
ON CREDIT SPREADS:
AN AUTOREGRESSIVE MODEL APPROACH

MASTER'S THESIS IN MATHEMATICAL STATISTICS

By

ANTON SCHÖLIN & FILIP PERSSON MÖRK

SUPERVISOR:
PROF. ERIK LINDSTRÖM



LUNDS
UNIVERSITET

FACULTY OF ENGINEERING
2017

Abstract

This thesis proposes an autoregressive credit spread model to make long term simulations of credit spreads and credit indices in the Investment grade and High yield bond segments. Several models are tested, and the final spread model produces simulations with statistics consistent with historical data, even though the model itself is relatively parsimonious. A transition from spread to index is proposed, which gives simulated indices with characteristics that match historical indices reasonably. Also, dependence between asset classes is introduced with a grouped t -copula.

Acknowledgements

We should like to express our gratitude towards our supervisor professor Erik Lindström for his guidance, support and valuable advice throughout this thesis. Furthermore, our thankfulness is extended to our supervisors at Kidbrooke Advisory - Edvard Sjögren and Ludvig Hällman - for their never-ending input and ideas. Lastly, we should like to thank our families and friends, without whose companionship this process would have been a less enjoyable experience.

Contents

- 1 Introduction** **1**
 - 1.1 Background 1
 - 1.2 Literature 5
 - 1.3 Thesis contribution 6
 - 1.4 Objectives 6
 - 1.5 Scope and limitations 6
 - 1.6 Thesis outline 7

- 2 Theory and concepts** **8**
 - 2.1 Pricing bonds 8
 - 2.2 Modeling defaults 8
 - 2.3 Time series and volatility modeling 11
 - 2.4 Modeling a credit index 14
 - 2.5 Retrieving \mathbb{P} -probabilities for default 15
 - 2.6 Dependencies across asset classes 16
 - 2.7 Statistical inference 17
 - 2.8 Model validation 18

- 3 Method** **23**
 - 3.1 Data selection 23
 - 3.2 Data analysis 24
 - 3.3 Model building 27
 - 3.4 Model evaluation 28
 - 3.5 Index model 28
 - 3.6 Asset class dependence 31
 - 3.7 Model validation 33

- 4 Results** **36**
 - 4.1 Data analysis 36
 - 4.2 Model comparison 38
 - 4.3 Model evaluation 41
 - 4.4 Models with input 44
 - 4.5 Parameter estimates and stability for final models 47
 - 4.6 Historical spreads with confidence interval 51
 - 4.7 Index model 51
 - 4.8 Dependence between asset classes 56

- 5 Discussion** **61**
 - 5.1 Objectives summary 61
 - 5.2 Credit spread model 62
 - 5.3 Credit index 62
 - 5.4 Dependence 63
 - 5.5 Future research 64

A	Plots	67
A.1	Residual and distributional analysis plots	67
A.2	Bootstrap analysis plots	73
B	Grouped t-copula	74
B.1	Derivation of PDF	74
C	Risk premium	75
C.1	Risk premium estimation	75
D	Rolling window	76
D.1	Rolling windows variance of estimates	76

Chapter 1

Introduction

1.1 Background

In the financial world, risk models are currently in high demand. Even though the market for credit derivatives has decreased substantially after the record high values of 2007, it is still a multi-trillion dollar market. However, nowadays it is not only speculation that drives the demand for good credit risk models. After the most recent financial crisis new rules and regulations have been developed and the requirements for banks and insurance companies have been strengthened with respect to risk management. This means that financial institutions require good credit risk models, whether they want to or not. There are several measures of credit risk, being the risk of default on a debt due to the inability of the borrower to make required payments. One is the so-called credit spread, which is the spread between the yield of a risk-free bond and a bond associated with risk. This thesis will focus on the modeling and simulation of this spread, and also credit indices to which the spread is tightly connected. Its composition, characteristics and importance will be discussed below and in the chapters to come. This introduction serves to briefly introduce the reader to the concept of financial risk, credit ratings and, of course, credit spreads. The introduction will continue with the thesis objectives, its scope and limitations and an outline of the paper.

1.1.1 Financial risk

Risk is a core concept in finance. If risk was not present speculative trading would not exist and the possibility to make money on any financial market would be diminished. Since risk is such a wide concept, there are several different types of risk present in finance. The main topic of this thesis is credit risk, which is one type of risk, but far from the only one. Some of the most common types of financial risk are listed below.

Market risk The risk of losing money because of factors affecting the entire market is called market risk. This risk is also called systematic risk and it can not be removed with diversification. Major market shifts can be caused by, for example, politics, recessions, depressions or natural disasters.

Liquidity risk If a particular asset can not be traded due to lack of buyers or sellers, a market participant might lose money because they can not buy or sell the asset quickly enough. This is called liquidity risk.

Interest rate, Currency and Equity risk The names of these types of risk are quite straight-forward, and are risks associated with the volatility inherent in interest rates, currencies and equity, respectively. The driving factor of the volatility is the different views market participants have of future values of stocks, exchange rates and interest rates. Without this difference in views of the future, financial markets would probably not exist, at least not to the extent they do today.

Credit risk Credit risk is the risk of not being paid back on a debt due to the borrowers inability to make the required payments, also known as default. This risk is associated with credit quality, and the less

credit quality of the borrower, the more compensation the investor wants in order to lend money. There are some institutions that are considered more or less risk free, that is, the probability of default is so low it is considered as zero. The United States government is usually considered a risk free part to lend money, and thus the rate at which money is lent to the US Treasury is considered the risk free rate. Crudely, one could say that they will always be able to pay back their debts, since they can always print more money. The government can also increase taxes to increase income. However, lending money to a government is not entirely risk free. For example, the political stability plays an important role when assessing sovereign credit risk. Credit rating agencies look at several different factors when assigning credit ratings to governments, including government debt and budget deficit. Companies seeking to borrow money issues what is called Corporate bonds. Since lending to companies is generally more risky than lending to a government, the compensation is higher. This is reflected in the price of the bond, which is lower for a bond issuer with a higher risk for default.

1.1.2 Credit ratings

Different parties on the financial market have different prerequisites to be able to pay back a loan. Assessing the credit worthiness of a counter-party is of great importance for most market participants, and this is why there exist so-called credit rating agencies. Their job is to rate a debtor's ability to pay back debt. These ratings look different for different credit rating agencies, but are generally ranging from AAA to D. AAA bonds are very likely to pay back borrowed money and D are bonds in default. The most significant credit rating agencies are Standard & Poor's, Moody's Investors Service, and Fitch Ratings. They are called 'The Big Three' and control approximately 95 % of the credit ratings business.

Criticism has been raised to the accuracy and responsiveness of credit ratings, especially in connection to the subprime mortgage crisis (Levin and Coburn, 2011). Also, several papers have explored the bias that can arise when the rating fees are paid by the asset issuers, among others Damiano et al. (2008) and Bolton et al. (2012). The tight connection between credit ratings and asset prices means that asset issuers could potentially affect the prices of its own assets by shopping for a good credit rating.

Bonds can be divided into different segments depending on their credit rating, visualized in Table 1.1. Investment grade (IG) bonds are those with a rating of BBB-and higher, thus being bonds with low probability of default. High yield bonds (HY) are comprised of bonds with credit ratings lower than BB+. Since the objective probability of default is generally higher for HY bonds their credit spread is most often higher as well.

Credit rating	Rating description		
AAA	Prime	Investment grade	
AA+	High Grade		
AA			
AA-			
A+			
A	Upper medium grade		
A-	Lower medium grade		
BBB+			
BBB			
BBB-	Non-investment grade speculative		High yield
BB+			
BB			
BB-			
B+		Highly speculative	
B			
B-			
CCC+		Substantial risk	
CCC		Extremely speculative	
CCC-		Default imminent with little prospect for recovery	
CC			
C			
D	In default		

Table 1.1: Table describing bond credit ratings, with credit rating scale defined by S&P and Fitch.

1.1.3 Credit spreads

Credit spreads are defined as the difference in yield to maturity between a default-free zero coupon bond and a zero coupon bond with some default risk. This can of course be extended to coupon bonds as well, as they can be seen as multiple zero-coupon bonds. The credit spread is essentially the price premium for choosing to lend money to a borrower who might not pay back in full, or at all. Spreads can be calculated for individual corporate bonds, but also for portfolios of bonds in a given credit class.

High values of credit spreads are associated with times of economic turmoil. Since the spread measures the default probability, a high value indicates that the market perceives the bond as riskier than before. In Figure 1.1 the historical path of credit spreads can be seen. The spreads are from the IG and HY segments, with corporate debt publically issued in the US domestic market. The shaded areas in the figure corresponds to areas of economic turbulence. The first is located at around 1999, which is the Russian financial crisis. This is followed by the dot-com bubble burst, September 11, and the Enron, WorldCom and Tyco scandals. The highest peak at around 2008 is the economic crash that started with the housing market collapse in the U.S, and the last shaded area at around 2012 is the Euro area crisis, with deep recessions in among other countries, Greece and Portugal.

