I look at methods from both the ECB- and CreditMetrics-paper as well as some methods I develop myself. Given the model and the parameters, I calculate the risk measures using Monte Carlo-methods. Most of the theory for this is presented in the ECB- and CreditMetrics-paper, but I also add methods for calculating confidence bounds for all risk measures and extend the methods to include importance sampling. Being able to calculate the risk measures, I analyse some different portfolios, to see how different compositions of bonds affect the risk measures. I make a study on the convergence rates of the Monte Carlo-methods. I perform a sensitivity study to see how much the risk measures are affected by changes in the model parameters. Especially, I develop a method to see how the measures are affected when using a migration matrix that is not too improbable to have generated data, which results in some kind of confidence bound for the risk measures.

I also look at the problem of finding the optimal portfolio with respect to the risk measures, i.e. the portfolio with the lowest risk. There is a theoretical background for this problem in (Rockafellar and Uryasev, 2000) and another good background can be found in (Iscoe et al., 2012). These papers say that the problem is convex under some conditions and that it can be solved numerically.

I try to solve the problem myself, using a genetic algorithm.

1.3 Results

The main results of this master thesis show that the model I examine gives risk measures of low certainty. This should however come as no surprise, knowing how limited the historical sovereign bond data is. Any model that tries to say something about the tail of the loss distribution should face the same problems. The methods I tried to use to smooth to the migration matrix was unsatisfactory and can be discarded. The same goes for the method used to estimate the asset correlations from CDS spreads, which rather shows weakness in the assumption of constant correlation. The comparison of different portfolios shows that the credit risk is decreased by owning higher rated bonds and by diversification, which was also expected. Comparing different Monte Carlo methods show that the calculation speed is increased by the use of importance sampling. The error induced by the Monte Carlo methods is small, even for small sample sizes, as compared to the model error. The genetic algorithm returns low-risk portfolios, but is unable to converge to the optimal portfolio at an acceptable rate.

None of these results can be said to really contribute to improve the model. The problems with the model were known before and the improvements I have tried to come up with have not worked out. The thesis mainly establishes the fact that credit risk is really hard to quantify with certainty, especially for sovereign bonds.

1.4 Outline

The thesis starts in section 2 with a brief introduction to credit risk and risk measures, which can basically be skipped for readers familiar with these subjects. This is followed by a presentation of the Monte Carlo methods with importance sampling that will be used, in section 3. Section 4 introduces the model, first as one-dimensional in 4.1 and generalized to many dimensions in 4.2. Section 4.3 explains methods for calculating risk measures and in section 4.4 different methods for estimating the parameters is discussed. This is followed by results and validation in section 5 and the optimization problem in section 6. The thesis ends with a discussion in section 7. The reader is expected to be familiar with basic mathematical statistics and some basic economics.

2 Credit risk

Credit risk is the risk of a counterparty not fulfilling its financial obligations. The easiest way to think about it is as the risk a bank takes when it accepts to give a customer a loan. If the customer is unable to repay his loan in time, the bank will generally lose money. There are other risks associated with lending money as well. For example, if inflation rises and the interest rate of the loan is fixed, the money repayed to the bank may be worth less in terms of actual goods than the money lent, even if the customer repays the loan in time. This might be referred to as inflationary risk. Similarly, if a Swedish bank buys bonds from an American company, which is a way of lending the company money, the bank faces a currency risk. If the USD/SEK-ratio decreases the money received from the company, once changed back to SEK, might be worth less than the initial investment. Inflationary risk, currency risk and other risks that do not depend on the counterparty's ability to fulfill its obligations are **not** considered to be credit risks. Therefore, when modeling credit risk, it is important to keep all other types of risk outside of the model.

In this paper, credit risk will be defined as the possibility of losses due to *defaults* or *credit rating changes* among counterparties. That a counterparty *defaults* in this context simply means that it fails to repay the entire loan in time, e.g. that it misses a coupon payment. A *credit rating change* means that the market price of a bond changes because a major rating agency has changed its credit rating. More about this in the next section. For an introduction to credit risk, see for example (Bluhm et al., 2003) or (Duffie and Singleton, 2003).

2.1 Credit ratings

A credit rating agency is a company that issues ratings for different kind of loans. The ratings are meant to show what credit quality an investor can expect from different obligors. Most rating agencies use some kind of letter combination for different ratings, e.g. AAA for an obligor with a very small estimated probability of default. See fig 2.1 for an example.

<p>'AAA'—Extremely strong capacity to meet financial commitments. Highest Rating. 'AA'—Very strong capacity to meet financial commitments. 'A'—Strong capacity to meet financial commitments, but somewhat susceptible to adverse economic conditions and changes in circumstances. 'BBB'—Adequate capacity to meet financial commitments, but more subject to adverse economic conditions. 'BBB-'—Considered lowest investment grade by market participants. 'BB+'—Considered highest speculative grade by market participants. 'BB'—Less vulnerable in the near-term but faces major ongoing uncertainties to adverse business, financial and economic conditions. 'B'—More vulnerable to adverse business, financial and economic conditions but currently has the capacity to meet financial commitments. 'CCC'—Currently vulnerable and dependent on favorable business, financial and economic conditions to meet financial commitments. 'CC'—Currently highly vulnerable. 'C'—Currently highly vulnerable obligations and other defined circumstances. 'D'—Payment default on financial commitments.</p> <p>Note: Ratings from 'AA' to 'CCC' may be modified by the addition of a plus (+) or minus (-) sign to show relative standing within the major rating categories.</p>
--

Figure 2.1: Rating classes from Standard & Poor's. Source: www.standardandpoors.com.

To decide what rating to give an obligor, the rating agencies consider both qualitative and quantitative information about the obligor's economical situation and tries to predict its future ability to repay debt. This information could be both hard figures, for example the obligor's current total debt, and facts that are more difficult to interpret, like the current political situation in the obligor's country. It is a difficult and exhaustive job to analyze all this information, therefore many investors use credit ratings issued by rating agencies instead of doing this analysis on their own.

The three by far most well-known and used credit agencies Fitch, Standard & Poor's and Moody's are all American, but there exists other smaller rating agencies from other parts of the world as well. Since so many investors use credit ratings from these three agencies, their ratings affect the market prices of bonds and other tradable debt. For example, an investor owning bonds issued by a large AA-rated country can at any time sell them at the bond market, assuming some liquidity on the bond market. However, if the country is downgraded to A-rating the market price is likely to instantly drop, reducing the present value of the investor's portfolio. The risk the investor faces by owning the bonds can be considered as credit risk.

Example 1. A simple model for credit risk could look as follows. A Swedish investor owns a bond issued by a British company. The company is B-rated by a rating agency that only uses four different rating classes: A, B, C and D for default. The bond matures in three years and then pays the owner 100 GBP. The current market price of the bond is 86.38, which corresponds to a 5% yield since

$$\frac{100}{1.05^3} = 86.38.$$

It is assumed that the yield is constant for this rating and that the investor

has a one-year-risk horizon. In one year the bond will have the market value $86.38 \cdot 1.05 = 90.70$. However, it is also assumed that an upgrade to rating A gives an 10% increase in market value, a downgrade to C a 10% decrease and a default gives a 50% decrease resulting in the values in table 1. The probabilities of these different states are assumed to be known as well.

Credit rating	Market value GBP	Probability	L
A	99.77	0.05	-9.07
B	90.70	0.90	0
C	81.63	0.04	9.07
D	45.35	0.01	45.35

Table 1: Different scenarios for the bond one year from now.

Since it is in some sense expected that the bond will stay in rating class B, the random variable L describing the credit loss in one year from now, is defined as the difference between the forward value given that the bond stays B-rated and the actual forward value. This results in the different values of L that can be seen in table 1. Also, the probability function of L is shown in figure 2.2.




Figure 2.2: The probability distribution of L in example 1.

The shape of the probability function is typical for credit risk. There is a

large probability of small or no loss and a small probability of a very large loss. Notice how the model tries to capture the credit risk in isolation. By holding the bond, the company in reality faces an interest risk as well, since the yield of the bond might change even without a credit rating change. Fixing the yield and subtracting from the forward value of a B-rated bond is an attempt to exclude the interest rate from the model. In the same way, if the company wanted to evaluate the risk in SEK instead, they could eliminate the currency risk from the model by fixing the exchange rate.

2.2 Risk measures

In order to quantify the credit risk in a portfolio it is common practice to use risk measures. To get the full picture of what the risk structure looks like, it is of course best to study the full probability density function of the loss L . However, this may be inconvenient if one for example wants to compare two different portfolios. If it is possible to find a single number that in an appropriate way describes how big the risk is, it would be possible for a bank to analyze time series of how the risk in the portfolio changes, make policies about how big risks that are acceptable and to minimize the risk given some constraints. In the following sections some commonly used risk measures will be presented. A reference for risk measures is (Dowd, 2002, pp. 27-44), which especially introduces and discusses VaR and ES . EL and UL are pretty straight forward, but are also defined in (van der Hoorn, 2009, pp. 125-128).

2.2.1 Expected Loss

Expected loss (EL) is simply defined as the expected value of the loss L , i.e.

$$EL = E[L].$$

EL is a very basic measure of risk and has some obvious drawbacks since it does not say anything about the diffusion of L . However, for risk-neutral investors, portfolios with the same EL are considered equal. Theoretically, any rational investor without any other information about the loss distribution would always prefer the one with the lowest EL . In reality, however, investors are known to be risk-averse, i.e. they do not want to risk losing money unless they are compensated for it. For example, a bank would in general not lend 100 USD with a 25% interest rate if the probability of getting the money back is 0.8, even though

$$EL = 0.8 \cdot (-25) + 0.2 \cdot 100 = 0.$$

To accept giving the loan at all the bank would demand a higher interest rate. However, in a situation where the bank has a choice between two different borrowers with equal probability of paying back the loan, the bank will always choose the one which gives the best interest rate, or equivalently, the lowest EL . This illustrates in what way EL can be used as a risk measure. If two loss distributions are known to be otherwise more or less similar, the one preferred is usually the one with the lowest EL .

2.2.2 Unexpected Loss

Unexpected Loss (UL) is defined as the standard deviation of L . UL is commonly used in risk management to show the diffusion of the distribution. For risk-averse investors, UL is of course preferred low. UL is a good measure when the loss distribution is light-tailed or approximately Gaussian (of course, when the distribution is perfectly Gaussian, EL and UL describes the distribution entirely). Unfortunately, for credit risk the loss distribution is often heavy-tailed since defaults often imply large losses. Therefore, UL works less well for credit risk than it does for many other kinds of risk.

2.2.3 Value-at-Risk

Value-at-Risk (VaR_α) is a widely used risk measure that focuses on the highest percentiles of the loss distribution. It can be defined as

$$VaR_\alpha = \min\{l : P(L > l) \leq 1 - \alpha\}, \quad (2.1)$$

where α usually represents a large part of the probability mass and typical values are $\alpha = 0.95$ or $\alpha = 0.99$. The idea behind Value-at-Risk is that it should show a “worst-case scenario”, e.g. $VaR_{0.99}$ could be interpreted as the amount of money that the loss will only exceed in one year out of a hundred, in the long run. See fig 2.3 for a graphical example. If the probability density $f_L(l)$ of L is continuous and $\neq 0$ on one interval only, then a more easily-interpreted definition of VaR_α can be used

$$VaR_\alpha = \{l : P(L > l) = 1 - \alpha\}.$$

I will sometimes in this paper, somewhat sloppy, use this definition even when the distribution is discrete to simplify my reasoning.




Figure 2.3: This graph shows an example of a loss density. The red section represents 5% of the total area. Thus, $VaR_{0.95} = 4.74$.

A big reason to why VaR has become so popular is that it is in general easy for everybody in an organization to understand the concept, as compared to UL . It can also be argued that VaR tells more about the tail than UL does. However, there are some well-known problems with using VaR as a risk-measure. For example, it is difficult to choose α in a way that both shows how large the losses could be in a very extreme case and how large they would be on just a bad day. For this reason it is common to look at VaR for many α :s at the same time. Another problem with VaR is that it does not satisfy sub-additivity, i.e. for two random variables L_1 and L_2

$$VaR_{\alpha}(L_1 + L_2) \leq VaR_{\alpha}(L_1) + VaR_{\alpha}(L_2)$$

does not hold in general (Dowd, 2002, p. 34). This is a problem since it contradicts the intuition that diversification, i.e. to hold many different assets, lowers the risk. When optimizing a portfolio in terms of VaR , the properties of VaR could lead to absurdities, which can be illustrated with an example.