The credit spread for a given credit rating says something about the state of the economy. If the spread for IG bonds increases, the market thinks IG bonds are more likely to default, implying a recessing economy. In times of economic uncertainty the “flight-to-quality” phenomenon, meaning investors seeking safer investments, tends to widen credit spreads. Being able to predict spreads for both IG and HY bonds can help predict the risk of investments over the entire market. Simulation of credit spreads can give a view of the possible state of the economy at a future point in time, which can help in assessing the overall risk of a certain investment in bonds with different credit ratings.




Figure 1.1: *Barclays Bloomberg option-adjusted spreads for Investment grade and High yield segments. The shaded areas highlights periods of economic turbulence.*

1.1.4 Credit risk and the financial crisis of 2008

Credit risk was an important factor in the financial crisis of 2008, which caused millions to lose their jobs and homes, as well as the demise of one of America's largest investment banks. Of course, a lot of factors are at play when it comes to the financial market, but lack of modeling defaults has been pointed out as one of the reasons for the magnitude of the crash (Salmon, 2009). The crisis is believed to be rooted in America's housing market, especially in the so-called Subprime mortgages. In 2004 these risky loans became more and more popular, and in 2006/2007 the default rates on these loans began to increase heavily. Modeling of defaults so far had not included the expanse of tail-dependence that was now seen, and the banks were hit by surprise. Their models did not take into account the high correlation between defaults when default rates are high. This led to relatively cheap loans for issuers with low credit quality and, in the end, a big economic wipe-out.

1.1.5 Regulations

To prevent future economic crises due to risky loans, new laws and regulations have been developed to strengthen capital requirements. This affects banks and insurance companies, by making them increase liquidity and decrease leverage. A key part of this is the risk measures of future payments. Credit risk is thus an important part of the new rules and regulations. Being able to assess future distributions of credit risk is an integral part of a bank's overall risk assessment.

The most noteworthy regulatory framework for banks is Basel III, developed by the Basel Committee

on Banking Supervision. This committee has members from 28 countries all over the globe and has a large impact on the financial world. The Third Basel Accord follows, of course, the second but has strengthened the capital requirements for the financial institutions. The first change is an increase in the ratio of capital versus risk weighted assets. A leverage ratio is also introduced, controlling the size of the banks' balance sheets. A main aspect of the framework is the regulations for liquidity. Basel III has developed stress tests for financial institutions to see how well they handle periods of high defaults.

Insurance companies in Europe has to abide the Solvency II Directive, which regulates the amount of capital the companies must hold in order to reduce risk of default, or insolvency. Insurance companies in the EU need to be able to pay what they owe over the next 12 months with a probability of at least 99.5 %.

In order to comply with the requirements, as well as protecting themselves against lending more capital than they can handle, financial institutions need to be able to predict how the credit spreads, and thus default rates, can change. Knowing the distributions of future credit spreads is an important factor in risk management.

1.1.6 Credit risk and trading strategies

Assessing credit risk of portfolio constituents is an important part in setting the risk profile of a trading strategy. The future distribution of credit risk in different rating segments affects the way investors should invest their money, given their risk appetite. Therefore, it is not only because of rules and regulations that credit spreads are interesting for banks and insurance companies, but also because of the foundation it provides for making good investments. Credit spread models can provide a way of simulating credit indices, making it easier to forecast investment riskiness and profitability.

1.2 Literature

There are several papers describing approaches for modeling credit spreads. Most studies are on a firm specific level. These models differ from those on credit rating or index level. This might be due to the fact that the value of specific firms tend to vary more than the value of an index. In this section, some approaches of modeling credit spreads and indices are described.

An approach to model firm specific credit spreads, described in Arvanitis et al. (1999), is to model credit spreads for different ratings and allow the firm to jump between the different rating classes. The jumps between ratings are governed by a Markov process and the credit spreads for each rating is modeled by a jump diffusion process. There are also successful attempts to model the credit spread using a time continuous mean reverting process with stochastic volatility. This approach is described in Jacobs and Li (2008), and is inspired by the connection between the credit spread and the hazard rate, which is often modeled using a Cox-Ingersoll-Ross model. On firm specific level only time continuous models were found.

When it comes to credit spreads on index level these can be divided into two different groups - those that consider rebalancing of the portfolio, and those that do not. In Bierens et al. (2003) a method of modeling the option adjusted logarithmic spreads connected to portfolios of bonds with a specific rating was developed. An ARX(1,1) model with ARCH structured noise and jumps was proposed. The input to the spread model was the Russel 2000 stock index. Further input to the model was the CBOE VIX index which was significant in explaining the conditional jump probabilities. Rebalancing of the portfolio was considered by removing a memory term of the days of rebalancing.

A paper using multivariate GARCH models to model European corporate bond indices is Gabrielsen (2010). The paper focuses on the time varying correlation between corporate bond indices. Evidence of time varying correlation is found. It also suggests the Markov regime switching between different GARCH specifications and finds significant results for High volatility regimes. A similar methodology was proposed by Manzoni (2002), who also used time varying dependence and GARCH models on the sterling Eurobond market.

The credit spread reflect the risk involved in borrowing money to companies. The risk of default or other credit events ought to be correlated with macroeconomic variables, for instance the market volatility, the inflation and the interest rate. These are also correlated with each other. Empirical studies suggests that the shape of the risk free yield curve and the credit spread curve together with current risk free rates and credit spreads are enough to optimize predictions of future credit spreads Krishnan et al. (2007).

Given a credit spread and a risk free rate the problem of transitioning to an index remains. No articles describing this procedure were found. In Wang et al. (2009) a method for modeling defaults between dependent assets was proposed. To model a firm specific default, two random variables were considered, one accounting for the firm specific risk and one accounting for the market risk. A default would then occur when the sum of the random variables exceeded some threshold. By estimating the dependence between assets it was possible to simulate dependent credit events.

1.3 Thesis contribution

Some articles suggest an autoregressive approach for modeling credit spreads. To our knowledge all of these either consider the driving noise term being Gaussian or Gaussian mixture. Since the Gaussian distribution can not explain the fat tails of the innovations of the credit spreads a popular approach is to add a jump process. The way the problem of fat tails is tackled in our thesis is by using a t -distributed driving noise term instead. This reduces the complexity of the model while still explaining the fat tails.

To model the volatility process all articles considered used the GARCH setup. It was shown in our study that an EGARCH setup with a term accounting for skewness with respect to the sign of the innovation was significant in explaining the volatility. Models with the credit spread as input to the volatility yielded high log-likelihood of the residuals. However, these models were not successful in producing realistic simulations.

The transition from a credit spread to a credit index is not found in any paper. The price of the index given a specific credit spread can be calculated using the definition of a credit spread, and the problem remaining is that of modeling credit events within the index. One paper describing this was found. The method demanded that the correlation between assets was estimated. Since the indices considered in this thesis consist of so many assets and are re-balanced frequently this approach could not be used. Therefore a new way of modeling credit events within indices was proposed.

1.4 Objectives

This thesis is built around several objectives, which are listed below. Each objective builds towards the goal of making good simulations of credit spreads and credit indices.

Find an autoregressive credit spread model This thesis sets out to use the autoregressive structure in credit spreads to model spreads in the High yield and Investment grade segments. Several model structures are compared and tested.

Spread-to-index transition In order to simulate credit indices, a method of going from spread to index must be found. This is straight forward when dealing with company specific bond data, but not so for entire rating segments.

Dependence structure To understand the role of credit spreads in the financial world, dependencies with other asset classes must be examined. Interest rates is one particularly important asset class to analyze, since credit indices are dependent on both the credit spread and the risk-free rate. Thus, the dependence between the two will have an effect on simulating credit indices.

Simulation validation The overall goal with this thesis is simulating credit spreads and indices. Validating against historical data will provide a measure of how well the models perform. The validation method should test several statistics of the simulations to make sure that the models can produce realistic credit spreads and indices.

1.5 Scope and limitations

The scope in this thesis will be on modeling and simulating credit spreads, and in turn credit indices.

1.5.1 Macro-economic dependencies

Credit spreads are of course dependent on many macro-economic variables such as market volatility, economic recessions and depressions, market outlook on future economic states and many more. In this thesis we treat the credit spread as the factor deciding the state of the economy and not the other way around, which is why there are not many macro-economic variables present in the models proposed. Since the goal of the model is to predict a distribution of future states of the economy, a choice can be made what defines an economic state, and in which order the causality of different variables come. If one is able to predict the credit spreads, and the correlation of credit spreads with other economic factors are high, then it is just a matter of which variable is the easiest to predict and model.

1.5.2 Regime switching models

In order to keep the model parsimonious and easy to calibrate, regime switching models will not be considered. The model will thus be the same at all times, and the process will not switch between several different models based on some criteria. Regime switching could improve the results, but there is value in ease of calibration which will be much more difficult with a regime switching approach.

1.6 Thesis outline

The rest of the thesis is structured in the following way.

Chapter 2. Theory and concepts In Chapter 2 theories and concepts regarding credit risk, credit spreads and finance are reviewed. This will give the reader a good understanding of the underlying dynamics of credit spreads and default probabilities, and an overview of mathematical background involved when modeling and simulating credit spreads and indices.

Chapter 3. Method This chapter focus on the methodology used in the practical aspects of this thesis. It describes the methods used to derive the results in the following chapter. Firstly, the data selection process will be described, followed by the way models were built, and the method behind the spread-index transition. The methods behind handling dependence, model calibration and validation is depicted next.

Chapter 4. Results Here, parameter estimations will be listed, as well as other metrics providing a measure over how good the different models are in different aspects. Simulation results and validation against historical data will be shown, both for spreads and indices. This is essentially the results achieved from following the methodology from the previous chapter.

Chapter 5. Discussion Some aspects of the results in the previous chapter will be discussed. Problems encountered in the thesis will be surfaced, together with explanations to choices made throughout the thesis. There will also be a review of the objectives, to sum up the thesis.

Chapter 2

Theory and concepts

In this chapter, some mathematical and financial background will be provided. The concepts used in this thesis will be summarized, to make the reader better understand the results in the later chapters.

2.1 Pricing bonds

A default-free zero coupon bond, with $p(t, T) = 1$, can be priced as follows, according to Björk (2009).

$$p(t, T) = e^{-\int_t^T f(t, s) ds}, \quad (2.1)$$

where $p(t, T)$ is the price of the zero coupon bond with maturity T at time t , and $f(t, s)$ is the instantaneous forward rate with maturity s , contracted at t . The integral of the instantaneous forward rate can be seen as the yield of the bond, $y(t, T)$. For the price of a corporate bond, there is a risk of default, and therefore the price of such a bond will be cheaper, making the yield higher. The discounting of a corporate bond cash flow s_i will be on the form

$$s_i e^{(y(t, T) + c_i(t, T))(T - t)}, \quad (2.2)$$

where $c_i(t, T)$ is the credit spread of company i , corresponding to the maturity T , just as in McNeil et al. (2005). This spread will increase with the probability of default for the given company.

To assess the price of a corporate bond it is essential to know the probability of default for that particular bond. The higher the default probability, the lower the price of the bond. There are several ways of modeling the probability of default for a given company and we will present some different approaches: The classical models are Structural models and reduced-form models, see Section 2.2.1 and 2.2.2 below. Another way of modeling prices of corporate bonds are models based on credit ratings, which are described in Section 2.2.3.

2.2 Modeling defaults

2.2.1 Structural models

Structural models handle defaults as a relation between a company's assets and liabilities. Whenever the assets fall below a liability threshold value, L , representing liabilities a default has occurred. One of the most famous structural models is that pioneered by Merton (1974), where a company's assets are assumed to follow a geometric Brownian motion. The dynamics of the assets, $S(t)$, thus look like the following.

$$dS(t) = rS(t)dt + \sigma S(t)dW_t, \quad (2.3)$$

where W_t is a standard Brownian motion under risk-neutral measure \mathbb{Q} . Given a market containing a risk free asset, $B(t)$, as well as the assets of the company, $S(t)$, the price of the company bond will satisfy the Black and Schole's equation. Hence, it follows from the Risk Neutral Valuation Formula (Björk, 2009), that the price of the company bond is given by

$$P(t, T) = E^{\mathbb{Q}} \left[\frac{B(t)}{B(T)} \cdot \phi(S(T)) | \mathcal{F}_t \right], \quad (2.4)$$

where $\phi(S(T))$ is given by $\min(L, S(T))$. The solution to equation 2.4 is, with a fixed interest rate given by

$$P(t, T) = e^{-r\tau} L\Theta(-d) + S_t\Theta(d - \sigma\sqrt{\tau}),$$

$$d = \frac{\log(\frac{L}{S_t}) - (r - \frac{\sigma^2}{2})\tau}{\sigma\sqrt{\tau}}, \quad (2.5)$$

where $\tau = T - t$, σ is the standard deviation of the diffusion term in the geometric Brownian motion and $\Theta(a) = \int_{-\infty}^a f(x)dx$, where $f(x)$ is the standard normal probability density function (PDF).

A big drawback with the Merton model is that a company can only default at its debt maturity date, which is unrealistic. Also, the value of the firm is assumed to be a tradeable, observable asset. This is not true in all cases, especially for unlisted companies. It also assumes that assets of companies move as geometric Brownian motions. Empirically, increments of price processes of company bonds have fatter tails than what is possible to achieve with Gaussian noise.

A model which addresses the problem of the non observability of a company's assets is the KMV-model. This is a model which replaces the expected default frequency function given by the Merton model with empirical default data, and which estimates the company's assets from its stock value. It also replaces the default threshold D by a more realistic structure of the companies liabilities in case of a default situation. The model is maintained by Moody's and was in 2004 used by 40 out of the world's 50 biggest banks (McNeil et al., 2005).

2.2.2 Reduced-form models

In reduced-form models the time of default is modelled as some non-negative random variable. According to McNeil et al. (2005), the most basic form of reduced-form credit risk models are hazard rate models. These models treat defaults as jumps in a Poisson process and hence the time of default, denoted τ , is exponentially distributed with a cumulative distribution function (CDF) given by equation 2.6.

$$P(\tau < t) = 1 - \exp\left(-\int_0^t \lambda(s)ds\right). \quad (2.6)$$

This gives the *hazard function* $\lambda(t)$ as

$$\lim_{\Delta t \rightarrow 0} P(\tau \leq t + \Delta t | \tau > t) = \lambda(t). \quad (2.7)$$

For a given time t , the hazard function gives the instantaneous risk of default at that time, given survival up to t . We can use the hazard function under a risk neutral measure to calculate the price of a defaultable zero-coupon bond. Given that a default occurs, nothing will be paid out to the bond holder, the price will be a combination of the risk neutral hazard function and the risk-free rate. Recalling the price of a risk-free bond with face value one

$$p(t, T) = E^{\mathbb{Q}} \left[\exp\left(-\int_t^T r(s)ds\right) \right], \quad (2.8)$$

where $r(s)$ is the short rate. Further, the price of the defaultable bond can be written as

$$p_d(t, T) = E^{\mathbb{Q}} \left[\exp\left(-\int_t^T r(s)ds\right) \mathbf{1}_{\{\tau > t\}} \right] = E^{\mathbb{Q}} \left[\exp\left(-\int_t^T r(s) + \lambda(s)ds\right) \right]. \quad (2.9)$$

It is unrealistic that the hazard rate is deterministic for all future points in time. Therefore it is often modelled as a stochastic process. The resulting Poisson process has a stochastic intensity and is called a Cox process. Changing $\lambda(s)$ to $\lambda(s)(X_s)$, where X_t are some state variables, the value of a defaultable bond at time t is given by the following equation.

$$p_d(t, T) = E^{\mathbb{Q}} \left[\exp \left(- \int_t^T r(s) + \lambda(s)(X_s) ds \right) \right]. \quad (2.10)$$

Adding the fact that a partial sum of the face value can be given to the bond holder in case of default, the pricing becomes somewhat more involved. This fraction is called the recovery rate. There are some different types of recoveries: Recovery of market value, face value, and treasury. Recovery of face value measures the value to the investors as some percentage of the face value of the bond. Recovery of market value measures the change in market value in case of default. This is convenient for modeling purposes, as seen in Lando (2004). With an assumption of a constant recovery rate δ , the price of the defaultable contingent claim at time t becomes

$$p_d(t, T) = E_t^{\mathbb{Q}} \left[\exp \left(- \int_t^T (r(s) + (1 - \delta)\lambda(s))(X_s) ds \right) f(X_t) \right]. \quad (2.11)$$

A δ of zero leads back to equation 2.10. The recovery of treasury approach replaces the defaultable bond with a risk-free bond in the case of default, with the same maturity, but reduced payment.

The hazard process is similar to a short rate process, especially if r_t and λ_t are independent and the expectations can be separated. Interest rates can be modeled in a variety of ways. A popular approach is the CIR-model, introduced by Cox et al. (1985). The model can be written in the following way

$$dr_t = \kappa(\theta - r_t)dt + \sigma\sqrt{r_t}dW_t. \quad (2.12)$$

Another popular model is the Vasicek model, defined below.

$$dr_t = a(b - r_t)dt + \sigma dW_t \quad (2.13)$$

This model can become negative, which is good for interest rates, but not for a hazard process.