Example 2. Consider two assets with loss distributions L_1 and L_2 . L_1 takes the value 0 with probability 0.95 and the value 1000 with probability 0.05. L_2 takes the value 900 with probability 0.991 and the value 10^9 with probability 0.009. Most people would agree that the first asset is less risky, but $VaR_{0.99}(L_1) = 1000$

and $VaR_{0.99}(L_2) = 900$, so $Var_{0.99}(L_1) > Var_{0.99}(L_2)$. This extreme example shows that VaR can lead to absurdity as a risk measure.

2.2.4 Expected Shortfall

Expected Shortfall (ES) is a risk measure that tries to fix the drawbacks with VaR . ES is known by many names in the risk literature, with Expected Tail-Loss, Conditional VaR and Tail VaR being the most common ones. The definition looks as follows

$$ES_\alpha = E[L|L > VaR_\alpha].$$

So, while the definition of $VaR_{0.99}$ could be interpreted as the amount that the loss will exceed on average in only one year out of a hundred, $ES_{0.99}$ can be interpreted as the expected value of the loss in those worst 1% years. For the loss density in figure 2.3, $ES_{0.95}$ is calculated as the center of mass of the red section, i.e.

$$ES_{0.95} = \frac{1}{0.05} \cdot \int_{VaR_{0.95}}^{\infty} l \cdot f_L(l) dl = 5.92,$$

where $f_L(l)$ is the probability density. Unlike for VaR , sub-additivity always holds for ES (Dowd, 2002, p. 37). Moreover, ES can handle cases like the one in example 2 in a better way, since the whole tail affects its value. The problem with choosing a good α still remains, but is perhaps made somewhat smaller, since ES is always affected by the outermost part of the tail.

3 Monte Carlo methods

Monte Carlo methods is used in statistics in order to make calculations. The basic idea is to use a mathematical model for a phenomenon that appears random in order to make simulations. If the simulations are easy to generate, then one can produce a large sample and use the law of large numbers to approximate for example the expected value of some process. Monte Carlo simulations works best, compared to other methods, when working with complex models of high dimension, since it may then be difficult to calculate things analytically or even numerically with deterministic methods. An introduction to Monte Carlo methods can be found in (Sköld, 2006).

In credit risk, when working with a portfolio of risky assets, the dimension is usually high and Monte Carlo simulations are therefore very useful. For now, set aside the underlying model and just consider the case when it is possible to generate samples from the random loss variable L . Let N be the number of samples generated and let L_1, L_2, \dots, L_N be the generated samples. An estimate of EL can then be calculated as the mean, i.e.

$$\widehat{EL} = \bar{L} = \frac{1}{N} \sum_{i=1}^N L_i.$$

The estimate is unbiased and by the law of large numbers $\widehat{EL} \rightarrow E[L]$ as $N \rightarrow \infty$. The central limit theorem makes it possible to construct a confidence interval for a given N . Since

$$\sqrt{N} \left(\widehat{EL} - E[L] \right) \sim N \left(0, \sigma^2(L) \right)$$

for large N a two-sided q -confidence interval for EL can be constructed as

$$\left(\widehat{EL} - \lambda_{q/2} \frac{\sigma(L)}{\sqrt{N}}, \widehat{EL} + \lambda_{q/2} \frac{\sigma(L)}{\sqrt{N}} \right), \quad (3.1)$$

where $\lambda_{q/2}$ denotes the the $q/2$ -quantile of the standard normal distribution and where $\sigma^2(L)$ can be approximated by

$$\sigma^2(L) \approx \frac{1}{N-1} \sum_{i=1}^N (L_i - \bar{L})^2. \quad (3.2)$$

Now, instead assume that EL is known. In this case UL can be estimated with basically the same method. Let $X_i = (L_i - EL)^2$, $i = 1, \dots, N$. Then

$$\widehat{V[L]} = \bar{X} = \frac{1}{N} \sum_{i=1}^N X_i$$

is an unbiased estimator of the variance of L since

$$E \left[\widehat{V[L]} \right] = E \left[\frac{1}{N} \sum_{i=1}^N X_i \right] = \frac{1}{N} \sum_{i=1}^N E[(L_i - EL)^2] = V[L].$$

Thus, an estimate of the variance of L is available and a confidence interval can be constructed in the same way as for EL . Since UL is defined as the standard deviation of L , taking the square root of this estimate and the confidence bounds will give an estimate and a confidence interval for UL . Of course, another estimate of UL is given by the square root of (3.2), but the method I present here makes it easier to construct a confidence interval.

To get an estimate of VaR_α , first order the simulated samples. Let $L_{(1)}, L_{(2)}, \dots, L_{(N)}$ denote the ordered samples, such that $L_{(1)} \leq L_{(2)} \leq \dots \leq L_{(N)}$. The definition of VaR_α (2.1) suggests that a good estimate should be the smallest value l such that at most $1 - \alpha$ of the samples are larger than l . Using the empirical distribution function

$$\widehat{F}_L(l) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{L_i < l\}} \quad (3.3)$$

as an approximation of the true loss distribution function would imply that the estimate should be $\widehat{VaR}_\alpha = L_{(\lceil \alpha N \rceil)}$ (using $\lceil \cdot \rceil$ as a notation for rounding to the nearest larger integer). Intuitively, this estimator is asymptotically unbiased. A

way to construct a confidence interval for VaR_α is presented in (Gupton et al., 1997, pp. 150-151). The method is based on the binomial distribution. Given the true value of VaR_α , the random variable Y representing the number of samples larger than VaR_α is binomially distributed, i.e. $Y \in Bin(N, 1 - \alpha)$. Notice, that for some integers A, B , such that, $0 < A \leq B < N$ the event

$$A \leq Y \leq B$$

is exactly the same event as

$$L_{(N-B)} < VaR_\alpha < L_{(N+1-A)}.$$

For example, if it is known that the number of simulated samples that exceed VaR_α is greater than or equal to 3, but smaller than or equal to 6, then one knows that third largest sample must be greater than VaR_α and that the seventh largest sample must be smaller than VaR_α and vice versa. Because of this, a confidence interval for Y gives confidence interval for VaR_α as well. If N is fairly large and α is not too close to one (rule of thumb: if $N\alpha(1-\alpha) > 10$ (Blom et al., 2005)) a normal approximation can be used for Y , i.e. $Y \sim N(N\alpha, N\alpha(1-\alpha))$. A q -confidence interval for Y is thus

$$\left(N\alpha - \lambda_{q/2} \sqrt{N\alpha(1-\alpha)}, N\alpha + \lambda_{q/2} \sqrt{N\alpha(1-\alpha)} \right)$$

and an interval for VaR_α can be constructed as

$$\left(L_{(\lfloor N - (N\alpha + \lambda_{q/2} \sqrt{N\alpha(1-\alpha)}) \rfloor)}, L_{(\lceil N+1 - (N\alpha - \lambda_{q/2} \sqrt{N\alpha(1-\alpha)}) \rceil)} \right),$$

where the rounding has been chosen to make the interval as large as possible. This bias will decrease as N increases.

For estimating ES_α we can use the fact that we already have an estimate of VaR_α . If we insert \widehat{VaR}_α into the definition of ES_α we see that ES_α is the expected value of losses greater than \widehat{VaR}_α . Thus, we can use the ordered samples to get the estimate as

$$\widehat{ES}_\alpha = \frac{1}{N - \lceil \alpha N \rceil} \sum_{i=\lceil \alpha N \rceil + 1}^N L_{(i)}$$

, i.e. as the mean of the losses greater than \widehat{VaR}_α . In the same way a confidence interval for ES_α can be constructed by taking the mean of losses greater than the confidence bounds for VaR_α . Another, equivalent way of estimating ES_α is as

$$\widehat{ES}_\alpha = \widehat{VaR}_\alpha + \frac{1}{(1-\alpha)N} \sum_{i=1}^N \max(L_i - \widehat{VaR}_\alpha, 0), \quad (3.4)$$

which does not need the samples to be ordered. To show why this estimator works we consider the distribution function

$$F_{L|L>x}(l) = P(L \leq l | L > x) = \frac{P(x < L \leq l)}{P(L > x)} = \frac{F_L(l) - F_L(x)}{P(L > x)}.$$

Taking the derivative of both sides gives

$$f_{L|L>x}(l) = \frac{f_L(l)}{P(L > x)}.$$

Now

$$\begin{aligned} E[L|L > x] &= \int_x^\infty l \cdot f_{L|L>x}(l) dl = \frac{1}{P(L > x)} \int_x^\infty l \cdot f_L(l) dl = \\ &= \frac{1}{P(L > x)} \int_x^\infty (l+x-x) \cdot f_L(l) dl = \frac{1}{P(L > x)} \left(xP(L > x) + \int_x^\infty (l-x) \cdot f_L(l) dl \right) = \\ &= x + \frac{1}{P(L > x)} \int_0^\infty \max(l-x, 0) \cdot f_L(l) dl = x + \frac{1}{P(L > x)} E[\max(L-x, 0)]. \end{aligned}$$

Inserting $x = VaR_\alpha$ on both sides finally gives

$$\begin{aligned} E[L|L > VaR_\alpha] &= VaR_\alpha + \frac{1}{P(L > VaR_\alpha)} E[\max(L - VaR_\alpha, 0)] \\ \Leftrightarrow ES_\alpha &= VaR_\alpha + \frac{1}{1-\alpha} E[\max(L - VaR_\alpha, 0)], \end{aligned} \quad (3.5)$$

which explains the estimator in (3.4). We will return to this estimator in section 6 on optimization.

3.1 Importance sampling

Importance sampling (IS) is used in Monte Carlo methods as a so called variance reduction technique. The idea behind IS is that some possible outcomes of the random variable might be more important for the estimation than others and to make sure that these outcomes are included in the samples one can tweak their probabilities. For example, in credit risk there is typically a small chance of very large losses. As the estimates of VaR and ES depend strongly on the largest outcomes of the simulation these estimates can vary a lot depending on if the simulated sample contains many or few of these largest outcomes. Since the probability of the largest losses is typically very small, the estimators will have a large variance.

With IS the probability density $f_L(l)$ of L is replaced by an instrumental distribution $g_L(l)$ that will typically increase the probability of the more important outcomes. To be able to get correct estimates using the instrumental distribution, the estimators presented in the last section will have to be changed

to compensate for this. When calculating the expected value of an arbitrary function $\phi(L)$ we used in the last section that

$$E[\phi(L)] = \int \phi(l)f_L(l)dl \approx \frac{1}{N} \sum_{i=1}^N \phi(L_i),$$

where L_i are the simulated samples from $f_L(l)$. Now, when instead using the instrumental distribution for simulating we can use that

$$E[\phi(L)] = \int \phi(l)f_L(l)dl = \int \phi(l)\frac{f_L(l)}{g_L(l)}g_L(l)dl = \int \phi(l)\omega(l)g_L(l)dl \approx \frac{1}{N} \sum_{i=1}^N \phi(L_i)\omega(L_i),$$

where L_i is now simulated using $g_L(l)$ and $\omega(l) = \frac{f_L(l)}{g_L(l)}$ is called the weighting function. This gives us a way to make estimations using the instrumental distribution $g_L(l)$, using the weighting function of each sample as a way to compensate for its new probability.

EL can now be estimated as

$$\widehat{EL}^{IS} = \frac{1}{N} \sum_{i=1}^N L_i\omega(L_i).$$

A confidence interval can also be constructed similarly to that in the standard case. The only difference is that the estimate of σ in (3.1) is replaced by

$$\sigma^2 \approx \frac{1}{N-1} \sum_{i=1}^N \left(L_i\omega(L_i) - \frac{1}{N} \sum_{j=1}^N L_j\omega(L_j) \right)^2.$$

For the importance sampling to be meaningful this new estimate of σ should be lower than the old one. This will maybe not be the case when estimating EL because EL does not depend that much on the tail events, but it will be for the other risk measures. Assuming as before that EL is now known, we can also make an estimate of UL as

$$\widehat{UL}^{IS} = \sqrt{\frac{1}{N} \sum_{i=1}^N (L_i - EL)^2\omega(L_i)}$$

and construct a confidence interval similarly to before.