If discretized, both models will become autoregressive processes, and CIR will have state-dependent volatility. This suggests that the autoregressive approach of modeling credit spreads could work well.

2.2.3 Credit ratings

Credit ratings are given by third-party companies, as discussed in Section 1.1.2. The rating is set by analyzing some credit metrics, both historically and future trends. The credit metrics consists of both macro-economic variables and financial variables such as profitability, asset risk, and funding structure (Hill et al., 2016). The rating agencies also provide rating transition matrices containing historical information about the probability of migrating from one rating to another. Since credit ratings offer a way to retrieve a probability for default through the transition matrix, they can be used in pricing purposes. However, no information about the recovery rate is included in the rating and therefore this will be have to be dealt with in order to retrieve the price of a bond in a specific rating class. Of course, the risk-free rate is also needed to compute the price of a bond.

Since credit ratings use historical default data over long time horizons both the transition matrix along with the probability of default for a specific rating class will not reflect the current market information but the average historical information. The ratings are updated rather infrequently, typically quarterly or yearly. In volatile times this could result in bad performance. However, the risk of following an incorrect market opinion is diminished.

		Five Years Later																						
		Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa1	Caa2	Caa3	Ca-C	WR	DEF	
Current Rating	Aaa	56	7	10	3	1	1	0	0	0	0												20	
	Aa1	9	46	10	9	3	3	1	2	1	0						0						17	
	Aa2	4	6	32	16	5	6	3	2	0	0	0	0										26	
	Aa3	1	4	8	35	16	9	4	2	1	0						0	0					20	
	A1	1	2	3	9	33	15	8	4	2	1	1	1	0	0			0		0			20	0
	A2	0	1	1	3	11	34	14	8	4	2	1	1	1	0	0	0	0				0	20	0
	A3	0	0	0	2	4	19	24	13	8	5	2	2	1	1	0	0	0	0			0	17	1
	Baa1	1	1	0	1	2	6	12	29	13	7	2	2	1	1	0	0	0	0	0	0	0	19	1
	Baa2	0	0	0	1	1	3	6	11	31	13	3	2	2	2	1	1	1	0	0	0	0	20	1
	Baa3	1	0	0	0	1	2	3	6	15	26	7	4	4	2	2	1	1	0	0	0	0	23	3
	Ba1	0	0	0	0	1	2	3	3	7	11	18	5	4	5	4	1	1	0	0	0	0	31	4
	Ba2		0		0	0	0	1	2	3	8	9	13	7	7	5	2	1	0	0	0	0	36	5
	Ba3	0		0	0	0	0	1	1	2	3	4	6	13	7	6	3	1	1	0	1	1	39	11
	B1	0	0	0	0	0	0	0	0	0	2	2	3	6	13	7	5	2	1	1	1	1	39	17
	B2	0			0	0	0	0	0	0	0	1	1	2	6	15	8	4	2	1	1	1	35	23
	B3						0	0	0	0	0	0	1	1	4	4	10	3	2	1	1	1	38	32
	Caa1									0	0	1	1	1	3	3	4	8	3	1	1	1	36	39
	Caa2								0	0	0	0	1	1	2	1	2	2	5	1	3	4	41	40
	Caa3											1	1		1	1		1	0	2	0	3	62	
	Ca-C											0			0	2	2	1	1	1	4	4	43	46

Figure 2.1: Moody's five year transition probability matrix. The matrix shows the historical five year transition probabilities between different rating classes. The state WR corresponds to 'Withdrawn Rating', meaning that the company is no longer rated by Moody's.

2.3 Time series and volatility modeling

Financial time series are typically modelled as random walks, with or without a drift and with independent increments. They often display no autocorrelation in returns but tend to show autocorrelation when it comes to absolute values of the returns (Lindström et al., 2015). A popular way to capture this behaviour is by using the GARCH-models. More about GARCH models can be found in Section 2.3.2. The multivariate case is found in Section 2.3.4.

Another stylized fact for financial time series is that they tend to behave in an unsymmetrical way - negative shocks tend to have more impact on the volatility than positive shocks. This is not captured by the GARCH model. An extension of the GARCH model is the EGARCH family of models, see Section 2.3.3, which makes it possible to capture this behaviour.

Drawbacks with the GARCH and EGARCH models are that they are made in such a way that it is assumed that the volatility at time $t + 1$ is known at time t , which is unrealistic. Therefore an extension of those models are the stochastic volatility models, see Section 2.3.5.

Credit spreads show empirically, unlike financial time series but similar to interest rates, heavy dependence upon previous values. This calls for an autoregressive structure. However, just like financial time series the volatility of the increments seem to be dependent. Popular models to capture autocorrelation in time series data are the AR and ARMA models, see Section 2.3.1. To capture the autocorrelation of the volatility a GARCH-type model will have to be used in combination with the ARMA model.

For financial time series it is reasonable to assume that either the model or the model parameters are time dependent. Therefore either the model parameters or the model itself will have to be updated through time. One way to update either the structure of the model or the model parameters is by using a filter. By treating different models or parameters of the process as states in a hidden Markov chain, and the filter is used to estimate which state is most likely at a given point in time. Parameter updating procedures as well as parameter estimation is explained in Section 2.7.

2.3.1 ARMA

An ARMA(p, q) model, where ARMA denotes Auto Regressive Moving Average, is given by equation 2.14. If the p -term is zero then the model is called a moving average (MA) model and if the q -term is zero then the model is called an autoregressive (AR) model.

$$y_t + \sum_{k=1}^p a_k y_{t-k} = \sum_{k=1}^q c_k \epsilon_{t-k} + \epsilon_t. \quad (2.14)$$

The model is often written on the form

$$A(z)y_t = C(z)e_t, \quad (2.15)$$

where $A(z)$ and $C(z)$ are monic polynomials. The process is stationary if the roots of $A(z)$ lie within the unit circle and invertible if the roots of $C(z)$ do (Jakobsson, 2015).

2.3.2 GARCH

The GARCH model, where GARCH means general auto regressive conditional heteroscedasticity, is used to model time series with dependent volatility. The GARCH(p, q)-model is given by equation 2.16. If the q -term is zero then it is called an ARCH-model

$$y_t = \sigma_t z_t, \quad (2.16)$$

$$\sigma_t^2 = \omega + \sum_{k=1}^q \alpha_k y_{t-k}^2 + \sum_{k=1}^p \beta_k \sigma_{t-k}^2.$$

For the ARCH-process $\sum_{k=1}^q \alpha_k < 1$ is required for stationarity and $\alpha_i > 0$ is sufficient to ensure positive variances. For the GARCH-process $\sum_{k=1}^q \alpha_k + \sum_{k=1}^p \beta_k < 1$ is required for stability and (α_i, β_i) have to be positive to ensure positive variances (Lindström et al., 2015).

2.3.3 EGARCH

The EGARCH model is an extension of the GARCH model. It addresses two major drawbacks with the GARCH model - the symmetry regarding positive and negative shocks and restriction of the parameters. Essentially it has the same structure as a GARCH model except that the volatility process is log transformed and that the previous innovations are not squared (Lindström et al., 2015). The EGARCH(p, q) model is seen in equation 2.17.

$$y_t = \sigma_t z_t,$$

$$\log(\sigma_t^2) = \omega + \sum_{k=1}^q \alpha_k z_{t-k} + \sum_{k=1}^p \beta_k \log(\sigma_{t-k}^2). \quad (2.17)$$

An extension of the EGARCH model can be seen in equation 2.18.

$$\log(\sigma_t^2) = \omega + \sum_{k=1}^q \alpha_k h(z_{t-k}) + \sum_{k=1}^p \beta_k f(\sigma_{t-k}), \quad (2.18)$$

where h and f are functions. These can for instance be chosen such that the volatility is mean reverting or differently dependant on previous shocks regarding to their size.

2.3.4 Multivariate GARCH models

Consider multiple stochastic time series $\mathbf{y}_t = \{y_{1,t}, y_{2,t}, \dots, y_{n,t}\}$, which consist of some conditional mean vector, $\boldsymbol{\mu}$, and some noise vector, $\boldsymbol{\varepsilon}$, as

$$\mathbf{y}_t = \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_t, \quad (2.19)$$

with $\boldsymbol{\varepsilon}_t = \mathbf{H}_t^{1/2} \mathbf{z}_t$, where \mathbf{H}_t is the covariance matrix of \mathbf{y}_t and \mathbf{z}_t is a vector of independent noise.

The multivariate GARCH models generalizes the univariate models and takes into consideration the correlation between multiple processes. This is done by modeling the time evolution of \mathbf{H}_t , as described in Gabrielsen (2010).

VEC, DVEC and BEKK MGARCH These are all generalizations of univariate GARCH models, described in Section 2.3.2. The VEC-MGARCH model, from Bollerslev et al. (1988), describes the evolution of \mathbf{H}_t as

$$\mathbf{H}_t = \mathbf{C} + \sum_{i=1}^q \mathbf{A}_i \circ \boldsymbol{\varepsilon}_{t-i} \boldsymbol{\varepsilon}'_{t-i} + \sum_{j=1}^p \mathbf{B}_j \circ \mathbf{H}_{t-j}, \quad (2.20)$$

where \circ denotes the Hadamard product, that is element-wise multiplication. \mathbf{A}_i and \mathbf{B}_j are symmetric matrices. This model yields a large number of parameters to be estimated. A solution is the DVEC MGARCH model, also suggested by Bollerslev et al. (1988), which treats \mathbf{A}_i and \mathbf{B}_j as diagonal matrices, reducing the number of parameters. The drawback of these models is that they do not guarantee a positive definite covariance matrix. A model that does is the BEKK-MGARCH, suggested by Engle and Kroner (1995), which can be written in the following way

$$\mathbf{H}_t = \mathbf{C}\mathbf{C}' + \sum_{i=1}^q \mathbf{A}_i (\boldsymbol{\varepsilon}_{t-i} \boldsymbol{\varepsilon}'_{t-i}) \mathbf{A}_i' + \sum_{j=1}^p \mathbf{B}_j \mathbf{H}_{t-j} \mathbf{B}_j'. \quad (2.21)$$

This model can also be simplified by making \mathbf{A}_i and \mathbf{B}_j diagonal.

Dynamic Conditional Correlation MGARCH The *Dynamical Conditional Correlation M-GARCH* model is proposed by Engle and Sheppard (2001) and defines the covariance matrix as

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \quad (2.22)$$

where $\mathbf{D}_t = \text{diag}(\sqrt{h_{i,t}})$, and $\sqrt{h_{i,t}}$ is usually time varying standard deviations from univariate GARCH or EGARCH processes. The returns in these GARCH processes are assumed to be normally distributed, which give rise to a likelihood function. Without the normality assumption the estimator will still have the quasi-maximum likelihood interpretation. The estimation is done in two steps. First the univariate GARCH to estimate \mathbf{D}_t , and then using this result to construct

$$\mathbf{H}_t = \left(1 - \sum_{m=1}^p \alpha_m - \sum_{n=1}^q \beta_n \right) \bar{\mathbf{H}} + \sum_{i=1}^p \alpha_i (\boldsymbol{\varepsilon}_{t-i} \boldsymbol{\varepsilon}'_{t-i}) + \sum_{j=1}^q \beta_j \mathbf{H}_{t-j}, \quad (2.23)$$

where $\bar{\mathbf{H}}$ is the unconditional covariance of the residuals from the univariate GARCH estimations.

Comparison of Multivariate GARCH models Gabrielsen (2010) uses MGARCH models to estimate VaR for some credit portfolios based on maturity and rating. He finds that the Diagonal BEKK model outperforms the others in-sample goodness-of-fit, but that the RiskMetrics, which is essentially an IGARCH model developed by J.P. Morgan in 1989, was the model of choice according to a VaR loss function, measuring how well the models forecasted portfolio losses.

2.3.5 Stochastic volatility

Stochastic volatility models are a family of models where the volatility process has a stochastic term. All previously mentioned conditional variance models can be transformed to stochastic volatility models by adding a stochastic term. A simple example can be seen in equation 2.24.

$$\begin{aligned} y_t &= \sigma_t z_t, \\ \sigma_t^2 &= \eta_t, \end{aligned} \tag{2.24}$$

where z_t is a zero mean Gaussian random variable with unit variance and η_t is a positive random variable independent of z_t . This kind of model is also called "Normal Mixture Variance Model".

2.4 Modeling a credit index

All models described in Section 2.2 are firm specific. A straight forward approach to model an index, would be to model each underlying asset individually and then model their dependence, as proposed by e.g. Wang et al. (2009). Due to the number of underlying assets and the fact that the index might sometimes be rebalanced, the task of modeling the dependence is too big. Additional problems of time dependent dependence and the lack of data makes the task almost impossible. Therefore the proposed way of modeling the index is to view the entire portfolio of underlying assets as one company.

Merton models demand that the assets of the modeled company are observable. This is a manageable task if one firm is considered, but for an entire index this would be difficult. For Hazard rate models the fact that an entire index is modeled might actually make the problem easier, since simplifying assumptions regarding the recovery rates could be justified.

No matter what model is used, the problem of credit events within the index remains. A credit event will result in a lower price of the index without necessarily implying that the probability of default for the remaining portion of the assets have increased. Credit events within an index can of course be modeled. The problem with this is, as always, the lack of data. If the face value of the index would be known at all times then it would be possible to retrieve data of the impact of credit events from the index as equation 2.25 below.

$$D_t = 1 - \frac{\phi(T|t)}{\phi(T|t-1)}, \tag{2.25}$$

where D_t denotes the proportion of the index that did experience a credit event between time $t-1$ and t and $\phi(T|t)$ denotes the face value of the index with maturity T at time t . Since we can write the bond price difference as

$$\frac{P_t}{P_{t-1}} = \frac{\phi(T|t)}{\phi(T|t-1)} \left(\frac{e^{-(r_t+y_{t-1})(T-t)}}{e^{-(r_t+y_{t-1})(T-t+1)}} e^{(r_t+y_t \frac{1}{12})} \right), \tag{2.26}$$

where y denotes the credit spread and assuming monthly data, we can rewrite the expression for D_t . In the case of unknown face values but known credit spreads it would be possible to approximate the impact of credit events within the index as equation 2.27.

$$D_t = 1 - \frac{P_t^{\text{credit}} P_{t-1}^{\text{Treasury}}}{P_{t-1}^{\text{credit}} P_t^{\text{Treasury}}} \cdot \frac{e^{-y_{t-1}T}}{e^{-y_t(T-\frac{1}{12})}}. \tag{2.27}$$

There are theoretical ways of modeling index credit events as well. These are built on the fact that the price of the index is a function of the risk neutral probability of default. The procedure of finding the real world (\mathbb{P}) probability of default is described in Section 2.5. Using this probability it is possible to model the defaults using, for instance, the binomial distribution. Since the defaults are most likely dependent the number of trials of the binomial distribution should not equal the number of bonds within the index. Instead, it should be set to a number such that the variance of the binomial distribution equals that of the historical defaults. This way both the expected number of defaults as well as the variance of them are

matched. The drawback with the binomial approach is that it reduces the number of possible outcomes and can not replicate all possible theoretical scenarios.

A simplification could be to not model defaults within the index. Log excess returns, denoted ξ_t , of the index could be modeled instead of the credit spread. The excess returns captures changes in recovery rate, risk neutral probability for default, as well as credit events within the index. The log excess returns are given by equation 2.28 below.

$$\xi_t = \log \left(\frac{P_t^{credit} P_{t-1}^{Treasury}}{P_{t-1}^{credit} P_t^{Treasury}} \right). \quad (2.28)$$

2.5 Retrieving \mathbb{P} -probabilities for default

It is possible to calculate probabilities for a given event using different measures as long as the measures are equivalent. When pricing, all probabilities are calculated under the risk-neutral, or Martingale, measure \mathbb{Q} . Given the price of a risky bond under the assumption of no arbitrage, given by equation 2.4, it is possible to retrieve the \mathbb{Q} -probability for default. If one is interested in actual default probabilities one has to figure out how to go from this measure to the physical measure \mathbb{P} . For a complete market the risk-neutral measure is unique, as described in Björk (2009). Given a known \mathbb{Q} -probability it is therefore possible to theoretically retrieve the \mathbb{P} -measure, see Section 2.5.1. In reality, completeness is a strong assumption. If the market is not complete, a multitude of prices are possible, and thus also several \mathbb{Q} -measures. In one way this makes it easier to derive a \mathbb{Q} -measure, since there are more of them. Then it is just a matter of selecting one measure, corresponding to one selected price. This can be done using empirical methods for retrieving the \mathbb{P} -probability for default. Some of these methods are explained later in this section.

2.5.1 Theoretical relationship between \mathbb{Q} and \mathbb{P}

Assume that defaults occur as jumps in a Poisson process with intensity λ . Given that the process is started at $t=0$, the probability that a default has occurred before time T can be calculated as

$$\mathbb{Q}(\text{default}) = 1 - \mathbb{Q}(N(T) = 0) = 1 - e^{-\lambda^{\mathbb{Q}}T}, \quad (2.29)$$

where $N(T)$ denotes the number of jumps in the Poisson process after time T , and $\mathbb{Q}(\cdot)$ denotes the probability under the measure \mathbb{Q} . The \mathbb{P} -probability for default is calculated similarly.