To get an estimate of VaR_α I will take a slightly different approach since VaR_α is not defined as an expectation of a function of L , but a rather as a percentile to the probability density $f_L(l)$. Assume now that VaR_α is known but that α is unknown. We could then get an estimate of α using that

$$\alpha = P(L < VaR_\alpha) = E[\mathbf{1}_{\{L < VaR_\alpha\}}],$$

so

$$\hat{\alpha} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{L_i < VaR_\alpha\}} \omega(L_i)$$

is an unbiased estimate of α . Now, again assume that α is known and that VaR_α is not. If we could find a number VaR'_α s.t.

$$\hat{\alpha}' = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{L_i < VaR'_\alpha\}} \omega(L_i) \approx \alpha,$$

VaR'_α should be a pretty good estimate of VaR_α . A way to find such a number VaR'_α is the following. Order the samples L_1, \dots, L_N like in the last section to obtain the ordered sample $L_{(1)}, \dots, L_{(N)}$. By summing up the weights $\omega(L_{(1)}), \dots, \omega(L_{(N)})$ in reverse order we can find a number k s.t.

$$\frac{1}{N} \sum_{i=k+1}^N \omega(L_{(i)}) \leq 1 - \alpha,$$

but

$$\frac{1}{N} \sum_{i=k}^N \omega(L_{(i)}) > 1 - \alpha.$$

By comparing with the definition of VaR it can be argued that $\widehat{VaR}_\alpha^{IS} = L_{(k)}$ will be a good estimator. This method of course requires that $\frac{1}{N} \sum_{i=1}^N \omega(L_{(i)}) > 1 - \alpha$. To construct a confidence interval for VaR_α when using IS I will use the same reasoning as in the standard case. However, in this case we do not know the probability of a simulated sample being larger than $\widehat{VaR}_\alpha^{IS}$, but the best estimate is naturally $\frac{N-k}{N}$. Therefore I let $Y \in Bin(N, \frac{N-k}{N})$ and proceed as in the standard case.

For ES_α we once again use that we have already estimated VaR_α and use (3.5), which gives the estimator

$$\begin{aligned} \widehat{ES}_\alpha^{IS} &= \widehat{VaR}_\alpha^{IS} + \frac{1}{1-\alpha} \frac{1}{N} \sum_{i=1}^N \max(L_i - \widehat{VaR}_\alpha^{IS}, 0) \omega(L_i) = \\ &\widehat{VaR}_\alpha^{IS} + \frac{1}{1-\alpha} \frac{1}{N} \sum_{i=k+1}^N (L_{(i)} - \widehat{VaR}_\alpha^{IS}) \omega(L_{(i)}). \end{aligned}$$

The confidence interval for ES_α can be constructed in a similar way, using the confidence bounds for VaR_α .

4 Modeling

I will in this section introduce a portfolio model for credit risk. The model is very similar to one that ECB has used, presented in (van der Hoorn, 2009). This itself is clearly based on the framework model CreditMetrics™, presented in (Gupton et al., 1997). I start by explaining how the model works for a single bond. This is then extended to a model for a portfolio of bonds. I show how the different risk measures can be calculated and finally, I discuss how the different parameters of the model could be handled and estimated.

The main purpose of the model is to provide a way to make estimates of the different risk measures presented in section 2.2. This is mainly done by using Monte Carlo simulations, since the model provides a simple way to generate random portfolio credit losses from the loss distribution $F_L(l)$. These simulations can of course also be used to get a picture of what the whole $F_L(l)$ looks like using the approximation $\hat{F}_L(l)$, defined in (3.3). In section 6 we will also see how the model can be used for minimizing the risk.

4.1 One-dimensional model

If the portfolio consists of only one bond or bonds from a single obligor, the model will be one-dimensional. In this case the model will be much simpler since the joint probability functions of different bonds do not have to be considered. The one-dimensional model will in fact be very similar to the model described in example 1. For the model to work there must exist a rating system that links every obligor with a certain rating. In the ECB-model 9 different rating classes is used, but in the model I have implemented I have used 18 different rating classes from AAA to D and the ratings has been taken from Fitch¹. The number of rating classes does not matter for the explanation of the model and it is unclear how it effects the performance as well. The reason why I use 18 is mainly a practical matter of data availability. Sometimes the ratings will be referred to as numbers rather than letters in their natural order, i.e. AAA= 1, AA+= 2, ..., D = 18.

As mentioned before the risk horizon considered will be one year. The model could be generalized to work for different time horizons as well, but this will lie outside the scope of this paper. One of the most important parameters for the model will therefore be the probabilities of rating changes for different obligors in one year's time. A fundamental assumption will be that obligors that have the same rating today will have the same probability of a rating change to a given rating in one year's time. These probabilities can therefore be presented in a so called migration matrix. Migration is just another word for rating change. If the number of different ratings is r the size of the migration matrix M will be $r \times r$ and on position (i, j) the matrix will contain the probability that an obligor

¹The complete list of ratings is AAA,AA+,AA,AA-,A+,A,A-,BBB+,BBB,BBB-,BB+,BB,BB-,B+,B,B-,C,D.

of rating i migrates to rating j in one year's time. An example of a migration matrix can be seen in table 2. In practice migration matrices is usually based on historical data, which will be described in section 4.4.1.

From/To	A	B	C	D
A	0.86	0.119	0.02	0.001
B	0.05	0.90	0.04	0.01
C	0.02	0.05	0.91	0.02
D	0	0	0	1

Table 2: A fictional migration matrix with 4 rating classes. Every row contains the probabilities of migration to the rating classes represented by the columns. For example, the probability of an obligor upgrading from C to A is 2% and the probability of a B-rated obligor moving into default is 1%.

In addition to assuming that obligors with the same rating have the same probabilities of migration it will also be assumed that they are associated with the same yield. This means that all future cash flows coming from bonds issued by the same obligor will be discounted using the same discount factor. The discount factor used for an obligor with rating i will be denoted df_i . Now, suppose that we have a portfolio containing bonds issued by a single obligor with rating i_{old} and that we will receive cash flows at times t_1, t_2, \dots, t_k of sizes CF_1, CF_2, \dots, CF_k . Then the value of the portfolio one year from now will be

$$FV = \sum_{j=1}^k df_{i_{old}}^{t_j-1} CF_j,$$

assuming that the obligor's rating does not change and that $t_j \geq 1$ for all i . In the same way, if the obligors rating changes to i , the new value will be

$$\sum_{j=1}^k df_i^{t_j-1} CF_j.$$

If the obligor defaults no future cash flows will be received, but we instead assume a constant recovery rate RR , so that the new value will be $FV \cdot RR$. We are now ready to define the conditional credit loss CL_i given a rating change as

$$CL_i = \begin{cases} FV - \sum_{j=1}^k df_i^{t_j-1} CF_j & \text{if } i \neq D \\ FV(1 - RR) & \text{if } i = D. \end{cases} \quad (4.1)$$

L can thus be seen as a random variable that takes the values CL_1, CL_2, \dots, CL_r with probabilities from the migrations matrix, i.e. $M(i_{old}, 1), M(i_{old}, 2), \dots, M(i_{old}, r)$. The probability function $p_L(l)$ of L is thus very simple and it is straight forward

to calculate the risk measures EL, UL, VaR and EL analytically. For example,

$$\begin{aligned}
 EL &= \sum_{\tau=1}^r M(i_{old}, \tau) \cdot CL_{\tau} \\
 UL &= \sqrt{\sum_{\tau=1}^r M(i_{old}, \tau) \cdot (CL_{\tau} - EL)^2}. \tag{4.2}
 \end{aligned}$$

4.2 Multi-dimensional model

The issue when moving from the one obligor portfolio to a portfolio consisting of many obligors is how to model the joint distribution of rating changes, while keeping the marginal distributions for single obligors equal to the one-dimensional model. One cannot assume that the migrations of different obligors are independent since it is presumable that countries that are strongly linked economically will have a similar economical situation and will hence move similarly in credit ratings. That this is the case for corporations is shown empirically by Gupton et al. (1997, pp. 81-83), where they estimate confidence intervals for default correlations among firms, which do not cover zero.

Instead, a model in the fashion of the classical model by Merton (1974) will be used. The idea is that for every obligor there exist an underlying Gaussian random variable Z . If the obligor is a company, Z could be thought of as the total value of all the company's assets one year in the future. In my case, where the portfolio consists of obligors that are sovereign states, the interpretation of Z is not that simple. It may however be thought of in a similar way as the amount of money the state could make available in the near future. It is also assumed that there exists some fix numbers q_D, q_C, \dots, q_{AA+} such that $q_D \leq q_C \leq \dots \leq q_{AA+}$. These numbers are meant to represent certain levels of debt that corresponds to the different rating classes. Especially, q_D represents the total value of the payments that the obligor has to make one year from now. If the obligor is unable to make the payments, i.e. if $Z < q_D$, the obligor will default. In the same way, $q_D \leq Z < q_C$ means that the obligor gets rating C, $q_C \leq Z < q_{B-}$ means that the obligor gets rating B- etc. However, we already know the probabilities of these events from the migration matrix. So, if the obligor has rating i these probabilities, for the Z linked to that obligor, would be $P(q_{AA+} \leq Z) = M(i, 1), P(q_{AA} \leq Z < q_{AA+}) = M(i, 2), \dots, P(Z < q_D) = M(i, 18)$. Now, given the migration matrix the relationship between Z and q_D, \dots, q_{AA+} is fixed so Z can be normalized to a standard Gaussian variable, i.e. $Z \in N(0, 1)$ and then (the normalized version of) q_D, \dots, q_{AA+} can be uniquely determined. If Φ^{-1} denotes the inverse standard Gaussian distribution

function then

$$\begin{aligned}
 q_D &= \Phi^{-1}(M(i, 18)) \\
 q_C &= \Phi^{-1}(M(i, 17) + M(i, 18)) \\
 &\vdots \\
 q_{AA+} &= \Phi^{-1}\left(\sum_{j=2}^{18} M(i, j)\right).
 \end{aligned}$$

This implies that q_D, \dots, q_{AA+} will be the same for obligors of identical ratings. The relationship of M, Z and the q :s is visualized in fig 4.1. Now, a new matrix Q can be constructed. Q will have size $r \times r - 1$ and each row will contain the values q_D, \dots, q_{AA+} corresponding to the different initial ratings.




Figure 4.1: This graph shows the probability density of Z together with the values q_D, q_C, q_B derived using the second row in the fictional migration matrix M in table 2. The areas under the graph named A, B, C and D have sizes equal to the probabilities $M(2, 1), M(2, 2), M(2, 3)$ and $M(2, 4)$. For a B-rated obligor the rating one year in the future can thus be decided by comparing the outcome of Z to the different q :s. For example, the event $q_D \leq Z < q_C$ is equivalent to the obligor being downgraded to C.

Using the model with underlying Gaussians we can now construct the multi-dimensional model. Consider a portfolio with n obligors. We construct the

multivariate Gaussian vector

$$\bar{Z} = \begin{bmatrix} Z_1 \\ Z_2 \\ \vdots \\ Z_n \end{bmatrix},$$

where Z_1, \dots, Z_n are the underlying Gaussians linked to obligor $1, \dots, n$. Since all Z_1, \dots, Z_n are standard Gaussians they all have mean 0 and variance 1, but we introduce the covariance matrix

$$\Sigma = \begin{bmatrix} 1 & \rho_{1,2} & \cdots & \rho_{1,n} \\ \rho_{2,1} & 1 & \cdots & \rho_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{n,1} & \rho_{n,2} & \cdots & 1 \end{bmatrix},$$

where $\rho_{i,j}$ is the correlation between Z_i and Z_j implying that $\rho_{i,j} = \rho_{j,i}$ and $-1 < \rho_{i,j} < 1$ for all $i \neq j$. The different ρ :s will be input parameters to the model and how to estimate them will be discussed in section 4.4.3.

The underlying multivariate variable \bar{Z} describes the joint distribution of rating changes completely and the credit loss L for the n -obligor portfolio can now be defined as

$$L = \sum_{i=1}^n L_i,$$

where L_i now denotes the credit loss from obligor i , defined as in one-dimensional model. While it was fairly simple to compute the probability function $p_L(l)$ of L for the one-dimensional model, for the multi-dimensional model it quickly becomes computationally impossible as n grows. The Riksbank's portfolio typically consists of $n = 30$ obligors. Using the system with $r = 18$ rating classes implies that there is a total of $r^n = 18^{30} \approx 5 \cdot 10^{47}$ combinations of rating changes for the different obligors. To calculate the probability of each of these outcomes one has to integrate the multivariate Gaussian density function over an n -dimensional rectangle. For example, to calculate the probability that all obligors default one must integrate over the rectangle with corners in

$$\begin{bmatrix} -\infty \\ -\infty \\ \vdots \\ -\infty \end{bmatrix}, \begin{bmatrix} q_{D,1} \\ q_{D,2} \\ \vdots \\ q_{D,n} \end{bmatrix},$$

where $q_{D,1}$, etc. are the q_D :s associated with the ranking of obligor 1 etc. A quick test in Matlab using the built-in function `mvncdf` and $n = 25$ (which is the highest dimension that the Matlab function allows) shows that each such integration takes approximately 0.3 seconds. Needless to say, calculating $p_L(l)$ exactly using this method is not possible.