The price of a risky zero coupon bond with a zero recovery rate, constant risk free rate r and unit face value is given by

$$P_{ZCB}(0, T) = e^{-rT} \mathbb{Q}(N(T) = 0) = e^{-(r+\lambda^{\mathbb{Q}})T}. \quad (2.30)$$

Define $\mu = \frac{\lambda^{\mathbb{Q}}}{\lambda^{\mathbb{P}}}$, then it is possible to write the price of the zero coupon bond as

$$P_{ZCB}(0, T) = e^{-(r+\mu\lambda^{\mathbb{P}})T} = e^{-(r+\lambda^{\mathbb{P}}+(\mu-1)\lambda^{\mathbb{P}})T} = e^{-(r+(\mu-1)\lambda^{\mathbb{P}})T} \mathbb{P}(N(T) = 0). \quad (2.31)$$

By using equation 2.31 the relationship between $\lambda^{\mathbb{P}}$ and $\lambda^{\mathbb{Q}}$ can be empirically estimated as the difference between the average slope of the observed price process and the risk free rate as

$$\hat{E} \left[\frac{dP_{ZCB}}{dt} \right] - r = \lambda^{\mathbb{Q}} - \lambda^{\mathbb{P}} \Rightarrow \lambda^{\mathbb{P}} = r + \lambda^{\mathbb{Q}} - \hat{E} \left[\frac{dP_{ZCB}}{dt} \right], \quad (2.32)$$

where P_{ZCB} denotes the price process of the defaultable bond and $\hat{E} \left[\frac{dP_{ZCB}}{dt} \right] - r$ is estimated as the historical difference between the slope of the risk free asset and the defaultable bond.

2.5.2 Empirical methods for finding \mathbb{P} -probabilities for default

Credit ratings method It is of course possible to estimate default probabilities for companies by using historical default data. Summarized data depending on different criteria can be found in credit rating

matrices, see Section 2.2.3. Using this data and assuming a constant hazard rate it is easy to calculate $\lambda^{\mathbb{P}}$. The exact procedure can be seen below.

$$D = \mathbb{P}(\text{default}) = 1 - e^{-\lambda^{\mathbb{P}}T} \Rightarrow \lambda^{\mathbb{P}} = -\frac{\log(1-D)}{T}, \quad (2.33)$$

where D is the default probability given by the credit rating matrix.

Credit spread method Since the price of a risky bond can be expressed both as a function of the credit spread as well as a function of $\lambda^{\mathbb{Q}}$ it is possible to express $\lambda^{\mathbb{Q}}$ as a function of the credit spread. Assuming there is reliable credit spread data available, this is a good method of estimating $\lambda^{\mathbb{Q}}$. The exact expression for $\lambda^{\mathbb{Q}}$ as a function of the credit spread y can be seen below.

$$\lambda^{\mathbb{Q}} = \frac{1}{\tau} \log \left(\frac{1-\delta}{e^{-y\tau} - \delta} \right), \quad (2.34)$$

where δ denotes the recovery rate as a fraction of the face value of the bond. Define the risk premium as $r_p = \hat{E}^{\mathbb{P}} \left[\frac{dp}{dt} \right] - r$. Then by using equation 2.32 it is possible to calculate $\lambda^{\mathbb{P}}$ as

$$\lambda^{\mathbb{P}} = \lambda^{\mathbb{Q}} - r_p. \quad (2.35)$$

2.6 Dependencies across asset classes

An important aspect of credit spread modeling is the dependency structure to other asset classes. Consider a publicly traded company with an associated risk for default, which the price of a corporate bond for this company depends on. It is very plausible that an event triggering a drop in stock price for this company would also increase the probability of default. This would in turn mean that a drop in stock price increases the credit spread, and thus the credit spread and the stock price would be negatively correlated. On a larger scale, one can compare stock indices with credit spreads for entire segments.

To get accurate simulations across multiple asset classes, such as rates, spreads and stocks, the dependency structure is important. One way of introducing correlation is in the driving noise of the respective models. A popular approach to create the multivariate distribution for the noise is to use copulas.

2.6.1 Copulas

In Nelsen (2006), copulas are described as “functions that join or couple multivariate distribution functions to their one-dimensional marginal distribution functions” and as “distribution functions whose one-dimensional margins are uniform.” A copula takes care of all the dependencies between the marginal distributions. A joint distribution function $H(x_1, \dots, x_d)$ can be written as

$$H(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)), \quad (2.36)$$

where F is the distribution function for each marginal distribution and C is a copula. Sklar’s theorem states that if all F are continuous, then C is unique, and otherwise it is uniquely determined on the range of F . Also, if F_1, \dots, F_d are distribution functions and C is a copula, then H as defined by equation 2.36 is a d -dimensional distribution function, with F_1, \dots, F_d as marginal distributions.

Having a collection of copulas thus automatically yields a collection of multivariate distributions with whatever marginal distributions one might desire. This is a useful tool when it comes to modeling and simulation.

To decide which copula to use, likelihood based methods like AIC or BIC, see Section 3.4, can be applied. Bayesian approaches can also be used. This implies choosing a copula based on some desired characteristics and calculating the likelihood for the data, given that copula. A large number of samples of the same size as the original data are simulated. For each sample the corresponding likelihood is calculated. If the portion α of the samples have higher likelihood than the observed sample, the hypothesis that the copula have generated the observed sample can be rejected on the significance level α .

Gaussian copula A multivariate Gaussian distribution has the following density

$$f(\mathbf{x}) = \frac{\exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})\right)}{\sqrt{|2\pi\Sigma|}}. \quad (2.37)$$

To derive the Gaussian copula one can use the equation below

$$C(\mathbf{u}) = F\left(F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\right), \quad (2.38)$$

and integrate the density as

$$C_P^{Ga}(\mathbf{u}) = \int_{-\infty}^{\Phi^{-1}(u_1)} \dots \int_{-\infty}^{\Phi^{-1}(u_d)} \frac{\exp\left(-\frac{1}{2}\mathbf{x}'P^{-1}\mathbf{x}\right)}{\sqrt{|2\pi\Sigma|}} d\mathbf{x}, \quad (2.39)$$

where P is the correlation matrix and Φ^{-1} denotes the quantile function for a univariate standard normal distribution.

Student's t -copula The Student's t -copula, or just the t -copula, is derived from the multivariate t -distribution. A d -dimensional random vector \mathbf{X} has a multivariate t -distribution if it has the density

$$f(\mathbf{x}) = \frac{\Gamma\left(\frac{\nu+d}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)\sqrt{(\pi\nu)^d|\Sigma|}} \left(1 + \frac{(\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})}{\nu}\right)^{-\frac{\nu+d}{2}}, \quad (2.40)$$

where $\boldsymbol{\mu}$ is a mean vector and $\text{cov}(\mathbf{X}) = \frac{\nu}{\nu-2}\Sigma$. \mathbf{X} can also be written as

$$\mathbf{X} \stackrel{d}{=} \boldsymbol{\mu} + \sqrt{W}\mathbf{Z}, \quad (2.41)$$

where $\mathbf{Z} \sim N_d(0, \Sigma)$ and $W \sim Ig\left(\frac{\nu}{2}, \frac{\nu}{2}\right)$. Ig denotes the Inverse-Gamma distribution. The t -copula can be derived from equation 2.38 by taking the integral of the density of equation 2.40, in the same manner as in the Gaussian copula case. As $\nu \rightarrow \infty$ the t copula asymptotically converges to a Gaussian copula. A preferable behavior of the t copula when it comes to financial time series is the tail dependence, that is, the copula generates joint extreme events. This is not the case for the Gaussian copula, where the tail dependence coefficient is zero. This means that the multivariate distribution generated by the Gaussian copula generates independent extreme values. It has been debated that this was one of the reasons for worsening the financial crisis of 2008.

grouped t -copula A further generalization of the t -copula is the grouped t -copula. The idea is to make a copula that can have different tail dependence for different subvectors of \mathbf{X} . This means that \mathbf{X} is defined as follows

$$\mathbf{X} = \left(\sqrt{W_1}\mathbf{Z}_1, \dots, \sqrt{W_1}\mathbf{Z}_k, \sqrt{W_2}\mathbf{Z}_{k+1}, \dots, \sqrt{W_d}\mathbf{Z}_d\right)', \quad (2.42)$$

where $\mathbf{Z} \sim N_d(0, \Sigma)$ and $W_k \sim Ig\left(\frac{\nu_k}{2}, \frac{\nu_k}{2}\right)$. Different subvectors of \mathbf{X} has different ν related to the dependence structure, but the W_i 's are still perfectly positively dependent, with respect to Kendall's τ , see Section 3.6.1. It is important to separate this from the ν of the marginal distributions, which does not need to be t -distributed. A reasonable choice of subvectors can be different asset classes, which are now allowed to have different tail dependence. The grouped t -copula is the unique copula generated by multivariate distribution function of \mathbf{X} . A derivation of the pdf of the grouped t -copula can be found in Appendix B.1.

2.7 Statistical inference

In this section three ways of fitting the model parameters are described - least squares, maximum likelihood and filter estimation.

Maximum likelihood The most intuitive way of calibrating a model is by maximizing the likelihood of the residuals. Maximum likelihood estimates are consistent with the Cramer-Rao lower bound, which means that the resulting estimates have the smallest possible variance of all consistent estimates. Drawbacks with the method is the need for an expression of the pdf of the noise and that numerical optimization often is needed. Depending on the stiffness of the likelihood function there is a chance of converging to a local maximum and missing the true value. Numerical optimization can also be computationally inefficient.

Least squares Another way of estimating the parameters is by using the least squares method. This is constructed by finding those parameters which minimizes the sum of the squared residuals. In the case of Gaussian noise the least squares estimates are equivalent to the maximum likelihood estimates. However, for some distributions the least squares method will not yield consistent estimates. For example consider equation 2.43, where ϵ_i 's are the sorted residuals, it is possible to use least squares to find consistent estimates for any distribution.

$$\hat{\phi} = \underset{\phi}{\operatorname{argmin}} \sum_{i=1}^n \left(F_{\phi}(\epsilon_i) - \frac{i}{n+1} \right)^2. \quad (2.43)$$

The least squares method have the same drawbacks as the maximum likelihood method regarding the need for numerical optimization, but can often be solved in closed form.

Filter estimation The last way of estimating the parameters which will be covered in this section is to use filtering techniques for parameter estimation. The way to do this is to treat the parameters as hidden observations in noisy data. They are given a variance and are updated with every measurement as the parameters which are the most likely given the previous observations and the new one.

2.7.1 Parameter persistence

There are several ways of letting the parameters change with time. One way is to simply decide to only use a predetermined number of the latest observations when performing the parameter estimation. Another way is to use some kind of punishment function making old values have less influence of the estimation. With the filtering technique one way of letting the parameters change more with time is to increase their assumed variance.

2.8 Model validation

When trying to find a model for a data set, the data set is ideally very large. Optimally large enough so that the set can be split into three parts. The first part is for modeling and calibration and the second part is to evaluate whether the results attained from the first part are also valid outside the calibration set. That is, of course, assuming that the data is stationary and ergodic. If the results from the second part are unsatisfactory the model is rejected and the model building starts over again with the original first part of the data set. When a model passes the test on the second data set, the third and last set is used for evaluating how well the model works.

To evaluate whether a model is sufficient to explain the structure of a process there are several aspects to consider. The first being if the model assumptions are fulfilled on the set to which the model was calibrated. Does the residuals, representing the noise, fulfill the predetermined assumptions? The second is the predictive power of the model, that is, how well it performs outside the calibration set.

If the available data set is small, it could be better to use the entire set for calibration rather than wasting valuable observations on evaluation. An alternative way of validating the calibrated model is to assume that the model is correct and test this hypothesis on some confidence level using parametric bootstrap. This is the route taken in this thesis.

2.8.1 Residual analysis

Usual assumptions are that the noise is independent and identically distributed. The noise can be estimated as the residuals between the model predictions and the measurements. To statistically test whether these

assumptions can be rejected a number of methods for analysing the residuals are presented in this section. The first three paragraphs describe methods for testing the independence assumption and the remaining four examine the distributional assumption.

Visual autocorrelation test Asymptotically the estimated autocorrelations, denoted $\hat{\rho}(k)$, of the residuals are normally distributed as in Jakobsson (2015):

$$\hat{\rho}(k) \in N\left(0, \frac{1}{n}\right), \quad (2.44)$$

provided that $E[\varepsilon^4] < \infty$. By plotting $\hat{\rho}(k)$ for $k = 1, \dots, d$, where d is some arbitrarily chosen number dependant on what horizon it is relevant to examine the autocorrelation, it is possible to see whether there is a trend in the autocorrelation and if any $\hat{\rho}(k)$'s are outside a confidence interval at some significance level. If it is unclear whether there is structure in the residuals, maybe there is a small trend or most, but not all, $\hat{\rho}(k)$'s are inside the confidence interval, the visual autocorrelation test is insufficient and other tests will have to be used. However, it is always a good idea to plot the autocorrelation function and the confidence intervals as a first test. The autocorrelation function is further explained in Section 3.2.3.

Ljung-Box-Pierce, Monti and McLeod Li test Define the quantity Q as

$$Q = n \sum_{k=1}^d \hat{\rho}^2(k). \quad (2.45)$$

Since $\hat{\rho}(k)$ is asymptotically normally distributed, Q will asymptotically be $\chi^2(d)$ -distributed, where d is the number of examined lags. Unfortunately the convergence in distribution is slow and it has been shown that a better test quantity is

$$Q_{Ljung-Box} = n(n+2) \sum_{k=1}^d \frac{\hat{\rho}^2(k)}{n-k}, \quad (2.46)$$

as seen in Jakobsson (2015). This is a better approximation of a $\chi^2(d)$ -distribution. The rejection criteria of the hypothesis that there is no autocorrelation, is $Q > \chi^2(d)_{1-\alpha}$. This is known as the *Ljung-Box-Pierce test*.

Instead of testing the autocorrelation function, it is possible to test the partial autocorrelation function. This is done in exactly the same way as the *Ljung-Box-Pierce test* but the test quantity is instead being defined as

$$Q_{Monti} = n(n+2) \sum_{k=1}^d \frac{\hat{\phi}^2(k)}{n-k}, \quad (2.47)$$

where $\hat{\phi}^2(k)$ denotes the estimated partial autocorrelation function. The test is known as the *Monti test*.

To test for higher order auto correlation of the residuals the *McLeod Li test* can be used. The test quantity is the same as for the *Ljung-Box-Pierce test*, but the autocorrelation of the squared residuals is considered instead of that of the plain residuals.

Sign-change test For a white noise process the probability of positive shocks should equal that of negative shocks and the probability of one shock being succeeded by a shock of opposite sign should be fifty percent. Therefore the expected number of sign changes of the residuals should follow a binomial distribution with $n - 1$ number of trials and a fifty percent chance of success. Denote the number of sign changes with X . For large data sets the binomial distribution can be approximated with a normal distribution and it is asymptotically true that:

$$X \in N\left(\frac{n-1}{2}, \frac{n-1}{4}\right).$$

Reject the hypothesis that there is no autocorrelation in the residuals if

$$X \notin \left[\frac{n-1}{2} \pm \lambda_{\alpha} \sqrt{\frac{n-1}{4}} \right].$$

Probability plots An initial test to evaluate whether an independent random sample belongs to some distribution is to plot the empirical CDF versus the CDF of the test distribution. If the deviance is big the chances are that the sample has been generated by some other distribution. There are several versions of probability plots, where the quantile-quantile plot is the most common. The empirical quantiles should follow a uniform distribution and therefore lie on a straight line between zero and one.

No matter which type of probability plot is used, the question will always be how far from the theoretical line the empirical line can be without violating the assumption of the generating distribution. The *Kolmogorov-Smirnov test*, the *Brownian Bridge test for distribution* and the *Pearson's chi squared test* are statistical methods of deciding this.

Kolmogorov-Smirnov test The Kolmogorov-Smirnov test is used to test the null hypothesis that a random sample has been generated by a specific distribution. The test quantity is given by the maximal distance between the empirical CDF and the CDF of the assumed distribution. Denote this quantity D_n , where n is the size of the random sample. It is then true that D_n converges in distribution as

$$\lim_{n \rightarrow \infty} \sqrt{n}D_n \in K,$$

where K denotes the Kolmogorov distribution. The null hypothesis is rejected if

$$\sqrt{n}D_n > K_\alpha,$$

where K_α is the $1 - \alpha$ percentile of the Kolmogorov distribution.