4.3 Calculating risk measures

We saw in the last section that it was very hard to obtain the probability function of the credit loss L exactly as the dimension of the model became larger. One can however say much about the credit risk just by examining the risk measures EL, UL, VaR and ES . In this section I will show how these can be calculated. EL and UL can be calculated exactly analytically and all of them can be approximated using Monte Carlo methods. The Monte Carlo method will give an approximation of $p_L(l)$ as well.

To calculate EL is very simple since the joint distribution of rating changes does not matter. We have

$$EL = E[L] = E\left[\sum_{i=1}^n L_i\right] = \sum_{i=1}^n E[L_i],$$

where $E[L_i]$ is the expected loss from obligor i , calculated as in (4.2). To calculate UL is a little trickier and more time consuming, but can be performed in acceptable time. UL is computed as $\sqrt{Var[L]}$. We have that

$$Var[L] = Var\left[\sum_{i=1}^n L_i\right] = \sum_{i=1}^n Var[L_i] + 2 \sum_{i<j} Cov[L_i, L_j].$$

Here each $Var[L_i]$ can be obtained as the square of UL in (4.2). In the second term $Cov[L_i, L_j]$ can be calculated by first computing the joint loss probability function for the corresponding two-obligor portfolio, let us call it $p_{L',L''}(u,v)$. Then $E[L_i \cdot L_j]$ can be calculated as

$$E[L_i \cdot L_j] = \sum_{u=1}^r \sum_{v=1}^r uv p_{L',L''}(u,v)$$

and $Cov[L_i, L_j] = E[L_i \cdot L_j] - E[L_i]E[L_j]$. To get every $p_{L',L''}(u,v)$ one must do r^2 integrations over a bivariate normal density and there will be $\frac{n(n-1)}{2}$ covariances to calculate resulting in $\frac{n(n-1)}{2}r^2$ integrations in total. For the typical case $r = 18, n = 30$ this gives $\approx 10^5$ integrations in total. A Matlab test shows that each integration takes approximately 0.006 seconds in the bivariate case, so calculating UL for the whole portfolio would take something like $600 \text{ s} = 10$ minutes. This shows that UL is definitely computable analytically, but using Monte Carlo methods will be much faster and accurate enough so that is the method that I will use from now on.

To use the Monte Carlo methods explained in section 3, one needs to be able to generate pseudo-random samples from of the portfolio credit loss variable L . This can be done by first simulating the multivariate Gaussian vector \bar{Z} of underlying variables. This can be done for arbitrary mean and covariance matrix with the Matlab function `mvrnd`. I will not go into detail how this

is done mathematically, but it involves a Cholesky factorization of the covariance matrix Σ and is explained in (van der Hoorn, 2009, pp. 130-131). Once \bar{Z} is simulated each Z_i can be compared to the adequate row in the matrix Q that corresponds to that obligors current rating. This way one gets the rating changes of all obligors in the portfolio and the credit loss for each obligor can be obtained. Summing them up gives L . To make the simulation fast I have noticed that it is advantageous to first “guess” that each obligor will stay in the same rating class, when comparing the values in \bar{Z} to those in Q . This is because there is usually a high probability for the obligors staying, so making this guess reduces the number of comparisons.

A test on my computer shows that each sample from L takes approximately $2.5 \cdot 10^{-6}$ seconds to generate when $n = 29$, so a million samples or so can easily be generated. Now, when it is possible to generate such a large sample from L , the methods in section 3 can successfully be used to give good approximations of the risk measures. An approximation of the probability function can also be constructed as

$$\hat{p}_L(l) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{l=L_i\}},$$

where N is the number of samples (this is more or less equivalent to the EDF defined in (3.3)). The probability function can be visualized by plotting a histogram of the samples, for results see section 5.

Usually the probabilities of downgrade or default for obligors are very low. Therefore the probability function will naturally look skew with a fat tail, like in example 1. The risk measures VaR and EL that depend mostly on the tail of the distribution are therefore sensitive to the largest samples from the simulation. Since the worst outcomes will occur very seldom even in a million samples, this will result in a high variance for the estimates of VaR and EL and large confidence intervals. The variance can be decreased by the use of importance sampling as explained in section 3.1. For IS to work we need to find an instrumental distribution $g_L(l)$ that will simulate the larger losses with a higher probability. This can be done in a number of ways, but the method I will use is to change the mean for the underlying variable \bar{Z} . Originally the mean $\mu_{\bar{Z}} = \bar{0}$, but since low outcomes for the elements in \bar{Z} implies downgrades we can increase the probability of downgrades in the simulation by lowering the mean of \bar{Z} . How much the mean should be lowered for the different obligors to get a high variance reduction depends on the probabilities in the migration matrix. If we denote the original density function of \bar{Z} , $f_{\bar{Z}}$ and the new one $g_{\bar{Z}}$, then we can construct the weighting function ω needed for the IS estimators as

$$\omega(L) = \frac{f_{\bar{Z}}(\bar{Z})}{g_{\bar{Z}}(\bar{Z})}.$$

Now we have everything we need to calculate the risk measures using importance sampling. There is also a way to estimate $p_L(l)$ using the IS generated samples.

Since

$$p_L(l) = P(L = l) = E [\mathbf{1}_{\{l=L_i\}}],$$

then

$$\hat{p}_L^{IS}(l) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\{l=L_i\}} \omega(L_i).$$

4.4 Estimating parameters

So far in this section I have presented the model and showed how the different risk measures can be computed, given this model. However, for the model to be useful, we need a good way to estimate the parameters that the model is based upon. These parameters include the migration matrix, the conditional credit losses and the correlation matrix.

4.4.1 Migration matrix

The migration matrix M is supposed to contain the probabilities of obligors being in the different rating classes in one year's time given their current rating classes. The most basic way of estimating this is of course to look at historical data and register how many migrations there has been between each combination of ratings every year. The rating agencies all provide these historical matrices, see fig 4.2 for an example.

Fitch Sovereign IDR Transition Rates —1995–2011

(%, Modifier Level)
Average Annual

	AAA	AA+	AA	AA-	A+	A	A-	BBB+	BBB	BBB-	BB+	BB	BB-	B+	B	B-	CCC to C	D	Total
AAA	98.52	0.92	0.00	0.46	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
AA+	12.70	80.95	3.17	3.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
AA	0.00	10.00	84.44	3.33	1.11	1.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
AA-	0.00	0.00	13.79	75.86	5.17	0.00	0.00	1.72	1.72	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.72	0.00	100.00
A+	0.00	0.00	0.00	10.42	83.33	2.08	0.00	0.00	2.08	2.08	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
A	0.00	0.00	0.00	0.00	16.67	77.27	1.52	3.03	1.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
A-	0.00	0.00	0.00	0.00	0.00	14.49	79.71	5.80	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
BBB+	0.00	0.00	0.00	0.00	0.00	4.17	16.67	64.58	10.42	4.17	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
BBB	0.00	0.00	0.00	0.00	0.00	0.00	4.17	15.28	72.22	6.94	1.39	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00
BBB-	0.00	0.00	0.00	0.00	0.00	0.00	0.00	19.77	70.93	6.98	0.00	0.00	0.00	1.16	0.00	1.16	0.00	0.00	100.00
BB+	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.13	19.15	71.28	4.26	1.06	0.00	0.00	0.00	0.00	0.00	2.13	100.00
BB	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.49	2.99	20.90	67.16	5.97	0.00	0.00	0.00	0.00	0.00	1.49	100.00
BB-	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.38	10.71	72.62	8.33	4.76	1.19	0.00	0.00	0.00	100.00
B+	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.39	19.44	72.22	5.56	1.39	0.00	0.00	0.00	100.00
B	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	4.17	15.28	69.44	6.94	2.78	1.39	0.00	0.00	100.00
B-	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.17	0.00	0.00	2.17	19.57	69.57	6.52	0.00	100.00
CCC to C	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	6.25	18.75	50.00	25.00	100.00

Figure 4.2: A migration matrix from Fitch simply reflecting the average number of rating changes every year in the period 1995-2011 among the sovereign issuers that Fitch rates.

The matrix in fig 4.2 is typical for the migration matrices that the rating agencies provide. It has a lot of zeroes, reflecting that it is unusual for obligors to migrate more than one or two rating classes in one year. There are however some exceptions, for example the 1.72% that has moved from AA- to B-. The reason to why the matrix looks like this is of course that the data is very limited. It has only been collected for 17 years and there are presumably not that many issuers in every rating class. To use this matrix directly in the model is probably

a bad idea and doing so gives very unrealistic results. I did a quick test with this matrix and some different fictional portfolios of the same value. The test showed a much higher level of risk, in the sense of VaR and ES , for portfolios that contained obligors with rating AA-, than for example portfolios containing obligors rated A-. This is very contradictory, since the whole idea behind the rating system is to show which obligors that are more risky.

To make the model more realistic I would like to change or re-estimate the matrix in some way. I have looked at a few options for this. In the ECB model presented in (van der Hoorn, 2009) they do this by only focusing on the probability of default (PD) for every rating class. Since the default state gives the biggest losses this is the most important to change and they argue that even if it there are no observed defaults for the highest ratings, it could happen and therefore the probabilities in the last column should be positive and in ascending order. Hence, they assign positive PD:s for all initial rankings. What probabilities to use seems mainly subjective, but they try to use something that seems reasonable. For example they let the PD for AAA-rated obligors be 0.01%, for AA-rated 0.04% and so on. To make sure every row sums to one they lower the probabilities on the diagonal. This method is very ad hoc, but may generate results that seem more reasonable than when using the original matrix.

In (Gupton et al., 1997) they mention some properties that are desirable for a migration matrix, for example that better ratings never should have a higher chance of default or downgrade to a certain rating and that the chance of migrating to a given rating should be greater for more closely adjacent rating categories. They then try to find a “best” fit in some least squares sense to the original migration matrix, while keeping these constraints fulfilled. In addition to fitting to the original matrix they also try to make it fit some other data, for example historical migration averages for longer time periods than one year by assuming Markov properties, i.e. that the migration matrix for n years should be equal to M^n . They do not say exactly what method they use for finding the least squares fit, but I try to do something similar to their method using the build in Matlab function `lsqnonlin`. I use three different migration matrices as input, corresponding to one, three and ten years migration times. Unfortunately, I find it very hard to make the method robust and at the same time keep every probability greater than zero, guarantee that every row has sum one, etc.

Instead, I try something different, inspired by the method that we used to construct the multi-dimensional model, with the underlying Gaussian variables. Recall that underlying variable Z could be thought of as the total value of the obligors assets and q_D, \dots, q_{AA+} as some fixed levels of debt. When we constructed the multi-dimensional model we let all Z :s be standard normal and the q :s be different for every initial rating. Now, instead assume that there only exist one set of q_D, \dots, q_{AA+} and that there exists different Z :s for every initial rating and let us call them Z_{AAA}, \dots, Z_C . I will assume that each Z has a

different mean μ_{AAA}, \dots, μ_C and that all have the same variance $\sigma^2 = 1$, s.t.

$$q_D \leq \mu_C < q_C \leq \mu_{B-} < \dots \leq \mu_{AA+} < q_{AA+} \leq \mu_{AAA}. \quad (4.3)$$

This assumption would be consistent with their current rating. I fix $q_D = 0$, otherwise translating all the q :s and μ :s an equal distance would just result in the same thing. My idea is now to try to find q :s and μ :s s.t. the probabilities of the Z :s migrating to the different rating classes match the probabilities in the historical migration matrix. Given all the q :s and μ :s the probability of migration from one rating class to any other can be calculated. For example, the the probability of migration from BBB to B+ is

$$P(q_B \leq Z_{BBB} < q_{B+}) = F_{Z_{BBB}}(q_{B+}) - F_{Z_{BBB}}(q_B).$$

These probabilities can then be used to construct a new, re-estimated migration matrix. This method will have the benefit of keeping every row sum in the matrix equal to one and as well every element strictly greater than zero. In addition some order relations will automatically be fulfilled, for example that the probability of default is lower for better ratings. Unfortunately, Fitch does not include what number of obligors that has been used to calculate each percentage in the historical migration matrix. If this had been the case it would have been possible to treat the historical migration matrix as pure data and I could probably had come up with some Maximum Likelihood estimator, given this model. I would then find the q :s and μ :s that were most likely to generate the data.

Now, I will instead use a least squares method that minimizes the sum of the squared differences between each element in the historical migration matrix and each element in the migration matrix implied by the q :s and μ :s (except the last row). I use the built in Matlab function `lsqnonlin`, and as starting values I use $\mu_C = 0.5, q_C = 1, \mu_{B-} = 1.5, q_{B-} = 2 \dots$ and so on. To make sure that the order condition (4.3) is fulfilled I add an extra penalty to the function if it is not. The function runs in a few seconds. A numerical test of this method is presented in the following example.

Example 3. In this example I will use a historical migration matrix from Fitch with only eight different rating classes to make it easier to look at the result. The least squares method is used to estimate a new migration matrix. The results are presented in table 3.

From/To	AAA	AA	A	BBB	BB	B	C	D
AAA	0.9862	0.0138	0	0	0	0	0	0
AA	0.0379	0.9242	0.0273	0.0095	0	0.0047	0	0
A	0	0.0273	0.9235	0.0437	0.0055	0	0	0
BBB	0	0	0.0631	0.8931	0.0340	0.0049	0.0049	0
BB	0	0	0	0.0939	0.8449	0.0490	0	0.0122
B	0	0	0	0	0.1000	0.8684	0.0263	0.0053
C	0	0	0	0	0	0.2500	0.5000	0.2500
D	0	0	0	0	0	0	0	1

From/To	AAA	AA	A	BBB	BB	B	C	D
AAA	0.9862	0.0138	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
AA	0.0384	0.9375	0.0240	0.0000	0.0000	0.0000	0.0000	0.0000
A	0.0000	0.0275	0.9286	0.0439	0.0000	0.0000	0.0000	0.0000
BBB	0.0000	0.0000	0.0637	0.9019	0.0343	0.0000	0.0000	0.0000
BB	0.0000	0.0000	0.0000	0.0951	0.8553	0.0496	0.0000	0.0000
B	0.0000	0.0000	0.0000	0.0000	0.1004	0.8719	0.0272	0.0006
C	0.0000	0.0000	0.0000	0.0000	0.0001	0.2507	0.4985	0.2507
D	0	0	0	0	0	0	0	1

	μ_C	q_C	μ_B	q_B	μ_{BB}	q_{BB}	μ_{BBB}	q_{BBB}	μ_A	q_A	μ_{AA}	q_{AA}	μ_{AAA}
Starting value	0.5	1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5
Estimates	0.7	1.3	3.3	4.5	6.2	7.5	9.3	10.8	12.6	14.5	16.4	18.2	20.4

Table 3: Tables for example 3. The top table shows the historical migration matrix. The middle table shows the re-estimated migration matrix and the bottom table shows the starting values and estimates of the q :s and μ :s.

As can be seen in the results from example 3, the new migration matrix is not very different from the old one. The biggest difference is that the probability to move from rating AA to rating B has been significantly lowered as well as for some of the default probabilities and the probability mass has been moved towards the diagonal. I have tried this method for some other migration matrices as well and have got more or less the same result. Small probabilities outside of the main diagonal has been more or less erased. It is hard to say if the new matrix is more or less realistic than the old one. The new matrix does fulfill order relations between ratings, but this has been done by basically erasing probabilities that violates it. The new matrix also has positive probabilities for every migration, but the ones that are far from the diagonal are so small that they do not make any numerical difference anyway². One can also question how likely the observed historical migration matrix is given the new estimated one. One issue with this least squares method is that it treats all probability

²The PD for AAA-rated obligors in example 3 is even smaller than the smallest number Matlab can handle $\approx 10^{-325}$ so it is treated as identically zero.

differences equally. If the historical matrix would for example show a 0.1% PD for AAA this would be a significant risk factor, but the method would just erase this probability rather than raising the probabilities for downgrades. As a fix to this I tried to come up with a way to weigh the elements of the matrix differently in the least squares function, but I was not able to make much improvement.

For the calculations below I will use the historical matrix in fig 4.2. For comparison, I will also use the same matrix but with the PD :s tweaked in the same way as in the ECB paper.

4.4.2 Conditional credit losses

The conditional credit loss (CL_i) indicates how big the credit loss is for an investment in a single obligor i given the obligors rating one year from now. As defined in (4.1), the CL_i depends on four things:

- The time points of the future cash flows (t_j)
- The sizes of those cash flows (CF_j)
- The discount factors ($df_{cr(i)}$) associated with the different credit ratings
- The recovery rate (RR).

I have let $cr(i)$ denote the credit rating of obligor i . The first two points are usually known exactly, since that is how bonds work. Every bond pays specified amounts on specified payment dates. For the Riksbank's portfolio, however, I have not been able to get hold of this information, probably because the portfolio is very big and contains a lot different bonds that are registered in different systems, so it is a technical problem to get all t_j and CF_j in a single list. Instead, I have been given the current total market value of each position converted to SEK (TV_i) and the modified duration (MD_i) of each position. The MD_i is the center of mass of the future cash flows along the time axis and could be defined as

$$MD_i = \frac{\sum_j CF_j t_j}{\sum_j CF_j}.$$

Given the limitation of the data, I will instead view it as if there was only one future cash flow from each obligor, having a size that equals TV_i if discounted by $df_{cr(i)}$. Formally,

$$TV_i = (df_{cr(i)})^{MD_i} CF_{MD_i},$$

where CF_{MD_i} is the cash flow replacing the old ones.

The models assume that there exists a fixed discount factor for every rating class that can be used to discount all future cash flows. Looking at market bond yield data one sees that this assumption is not perfectly true. The yield can vary quite much for countries in the same rating class and for different maturities. The yield can also be lower for lower rated countries, for example on

the 1st of July 2012 the yield for Japanese 10-year bonds was 0.83% while the corresponding value for French bonds was 2.69%, even though France is rated AAA and Japan is rated A+ by Fitch³. This is non-intuitive, but is a sign of that the market cares about other things than Fitch's ratings and credit risk. However, this might not be such bad news for our model anyway, since we are mainly interested in the change of market value, when the ratings change. If France was downgraded to Japan's rating, we can be pretty sure that the yield would not decrease to the same yield as Japan, but instead increase due to the higher risk associated with the lower rating. Therefore, a good way to estimate the conditional credit losses might had been to look at how much the market value of different bonds has changed after rating changes, but I do not have that kind of data. Instead, I just choose different values for the different yields in an ad hoc way such that they reasonably match the current market yields of bonds from countries with the same rating and with maturities that are close to the *MD*:s. Each discount factor is then calculated as one divided by the corresponding yield.

For the recovery rates, there is very little data, since very few countries has defaulted during the last 15 years. The complete list can be seen in table 4. It is hard to come up with a single number for the recovery rate, based on this data. As can be seen the recovery rates varies a lot and the countries in this data set is generally of much lower rating than the ones in the Riksbank's portfolio. For the model I use something in the middle, so I let $RR = 0.55$.

³Source: www.tradingeconomics.com.

Year	Defaulting Country	Recovery rate (%)
1998	Russia	50
1999	Pakistan	65
1999	Ecuador	60
2000	Ukraine	60
2000	Ivory Coast	NA
2001	Argentina	30
2002	Moldova	95
2003	Uruguay	85
2003	Nicaragua	50
2004	Grenada	NA
2005	Dominican Republic	95
2006	Belize	NA
2008	Seychelles	NA
2008	Ecuador	NA
2010	Jamaica	80
	Issuer-Weighted RR	67
	Value-Weighted RR	36

Table 4: Recovery rates on sovereign defaults during the last 15 years. The last two rows shows the average RR weighted per issuer and weighted by the total value of bonds respectively. Source: Moody's.

I now have everything needed to calculate the conditional credit losses. Since there was some uncertainty in all the parameters used to calculate the CL 's it will probably result in some model error and to make this part of the model satisfactory, better data is needed. It will be important to update the CL 's frequently as for example the MD 's change.

4.4.3 Correlations

The correlation between obligors' underlying assets is an important parameter for the risk measures, since a high correlation makes it more probable that countries get downgraded at the same time. These events will result in the largest portfolio losses. The correlations are typically positive, since countries usually react similar to changes in the world economy. There exist some known methods to estimate the correlations in the correlation matrix Σ . In (Gupton et al., 1997), which mainly focuses on corporate bond, they look at the companies' stocks time series to make an approximation of the companies' assets and estimate the correlation from this. They also group the companies into different groups, depending on sector and geographical location and estimate the correlations in the different group. They then come up with a clever way to base the correlation between two companies on the groupings.

Since our portfolio consist mainly of sovereign issuers, we cannot use the same

technique. Countries do not issue stocks, so we cannot use that method to estimate the correlations. Maybe the method with groupings could be used, example by raising the correlation between countries in the euro zone, but we still need something to base the estimations on. The ECB approach is to just use 24% for all correlations, since this is suggested by Basel II⁴ to be an upper limit for the correlation. This method would result in estimates of the risk measures that will be more like upper limits of the risk measures.

I will try a different method for estimating the correlations, based on CDS spread data and inspired by Friewald (2009). A Credit Default Swap (CDS) is a credit derivative that works like an insurance against default. An illustration of the CDS contract can be seen in figure 4.3.

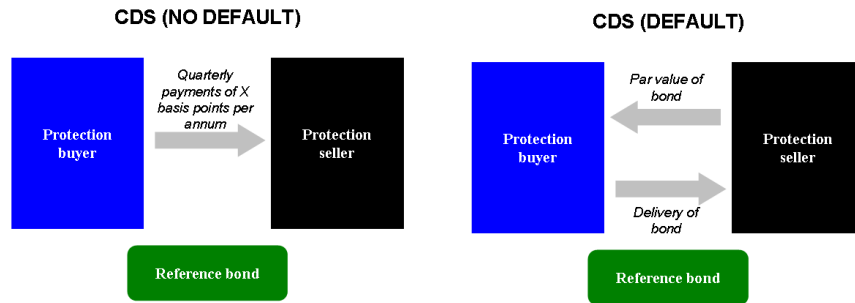


Figure 4.3: An illustration of how a CDS works.

For every CDS contract, there will exist a reference bond. The buyer of the CDS contract, will pay an insurance rate s per year of the value V of the bond to the protection seller. If the issuer of the bond defaults, the protection seller pays the buyer an amount equal to the value of the bond and receives the bond in return, retrieving whatever the defaulted bond pays. In this way, in case of a default the protection buyer reduces his losses by $V(1 - RR)$, where RR is the recovery rate. CDS:s are quoted daily which means that there exists market prices, quoted in the insurance rate s . One can thus assume that the prices will be fair. At least one can assume that the market is liquid enough to make prices for CDS:s on different bonds equally close to their fair price and since our main intention is to estimate correlations, this approximation should be adequate. If we assume a constant hazard rate, i.e. the probability of defaulting is the same at any time point in the next year, then a fair price of the CDS would be such that the insurance rate s equals the hazard rate λ of the bond multiplied by $(1 - RR)$, i.e.

$$\lambda = \frac{s}{1 - RR}.$$

⁴Basel II are recommendations on banking laws and regulations issued by the Bank of International Settlement.

Hence, the probability of default in one years time should be

$$PD = 1 - \exp(-\lambda \cdot 1) = 1 - \exp\left(-\frac{s}{1 - RR}\right).$$

Given the recovery rate RR (we will assume as before that $RR = 0.55$) we now have a way of estimating the probability of default in one year's time based on the market CDS prices. The market CDS prices will of course only reflect investors' opinions on what the PD is and not the true PD itself, but this can be seen as an approximation.

As input to this method I will use daily CDS prices from the time period 2008-2012 for 5-year bonds from some different countries. I denote the price of a CDS $s_i(t)$ for obligor i at time t . Similarly I denote the PD for the next year implied by the CDS price $PD_i(t)$. The prices are quoted about 260 times per year, so I let the length of one day be $d = \frac{1}{260}$ and $n_d = 260$. Now, recall the model with the underlying process that we used to construct the multi-dimensional model. To be able to use the CDS data for estimating correlation I will for each obligor i define an underlying process $Z_i(t)$ by

$$Z_i(t + d) = Z_i(t) + \varepsilon_i(t + d),$$

where $\varepsilon_i(t) \in N(0, d)$ and independent for $t = d, 2d, 3d, \dots$. This implies $Z_i(t + 1) - Z_i(t) = \sum_{k=1}^{n_d} \varepsilon_i(t + kd) \in N(0, 1)$. I will as before assume that there exists fixed constants q_D^i for each obligor such that if the process is below q_D^i a year from now the obligor defaults. Thus

$$\begin{aligned} & \begin{cases} P(Z_i(t + 1) < q_D^i | s_i(t)) & = PD_i(t) \\ P(Z_i(t + d + 1) < q_D^i | s_i(t + d)) & = PD_i(t + d) \end{cases} \Leftrightarrow \\ & \begin{cases} P\left(\sum_{k=1}^{n_d} \varepsilon_i(t + kd) < q_D^i - Z_i(t) | s_i(t)\right) & = PD_i(t) \\ P\left(\sum_{k=2}^{n_d+1} \varepsilon_i(t + kd) < q_D^i - Z_i(t + d) | s_i(t + d)\right) & = PD_i(t + d) \end{cases} \Leftrightarrow \\ & \begin{cases} q_D^i - Z_i(t) & = \Phi^{-1}(PD_i(t)) \\ q_D^i - Z_i(t + d) & = \Phi^{-1}(PD_i(t + d)) \end{cases} \Rightarrow \\ & Z_i(t + d) - Z_i(t) = \Phi^{-1}(PD_i(t)) - \Phi^{-1}(PD_i(t + d)). \end{aligned} \quad (4.4)$$

By using (4.4), we have a method to estimate the increments of the underlying process. Using this method for different obligors will generate different processes which can be used to estimate the correlation between them. What we are interested in is the correlation ρ between two assets over a one year time period.

$$\begin{aligned} \rho &= \text{corr}(Z_i(t + 1) - Z_i(t), Z_j(t + 1) - Z_j(t)) = \text{Cov}(Z_i(t + 1) - Z_i(t), Z_j(t + 1) - Z_j(t)) = \\ & \text{Cov}\left(\sum_{k=1}^{n_d} \varepsilon_i(t + kd), \sum_{k=1}^{n_d} \varepsilon_j(t + kd)\right) = \sum_{k=1}^{n_d} \text{Cov}(\varepsilon_i(t + kd), \varepsilon_j(t + kd)) = n_d \text{Cov}(\varepsilon_i(t), \varepsilon_j(t)) \end{aligned}$$

$$n_d \text{corr}(\varepsilon_i(t), \varepsilon_j(t)) \cdot \sqrt{\text{Var}(\varepsilon_i(t)) \text{Var}(\varepsilon_j(t))} = n_d \text{corr}(\varepsilon_i(t), \varepsilon_j(t)) \cdot d = \text{corr}(\varepsilon_i(t), \varepsilon_j(t)) \\ = \text{corr}(Z_i(t+d) - Z_i(t), Z_j(t+d) - Z_j(t)).$$

This means that the correlation can be estimated by simply using the standard estimator for correlation on the asset increment time series.

	Germany	France	Greece	Japan	New Zealand	Sweden	USA
Germany	1	0.73	0.26	0.09	0.13	0.57	0.41
France	0.73	1	0.31	0.11	0.14	0.61	0.42
Greece	0.26	0.31	1	0.06	0.04	0.23	0.18
Japan	0.09	0.11	0.06	1	0.06	0.09	0.09
New Zealand	0.13	0.14	0.04	0.06	1	0.12	0.12
Sweden	0.57	0.61	0.23	0.09	0.12	1	0.34
USA	0.41	0.42	0.18	0.09	0.12	0.34	1

Table 5: Asset correlations between some countries estimated from CDS prices.

Table 5 shows a few correlation estimates. The estimated correlations seem reasonable in the sense that countries that are closer to each other geographically have a higher correlation, but some correlations seem maybe too high in the light of the Basel II suggestion of a maximum correlation of 24%. One reason to this could be that the data is taken from a very special time period, in the aftermath of the financial crisis. However, since we have earlier assumed that the correlation is constant over time, this would be inconsistent with the model.

To test whether the constant correlation is a good model I test what happens if the correlation is estimated just using one year of data at a time. I use a sweeping time window of length 260 and make a new estimation for every time lag. How the correlation changes over time for some different combinations of countries is plotted in fig 4.4. As can be seen the correlation seems to be changing quite much over time, which indicates that either the assumption of the constant correlation is bad or the method of estimating the correlation is.




Figure 4.4: Correlations estimated from CDS data using a one year time window for some different combinations of countries.

Another problem when using a correlation matrix estimated like this is that there is no guarantee that the matrix is positively semi-definite, which is required for a correlation matrix. To be able to use the matrix for calculations, this must be fixed in some way. For this reason, and because I do not have CDS data for all obligors in the Riksbank's portfolio, I will just use the ECB ad hoc approach with all correlations equal to 24% from now on.

5 Results and Validation

In this section I will show some calculations done based on the model, with some different parameters to investigate how these affect the model. I would very much like to validate the model by looking at real credit loss data. The problem is of course that the risk measures, VaR_α and ES_α , states something about what happens in extreme cases and therefore requires a huge amount of data to be validated. For example, if we have an estimation of $VaR_{0.999}$ that is too high we would need 5000 years of data with no losses greater than our estimation to be able to reject it statistically with 95% significance! Needless to say, this is impossible and instead I will try to estimate how wrong the model could be, by looking at the input data.

When calculating the risk measures I will look at two different portfolios. One real portfolio that the Riksbank held some time ago consisting of $n = 29$ positions of high rating, 14 of them AAA. The portfolio is chunky, a few large positions represent more than 80% of the total value. I will call this portfolio “portfolio A”. The other portfolio is completely fictional and is given the same n , value and durations as portfolio A. I will call this portfolio “portfolio B” and it will contain positions of equal value and of all different ratings.

To begin with, I looked at the loss distribution for the different portfolios using the standard Monte Carlo method, the historical migration matrix and the correlation matrix with all off-diagonal correlations equal to 24%. The result is shown in fig 5.1. The loss distribution is much wider for portfolio B as expected since it contains much lower credit quality. One can clearly see jumps in portfolio A’s tail, which comes from the chunkiness, while portfolio B’s tail is much more smooth.




Figure 5.1: Histograms of simulated losses and an approximation of the probability function for portfolio A and B respectively. The figures on the right are close-ups on the tails of the distributions, showing the 5% largest losses.

I continued with estimating the different risk measures for the portfolios, for some different values of α . The results can be seen in table 6 and fig 5.2.

Observe that portfolio B has lower EL than portfolio A, but that VaR and ES are about 20 times higher. Observe also how VaR and ES increase rapidly as α approaches 1.

% of portfolio value	EL	UL	$VaR_{0.999}$	$ES_{0.999}$
Portfolio A	-0.0091	0.0281	0.3016	0.3453
Portfolio B	-1.0325	1.4054	6.0330	7.0727

Table 6: Risk measures for the two portfolios, calculated using standard Monte Carlo and the historical migration matrix.




Figure 5.2: VaR and ES for the two portfolios and for some different values of α , calculated using standard Monte Carlo and the historical migration matrix.

Next, I looked at the uncertainty of the risk measures that comes from the Monte Carlo methods. In fig 5.3 I have plotted the 95% confidence intervals for $VaR_{0.999}$ and $ES_{0.999}$ as functions of the sample size. As can be seen the intervals are much narrower when using IS. It should be mentioned, however, that the calculations take about twice the time when using IS, but the interval at half the sample size for IS is much narrower than that for full sample size standard Monte Carlo. Since the biggest problem is memory rather than time, IS is definitely to prefer. The sizes of the confidence intervals for IS are equivalent to an error of less than one percent, which is acceptable. The model error and errors that come from the parameters are probably much larger.

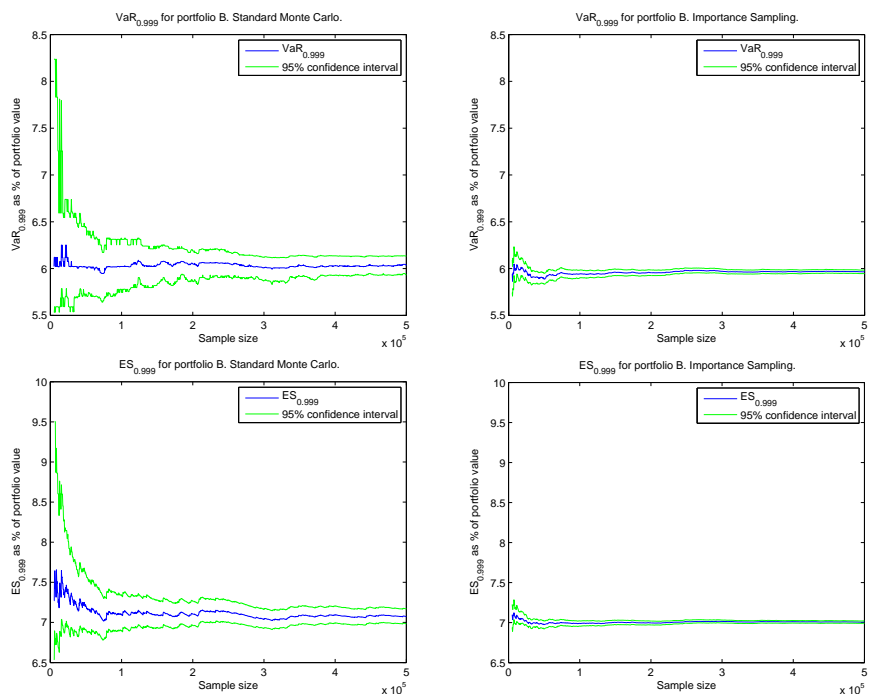


Figure 5.3: Comparison of the sizes of 95% confidence intervals for $VaR_{0.999}$ (above) and $ES_{0.999}$ (below), when using standard Monte Carlo (left) and Monte Carlo with importance sampling (right).

I wanted to validate the model by looking at how big the error from the parameters could be. I started with the migration matrix. First, I looked at the method proposed by ECB where they tweak the PD:s to numbers that they think are more reasonable. Table 7 shows the same results as table 6, but with the tweaked migration matrix. As can be seen, all risk-measures increase some, especially ES which increases about 20%.

% of portfolio value	EL	UL	$VaR_{0.999}$	$ES_{0.999}$
Portfolio A	-0.0091	0.0423	0.3042	0.4036
Portfolio B	-0.9064	1.5724	7.2338	8.4727

Table 7: Risk measures for the two portfolios, calculated using the tweaked migration matrix.

Next, I wanted to see how wrong the historical migration matrix could be from a more statistical point of view. The idea was to find a “worse” migration matrix that was still not too unlikely to have generated the data. As mentioned before, Fitch does not say how big the data set is that their migration matrices are based on, but I found a migration matrix from S&P with 8 rating classes

that has this included. For validation purposes I therefore made this analysis for the S&P matrix instead, assuming that the error for the Fitch matrix would be of roughly the same size. The S&P data can be seen in table 8.

From/To	AAA	AA	A	BBB	BB	B	C	D	n
AAA	455	13	0	0	0	0	0	0	468
AA	9	262	6	0	1	1	0	0	279
A	0	12	272	14	0	0	0	0	298
BBB	0	0	16	230	9	2	0	0	257
BB	0	0	0	20	275	15	4	2	316
B	0	0	0	0	23	256	9	5	293
C	0	0	0	0	0	7	7	8	22

Table 8: Migration history from S&P. Each element in the n-column shows the sum of the corresponding row in the matrix.

I denoted the matrix H , the elements in the matrix n_{ij} and elements in the n-column n_i for $i = 1, \dots, 7$, $j = 1, \dots, 8$. Each element p_{ij} in the usual historical migration matrix would then be estimated as $\hat{p}_{ij} = \frac{n_{ij}}{n_i}$, which is in fact the ML-estimator (Rydén and Lindgren, 2000, pp. 62-66). I would now like to construct a confidence interval for the elements p_{ij} . The idea could then be to use the upper bounds for the elements to the right of the diagonal and lower bounds for the elements to the left of and on the diagonal. This would create a migration matrix that would generate larger losses. The elements n_{ij} can be seen as outcomes of binomially distributed variables $X_{ij} \in Bin(n_i, p_{ij})$. In the same way each row in the matrix can be seen as outcomes of a multinomial distribution. Confidence intervals based on the normal distribution can unfortunately not be constructed, since the normal approximation is invalid for the elements in the matrix where $n_{ij} = 0$. Instead I used the binomial distribution directly by implementing a numerical method in Matlab and looked for numbers p_{ij}^U and p_{ij}^L such that

$$\max \left(p_{ij}^U | P \left(\tilde{X}_{ij} \leq n_{ij} \right) \geq \alpha \right), \quad \tilde{X}_{ij} \in Bin(n_i, p_{ij}^U)$$

$$\min \left(p_{ij}^L | P \left(\tilde{X}_{ij} \geq n_{ij} \right) \geq \alpha \right), \quad \tilde{X}_{ij} \in Bin(n_i, p_{ij}^L),$$

for some small number α . p_{ij}^U and p_{ij}^L could then be used as confidence bounds for each element⁵. There are two problems with the approach of using the bounds to construct a new matrix. The first problem is how to make sure that each row sums to one. The second problem is how to treat the fact that I am looking at many variables at the same time. If I would use the confidence bound for a small α for every p_{ij} , then each n_{ij} would not be too improbable, but all of them at the same time would be. Therefore I compared the probability of each

⁵This will be one-sided confidence intervals, for different sides depending on what side of the diagonal the element is on.

row of H given the new matrix to the probability of each row in H given the old matrix⁶. Since each row of the matrix in table 8 is multinomially distributed, the probability function for row i is

$$mnpdf_i(n_{i,:}, p_{i,:}) = \frac{n_i!}{n_{i1}!, \dots, n_{i8}!} p_{i1}^{n_{i1}} p_{i2}^{n_{i2}} \cdots p_{i8}^{n_{i8}}.$$

The ratios r_i between the probability of each row given the new and the old matrix with elements p_{ij}^{new} and p_{ij}^{old} can thus be calculated as

$$r_i(p_{i,:}^{new}, p_{i,:}^{old}) = \frac{mnpdf_i(n_{i,:}, p_{i,:}^{new})}{mnpdf_i(n_{i,:}, p_{i,:}^{old})}.$$

The nice thing here is that a lot of factors cancel (all that are taken factorial) and r_i is actually not that troublesome to calculate numerically, especially if one calculates the logarithm of the expression first. To get an idea of what a suitable size is for the ratios I compared to what the corresponding value is for 99%-confidence bounds for the standard Gaussian distribution.

$$\frac{\varphi(q_{0.99})}{\varphi(0)} = 0.067,$$

where φ is the standard Gaussian distribution and $q_{0.99}$ the 99%-quantile. I created a new matrix by using different α :s for each element and by simply normalizing each row to make it sum to one (the row sum for each row was already close to one so this should not make a big difference), making the ratio for each row close to 0.067. This required α :s in the interval 45 – 95% for the different rows. The matrix can be seen in table 9 under the name M_{CB1} .

I constructed one more confidence bound for the migration matrix, this time only focusing on the diagonal elements and the PD:s. Concerning only two elements in each row I could see them as binomially distributed and I did not have to use the messy multinomial distribution. This way I was also able to use higher confidence levels for the single elements. I used the ML estimation of the migration matrix as a base and moved probability mass from the diagonal elements to the PD:s according to a 99% binomial confidence interval for the two probabilities. The resulting matrix can also be seen in table 9 under the name M_{CB2} .

⁶What I really wanted to do was to compare the probability of each row of H or worse, but the cumulative distribution function for the multinomial distribution becomes very messy even for only 8 dimensions. The method I use instead can be seen as an approximation of this method.

M_{ML}	AAA	AA	A	BBB	BB	B	C	D
AAA	97.22	2.78	0	0	0	0	0	0
AA	3.23	93.91	2.15	0	0.36	0.36	0	0
A	0	4.03	91.28	4.70	0	0	0	0
BBB	0	0	6.23	89.49	3.50	0.78	0	0
BB	0	0	0	6.33	87.03	4.75	1.27	0.63
B	0	0	0	0	7.85	87.37	3.07	1.71
C	0	0	0	0	0	31.82	31.82	36.36

M_{CB1}	AAA	AA	A	BBB	BB	B	C	D
AAA	96.44	2.80	0.13	0.13	0.13	0.13	0.13	0.13
AA	2.94	92.43	2.47	0.28	0.65	0.65	0.28	0.28
A	0	3.85	90.29	4.91	0.24	0.24	0.24	0.24
BBB	0	0	5.51	88.17	4.21	1.29	0.41	0.41
BB	0	0	0	5.02	85.12	6.21	2.22	1.43
B	0	0	0	0	6.12	86.06	4.74	3.07
C	0	0	0	0	0	19.40	19.40	61.20

M_{CB2}	AAA	AA	A	BBB	BB	B	C	D
AAA	96.24	2.78	0	0	0	0	0	0.98
AA	3.23	92.28	2.15	0	0.36	0.36	0	1.63
A	0	4.03	89.75	4.70	0	0	0	1.53
BBB	0	0	6.23	87.72	3.50	0.78	0	1.77
BB	0	0	0	6.33	85.03	4.75	1.27	2.63
B	0	0	0	0	7.85	84.68	3.07	4.40
C	0	0	0	0	0	31.82	12.24	55.94

Table 9: The first table shows the ML-estimation of the migration matrix. The second shows a confidence bound for the matrix using the method with the multinomial distribution above. The last table shows another confidence bound that only focuses on the PD:s.

I calculated some risk measures based on the matrices in table 9, which are presented in table 10. To be able to do this I first changed the ratings in the portfolio to match S&P rating system. As can be seen, the risk measures become very large when using the confidence bound migration matrices. This was expected, since the risk measures says something about what happens in extreme cases, they basically say what happens when most of the obligors default using the new PD:s. The reason why portfolio B's risk measures are lower is probably because it is more diversified than portfolio A.

% of portfolio value	A $VaR_{0.999}$	A $ES_{0.999}$	B $VaR_{0.999}$	B $ES_{0.999}$
M_{ML}	0.25	0.30	6.99	8.22
M_{CB1}	23.53	24.428	12.00	14.31
M_{CB2}	26.64	32.15	17.31	20.29

Table 10: Risk measures calculated for portfolio A and B using the different migration matrices in table 9.

To estimate how large the error from the conditional losses could be is very hard because of lack of data. I did however do a short test to see how the recovery rates affect the risk measures. I looked at recovery rates in the interval of the recovery rate data in table 4. The table shows that for portfolio A, that contains high quality credit, the recovery rates makes a very small difference, but for portfolio B it is quite decisive.

% of portfolio value	A $ES_{0.999}$ hist	B $ES_{0.999}$ hist	A $ES_{0.999}$ tweaked	B $ES_{0.999}$ tweaked
$RR = 30\%$	0.34	10.22	0.44	12.33
$RR = 45\%$	0.34	8.29	0.42	10.02
$RR = 60\%$	0.34	6.37	0.40	7.70
$RR = 75\%$	0.34	4.51	0.38	5.42
$RR = 90\%$	0.34	2.80	0.36	3.24

Table 11: Shows how $ES_{0.999}$ varies for portfolio A and B for different levels of recovery rates and for the historical and tweaked migration matrix.

Finally, I did a test of how the risk measures are affected by the correlations. I used a correlation matrix with the same value on every off-diagonal position and let it vary between 0 – 99%. The result is presented in figure 5.4. It is clear from the plots that the risk increases with correlation. The maximum correlation suggested by Basel II, 24%, gives around twice the $ES_{0.999}$ as the uncorrelated case. If the correlations I estimated from CDS:s is correct the true $ES_{0.999}$ could be even higher.




Figure 5.4: $ES_{0.999}$ as a function of correlation, for portfolio A and B and for the historical and tweaked migration matrix.

6 Optimization

Given the model and the possibility to calculate risk measures for different portfolios, it is natural to wonder what portfolio one should hold to minimize the risk. This topic will be discussed in this section and will be based on (Rockafellar and Uryasev, 2000). As we saw in example 2, VaR can lead to absurdities as a risk measure. We will therefore focus on trying to find the portfolio with the lowest ES . Since $VaR < ES$, optimizing with respect to ES will give us a portfolio with a low VaR as well. We consider the case when there are n obligors that it is possible to buy bonds from. Define the portfolio x as

$$x = [x_1, x_2, \dots, x_n].$$

The optimization problem that we want to solve can be formulated as

$$\begin{aligned} \min_x \quad & ES_\alpha(x) \quad s.t. \\ \sum_{i=1}^n x_i & = B \\ 0 \leq x_i & \leq c_i, \forall i, \end{aligned}$$

for some fixed budget B and upper limits c_i .

Since the random credit loss L_i for each position i in the portfolio is directly proportional to the market value of that position TV_i , we now define the ratio between them to be $Y_i = \frac{L_i}{TV_i}$. The column vector containing all ratios is denoted Y and we denote its density function $f_Y(y)$. We can now define the portfolio credit loss as the function $L(x, Y) = xY$, where x is the portfolio. The probability of the loss being smaller than a certain value γ is defined by the function

$$\Psi(x, \gamma) = P(L(x, Y) < \gamma) = \int_{L(x, y) < \gamma} f_Y(y) dy.$$

For a fixed x this can be seen as the distribution function of losses. In (Rockafellar and Uryasev, 2000), they assume Ψ to be continuous everywhere with respect to γ . This is of course not always true in our case. Consider for example a portfolio consisting of a single bond, then Ψ will have jumps in the points where gamma equals the conditional credit losses for different rating changes. For larger, more well-diversified portfolios however, the jumps will be smaller and Ψ closer to continuous in some sense. The theory the authors present is based on the assumption of Ψ being continuous, but they hint that the theory is in fact valid even without this assumption, by saying that it is only made for simplicity. It might be possible to change my model so that Φ would be continuous, by for example replacing the migration matrix and the conditional credit losses with some continuous equivalent, using curve-fitting. For now, I will however just assume that the theory works anyway.

Now recall the expression for ES_α given in (3.5). In (Rockafellar and Uryasev, 2000), they define a similar function

$$F_\alpha(x, \gamma) = \gamma + \frac{1}{(1 - \alpha)} \int_{y \in \mathbb{R}^n} \max(L(x, y) - \gamma, 0) f_Y(y) dy$$

and then prove the following theorem.

Theorem 4. *The expected shortfall $ES_\alpha(x)$ for any fixed portfolio $x \in \mathbb{R}^n$ can be obtained as*

$$ES_\alpha(x) = \min_{\gamma} F_\alpha(x, \gamma).$$

Similarly, $VaR_\alpha(x)$ can be obtained as the smallest value γ that minimizes $F_\alpha(x, \gamma)$.

This gives us a new way of estimating ES_α without having to know VaR_α first. In the IS Monte Carlo setting the estimator would be

$$\widehat{ES_\alpha}(x) = \min_{\gamma} \left(\gamma + \frac{1}{(1 - \alpha)N} \sum_{i=1}^N \max(xY_i - \gamma, 0) \omega(Y_i) \right),$$

where Y_i is now the i :th simulated sample from an instrumental distribution $g_Y(y)$. The authors then go on proving another theorem that will be of great use.

Theorem 5. *Minimizing $ES_\alpha(x)$ over x is equivalent to minimizing $F_\alpha(x, \gamma)$ over all (x, γ) , i.e.*

$$\min_x ES_\alpha(x) = \min_{(x, \gamma)} F_\alpha(x, \gamma).$$

Furthermore, $F_\alpha(x, \gamma)$ is convex with respect to (x, γ) and $ES_\alpha(x)$ is convex with respect to x when $L(x, y)$ is convex with respect to x .

$L(x, y) = xy$ is obviously convex with respect to x also with the constraints $0 \leq x_i \leq c_i, \forall i$. Thus, we have convex optimization problem. This means that if we can find a local minimum, we know that we have found a global minimum. We will once again use Monte Carlo methods to approximate the loss distribution. This could of course result in non-optimal results, because the simulated sample may for example contain unexpected high losses for some obligor, resulting in a lower investment in that position than optimal. Therefore it is important to use IS and a large N . The optimization problem can now be formulated as

$$\min_{(x, \gamma)} \left(\gamma + \frac{1}{(1 - \alpha)N} \sum_{i=1}^N \max(xY_i - \gamma, 0)\omega(Y_i) \right), \text{ s.t.} \quad (6.1)$$

$$\sum_{i=1}^n x_i = B$$

$$0 \leq x_i \leq c_i, \forall i.$$

The authors suggest some further reading on methods on how to solve these kind of optimization problems numerically, but for a project of this size it would be too time consuming to understand and implement a new method, so I will instead try to use a method that I am familiar with since before and which will be described in the next section: a genetic algorithm. I also put some time on trying to solve the problem using Matlab's built in optimization toolbox without being able to reach convergence.

6.1 Genetic algorithm

A genetic algorithm is an optimization method that is inspired by the theory of evolution. The idea is to treat solutions to the optimization problem as individuals are being treated in nature: bad solutions are killed while good solutions live on and new solutions are created by randomly mixing the old good ones. Just like evolution has created well fitting species for this world the hope is that this method will produce optimal solutions for the optimization problem. I will now explain briefly how I have implemented a genetic algorithm for this problem.

First of all I constructed a pool of M portfolios x^1, \dots, x^M that formed the first generation of the population. These portfolios all fulfilled the constraints of the problem and were constructed in two different ways. Some portfolios were created as extreme case, for example portfolios that have the whole budget in a single position. The other portfolios were created randomly, using a uniform distribution to randomize the sizes of the different positions. γ was initiated as the estimated VaR_α of a random portfolio. Each portfolio was then evaluated in the function that is being minimized in (6.1) and the results were saved in the weight variables w_1, \dots, w_M . I had generated the first generation. The plan was then to construct a new, second generation out of the first one, then a third generation and so on until the weights no longer improved. The construction of

each new generation consisted of four steps: crossing, mutation, selection and update of γ .

- In the crossing step the portfolios were randomly paired two and two. Each pair constructed a new portfolio by randomly choosing one of the old portfolios' investment in each position for the new portfolios' investment in that position. The new portfolios were then normalized so that the total investment equals the budget.
- In the mutation step all positions in all portfolios, except for a few of the best ones, were increased or decreased randomly by a few percent. A variable made sure that the mutations became smaller the longer the algorithm ran. Afterwards it was ensured that the constraints were fulfilled.
- In the selection step the 50% of the old portfolios with the highest weights were discarded and replaced by the new ones and new weights were calculated. This became the next generation.
- A new γ was calculated as the estimate of VaR_α for the best portfolio.

I implemented the algorithm in Matlab and tested it for some different sets of possible positions, migration matrices and constraints. The first thing I noticed was that the algorithm was quite slow and it had to run for a long time to get close to anything that could be considered convergence. Given the complexity of the model, it was difficult to test the algorithm on something that had a known solution. One possibility was, however, to run it in the case when all obligors are identical. The positive correlation among obligors should then make it optimal to hold a diversified portfolio with equally much invested in each position. I therefore tested the algorithm with 10 identical fictional AAA-rated obligors, with the ECB-tweaked Fitch migration matrix, $\alpha = 0.95$ and $c_i = \infty$ for all i . The result was interesting. About 7 of the obligors were being treated as equally good by the algorithm and got the same investment, about one seventh of the budget. The other three got close to zero. I tested the same setup a few more times and got about the same result, except that different obligors were treated as bad every time. I compared the solution to the best theoretical solution, that each obligor gets the same. It turned out that the ES of the optimal portfolio from the algorithm was lower than for the optimal theoretical solution. My best explanation to this result is that the Monte Carlo-approximation of the loss distribution contained too many big losses for some of the positions, resulting in that these positions were reduced by the algorithm. Hence, for the Monte Carlo-approximation to be sufficient we would have raise the number of simulations N or come up with a smarter instrumental distribution for the IS. The first option is ruled out, since my computer runs out of memory for $N \approx 1$ million.

I also tested the algorithm for some real obligors and the original Fitch matrix (fig 4.2). The result was that the optimal portfolio was the one containing only obligors with rating A-. Looking at the migration matrix one understands

why. Obligors of rating A- has historically only been downgraded by one step maximally and the ES is therefore low. This result says more about the absurdity of using historical data only in the migration matrix.

Finally I did a test to check the convergence. I ran the algorithm twice for $n = 34$, with the same Monte Carlo-approximation of the loss distribution both times. Unfortunately, the result was not what I hoped for. The experiments generated two portfolios that were pretty similar, they differed less than 20% on each position. It seemed clear however that both portfolios represented some different local minima, because the optimum had been constant in both runs for several generations. A plot of the convergence rate for one of the runs is shown in fig 6.1.




Figure 6.1: Shows the $ES_{0.99}$ for the best and worst portfolios in the population in a run of the genetic algorithm.

For a conclusion, it seems like the genetic algorithm is able to produce a good portfolio with low risk, but it is not able to produce the optimal solution generally. I am not sure if this is because of theoretical or practical reasons. The algorithm seems to be performing better, the higher N and M one uses and one could probably improve the performance of the crossings, mutations and the IS as well. So, one cannot rule out that a genetic algorithm could work for this problem, but with my implementation and computer capacity it has not.

7 Discussion

7.1 Summary

After some initial background, I started out by explaining the model and how the different risk measures could be calculated. I then presented and discussed some different methods for estimating the parameters of the model: the migration matrix, the conditional losses and the correlation matrix. I continued by showing some calculations done on the model and I did some tests showing how sensitive the risk measures are to the parameters. Finally, I implemented a genetic algorithm trying to optimize the portfolio with respect to ES .

7.2 Conclusions

As a statistician, I cannot say that I have much confidence in this model. The model is based on several assumptions which are not undoubtedly correct. However, these problems would be present for most other credit risk models as well. I will now list some of the problems with the model.

- The first and biggest problem with the model is that it claims to say something about what happens in extreme cases, but it can only be based on a few years of data. Risk measures like $VaR_{0.999}$ and $ES_{0.999}$ are supposed to state something about losses that would occur on average just once in a thousand years and to base those statements on fifteen years history only, will of course make the statements very uncertain. Since the financial systems constantly changes, so will the credit rating companies' view on different types of obligors and therefore a much longer relevant history will probably never be available. Testing the effect of using the confidence bounds of the migration matrix when calculating the risk measures also showed how uncertain the measures are statistically. Measures like EL and UL can probably be estimated with higher certainty, but they are less interesting in terms of quantifying credit risk.
- Another problem with the model is the assumption of the underlying assets and the asset correlation. For sovereign obligors there exist no real value that corresponds to the underlying value which makes it impossible to validate the assumption of normality. Neither is it possible to estimate the correlations directly or to validate that the estimates are correct. Even if the normality assumption was correct, the assumption that correlations are constant in time seems wrong from fig 4.4. One could argue that the correlation estimates could be updated frequently to get the current correlation. However, it is not the correlation today that is interesting, but the correlation in the coming year. What says that the correlation cannot change rapidly and how could one in that case predict the change? What ECB (and myself in the end) did was to use an upper limit for the correlations to stay on the safe side. One can do this, but then one does not estimate measures like "expected shortfall" but rather "an upper

bound for expected shortfall". The same goes for when using the tweaked migration matrix.

- The credit rating agencies can be criticized in a number of ways. The agencies are private companies and their incentives might not always be to give ratings that best corresponds to the obligors' capabilities to live up to their financial obligations. For example, before the Enron crash the company had a high rating until just days before the bankruptcy, even though the agencies had been aware of the company's problems (The New York Times, 2002). It has often been noted that bond yields has risen for obligors already prior to downgrades, which means that the market sometimes has better information about the credit quality than the agencies. Therefore one should probably not trust the rating agencies blindly, but look at other data as bond yields or CDS spreads as well.

The main benefit of the model is that it is fairly simple and cheap. Not much data is needed to make the model work and given the parameters, I have shown that the risk measures can be calculated with high certainty using Monte Carlo methods. For the Riksbank's purposes the model could probably be sufficient to give an indication of how the credit risk in their portfolio develops over time. By using the upper limit for correlation from Basel II and the tweaked migration matrix they will get an upper limit for the risk measures, which will protect them somewhat from estimation errors showing a credit risk that is too low. However, they should not trust the figures from the model as absolute truths, but rather as numbers that can be used to compare the risk for different portfolios.

In the end, credit risk for sovereign bonds must be very hard to estimate for any model. The combination of having a limited amount of data and at the same time wanting to estimate a non-probable but long tail in the loss distribution is a problem that cannot be solved with certainty. The safest way to keep the credit risk low will always be to have a portfolio of high rating and, as I showed in table 10, a well diversified portfolio.

7.3 Further Development

- The primary output from the model was the risk measures VaR_α and ES_α for high levels of α . The problem with these was the uncertainty of input data, which made the risk measures very uncertain as well. A better risk measure might be one that instead of using a high α uses high confidence levels for the input data or maybe a combination of the two. This would be somewhat similar to what I did with confidence bounds for the migration matrix in section 5. Perhaps this could be achieved with some Bayesian approach, where prior distributions are put on the input parameters of the model and then posteriors are calculated using the input data. All this could be inserted into the Monte Carlo framework to calculate a loss distribution that is more based on the likelihood of

different input parameters, rather than interpreting the input data as the absolute truth. In (Gupton et al., 1997), they do something similar by treating the recovery rates as beta distributed. Using this method would probably be more computationally troublesome however.

- Using the underlying Gaussian variables and their correlations is very ad hoc and it would be nice to use something else that is measurable to base the joint distribution of rating changes on. Like I mentioned earlier, for corporations one can use stock prices, but this does not work for countries. I assume this problem is hard, but to make the model trustworthy one has to come up with something else.
- In this paper, I have only looked at a risk horizon of one year. Other risk horizons are of course interesting as well and to generalize the model to make it work for different horizons would not be too difficult. A lot of work has been done on this matter and in for example (Truuck and Rachev, 2009), they discuss how a time-continuous migration matrix, with migration intensities instead of migration probabilities, can be estimated by assuming Markov properties. The method basically requires data of the time points of all historical migrations, but given a one-year migration matrix M the corresponding time-continuous matrix can be calculated as

$$Q = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{(P - I)^k}{k},$$

if the sum converges. Given the time-continuous matrix, migration matrices for any horizon can be calculated.

- It would be interesting to insert a parameter into the model that describes the current state of the world economy. If the world economy is unstable more countries are likely to get economical problems and face downgrades or defaults. This parameter could for example be multiplied with the correlation matrix to increase the total volatility of the model in bad times. The current model is more time-stationary, which is a bit unrealistic.
- The optimization part can definitely be improved. The biggest problem seems to be that the Monte Carlo approximation of the loss distribution is not good enough, since the algorithm gave different solutions when using different random numbers. Maybe one could come up with some method that first runs the algorithm a few times and then bases a new instrumental distribution on the area of the loss distribution which is most important for the portfolios near the optimums of the first runs. This way one would get a higher resolution in the important part of the loss distribution, which could produce a more accurate optimum. One could also consider using some other algorithm than the genetic one.

References

- Blom, G., Enger, J., Englund, G., Grandell, J., Holst, L., 2005. Sannolikhetsteori och statistikteori med tillämpningar. Studentlitteratur.
- Bluhm, C., Overbeck, L., Wagner, C., 2003. An Introduction to Credit Risk Modeling. CRC Press LLC.
- Dowd, K., 2002. Measuring market risk. John Wiley & Sons Ltd.
- Duffie, D., Singleton, K., 2003. Credit Risk: Pricing, Measurement, and Management. Princeton University Press.
- Friewald, N., 2009. Estimating asset correlation from cds spreads: An analysis of the hedge effectiveness of ftd baskets. Working paper. Available at SSRN.
- Gupton, G., Finger, C., Bhatia, M., 1997. Creditmetrics-technical document. Tech. rep., J.P. Morgan & Co. Incorporated.
- Iscoe, I., Kreinin, A., Mausser, H., Romanko, O., 2012. Portfolio credit-risk optimization. Journal of Banking & Finance 36, 1604–1615.
- Merton, R. C., 1974. On the pricing of corporate debt: The risk structure of interest rates. Journal of Finance 29, 449–470.
- Rockafellar, R., Uryasev, S., 2000. Optimization of conditional value-at-risk. The Journal of Risk 2(3), 21–41.
- Rydén, T., Lindgren, G., 2000. Markovprocesser. Lund University.
- Sköld, M., 2006. Computer Intensive Statistical Methods. Lund University.
- The New York Times, ., 2002. Credit agencies waited months to voice doubt about enron. <http://www.nytimes.com/2002/02/08/business/enron-s-many-strands-warning-signs-credit-agencies-waited-months-voice-doubt.html?pagewanted=all&src=pm>.
- Trueck, S., Rachev, S. T., 2009. Rating Based Modeling of Credit Risk. Academic Press.
- van der Hoorn, H., 2009. Risk Management for Central Banks and Other Public Investors, Chapter 3. Cambridge University Press.
- Wilmott, P., 2007. Paul Wilmott Introduces Quantitative Finance 2nd ed. John Wiley & Sons Ltd.