The Kolmogorov-Smirnov test can also be used to test whether two random samples have been generated by the same distribution. In this case D_{n_1, n_2} is defined as the maximal distance between the empirical CDF's of the two samples and n_i is the size of sample i . The null hypothesis is rejected on significance level $1 - \alpha$ if

$$D_n > \sqrt{-\frac{1}{2} \log\left(\frac{\alpha}{2}\right)} \cdot \sqrt{\frac{n_1 + n_2}{n_1 n_2}}.$$

Brownian Bridge test for distribution Define the empirical distribution function $F_n(x)$ as

$$F_n(x) = \frac{1}{n} \sum_{i=1}^n I_{[X_i \leq x]},$$

where X_1, X_2, \dots, X_n are independent random samples from some distribution. If all X_i 's have been generated by the same distribution with distribution function $F(X)$, as a result of the *Glivenko-Cantelli theorem* and *Donsker's theorem*, (der Vaart and Wellner, 1996), the function $G_n(x)$ defined as

$$G_n(x) = \sqrt{n}(F_n(x) - F(x)),$$

converges in distribution as n goes to infinity as

$$\lim_{n \rightarrow \infty} G_n(x) \in N\left(0, F(x)(1 - F(x))\right). \quad (2.48)$$

This yields that asymptotically

$$F_n(x) \in N\left(F(x), \frac{F(x)(1 - F(x))}{n}\right). \quad (2.49)$$

By using this it is possible to test the hypothesis whether a random sample has been generated by some distribution. The test can be normalized by defining $U(t)$ as

$$U(t) = P(F(X) \leq t) = P(X \leq F^{-1}(t)) = F(F^{-1}(t)) = t,$$

and similarly

$$U_n(t) = \frac{1}{n} \sum_{i=1}^n I_{[F(X_i) \leq t]}.$$

Then equation 2.49 implies that $U_n(t)$ is distributed as

$$U_n(t) \in N\left(U(t), \frac{U(t)(1-U(t))}{n}\right) = N\left(t, \frac{t(1-t)}{n}\right).$$

Reject the hypothesis that the distribution with CDF $F(x)$ generated the sample at significance level $1 - \alpha$ if

$$U_n(t) \notin \left[t \pm \lambda_{\frac{\alpha}{2}} \sqrt{\frac{t(1-t)}{n}} \right],$$

where λ_α denotes the standard Gaussian α -percentile.

Pearson's chi-squared test The *Brownian Bridge test for distribution*, in the previous paragraph, is based on asymptotic results. Without detailed knowledge about the convergence rate asymptotic results come with a degree of uncertainty. There is a way of performing the same test but without any asymptotics involved.

For any distribution it is true that the probability that a random sample generated by that distribution belongs to the P :th percentile is P percent. Given n independent samples from the same distribution the number of samples belonging to the P :th percentile should follow a binomial distribution with probability P of success and n trials. This property makes the test well suited for testing entire random samples versus a distribution, but also for testing whether specific quantiles agree with the assumed distribution.

Define $N(P)$ as the number of observations in the P :th percentile. X_1, X_2, \dots, X_n are under the null hypothesis random samples from the distribution with CDF $F(x)$. This implies that under the null hypothesis $N(P)$ has the following distribution

$$N(P) = \sum_{i=1}^n I_{[F(X_i) \leq P]} \in \text{Bin}(n, P). \quad (2.50)$$

Reject the hypothesis that the distribution with CDF $F(x)$ has generated the random sample if

$$\left\{ P(k \leq N(P)) < \frac{\alpha}{2} \right\} \vee \left\{ P(k \geq N(P)) < \frac{\alpha}{2} \right\},$$

where k belongs to the binomial distribution in equation 2.50 and \vee denotes a logical "or".

To test not only the tails of a distribution but the entire distribution, one is interested in the performance of the distribution for every single part of the distribution function. A way to test this is to define

$$N(P_1, P_2) = \sum_{i=1}^n I_{[P_1 < F(X_i) \leq P_2]} \in \text{Bin}(n, P_2 - P_1). \quad (2.51)$$

Under the null hypothesis this should be true for all $P_1 \leq P_2 \in [0, 1]$. If the space $[0, 1]$ is split into k non overlapping subintervals covering the entire space there are k different test quantities. To test them all at once the following asymptotic result will be used

$$\text{Bin}(n, p) \xrightarrow[n \rightarrow \infty]{d} N(np, np(1-p)). \quad (2.52)$$

The interesting quantity to examine is the combined absolute deviance from the expected value for all the bins. By using that a sum of k different squared independent standard normal variables are distributed as a $\chi^2(k)$ -variable it is possible to create the test quantity Q for the entire distribution as

$$Q = \frac{1}{n \frac{1}{k} (1 - \frac{1}{k})} \sum_{i=1}^k \left(N\left(P_{(i-1)\frac{1}{k}}, P_{i\frac{1}{k}}\right) - n \frac{1}{k} \right)^2 \in \chi^2(k), \quad (2.53)$$

and reject the null hypothesis at significance level $1 - \alpha$ if $Q > \chi_\alpha^2(k)$. This test is commonly known as *Pearson's chi-squared test*.

Test of distributional moments To test whether an empirical set of independent random variables belong to some distribution, empirically estimated moments can be tested versus true moments. The test is designed such that the estimated moments belong to some distribution. If a confidence interval around the estimated moment does not cover the true value of the moment then the hypothesis, that the empirical set of random variables have been generated by the test distribution, can be rejected. An example is the *Jarque-Bera test* which tests whether sample data have the skewness and kurtosis to match a normal distribution.

2.8.2 Parametric bootstrap

To evaluate the performance of a model without using a testing set, it is possible to evaluate if it is possible that the model has generated the historical data used to fit the model. This can be done by simulating a large number of paths using the estimated model. For each path some statistic is calculated. The same statistic is calculated using the historical time series. This statistic is compared to the empirical distribution of statistics generated by the simulations. The hypothesis that the model have generated the historical data is rejected if the confidence interval created by the empirical distribution function of the generated samples does not cover the statistic calculated from the historical data. This method of testing a parametric model is called parametric bootstrap.

Chapter 3

Method

The method chapter will treat the methods used to derive the results in chapter 4. Starting with the data selection approach, the chapter continues with explaining the data analysis methods, followed by model building and evaluation. The spread-index transition is discussed, as well as dependence structure. The goal of the chapter is to make the reader understand the practical approaches of the thesis better.

3.1 Data selection

3.1.1 Spread data

In order to create a good model, it is vital to have good data to analyse and fit the model to. Choosing the right data, minimizing noise and approximations, is an important part in order to create a robust and realistic model. Our goal is to model credit spreads and indices, and thus data on the given portfolio of bonds are needed. There are plenty of choices when it comes to the type of data that could be used to achieve this. One could for example use market prices on the bonds, a bond index, or the spreads directly. The most convenient way when modeling spreads is, of course, to use the spread data directly if available. In order to move between the credit index and credit spread, one requires information about time to maturity and face value of the bonds, which is not often included in a credit index. It is much easier, and cheaper, to get hold of data on bond portfolios without information about every single bond. A lot of financial institutions offers data on credit spreads, for example Bloomberg and banks such as Barclays or Bank of America.

To ensure good results in statistical analysis, large enough sample size is of the essence. Since the US corporate bond market is the largest and most liquid, the data used is from the US Investment Grade and High yield corporate bond market. The data used in this thesis are taken from Bloomberg Barclays indices.

Most bond portfolios contain bonds that have embedded options, like callable or puttable bonds. One example are mortgages in the US, where the lender can pay their debts early and thus refinance their debt at a lower interest rate. When calculating the yield of the bond this needs to be accounted for, and the resulting yield spread is called an option-adjusted spread. The bond yield can be calculated with Monte Carlo-methods, by simulating when the bond will be sold and to what price, and taking the average. This is the most accurate measure of the spreads when a bond contains embedded options. In the US Investment grade and High yield bond segments there are embedded options present in several bonds, which is why the option-adjusted spread is used in this thesis. Another possibility could be to remove the bonds containing optionality, but this task is nearly impossible since almost every index contains these type of bonds and are rebalanced every month. Another aspect is that the amount of data would be smaller without the optionality bonds, which is not desirable.

The spread data used in this thesis is the following:

Barclays US Corporate High Yield Average OAS Data ranging from 1994 to 2016, containing average option-adjusted spreads for High yield corporate bonds denoted in US dollars. The bond issuers are corporate only and are rated Ba1/BB+/BB+ or below by Moody's, S&P, and Fitch, respectively.

Barclays US Aggregate Corporate Average OAS Data ranging from 1989 to 2016, containing average option-adjusted spreads for Investment grade corporate bonds denoted in US dollars. The issuing sectors are treasury, government-related, corporate, and securitized. The ratings from Moody's, S&P, and Fitch are Baa3/BBB-/BBB- or higher.

Each spread value, for both segments, also has a corresponding modified adjusted duration. This is used to discount the prices of the bonds. The duration data does not cover the entire period for which there is spread data. High yield has the least duration data ranging between 1995 and 2016. Since High yield and Investment grade spreads will be simulated simultaneously, only this period will be considered for recreating the credit indices.

3.1.2 Risk-free rate data

The interest rate data is taken from the US treasury. Both treasury indices and rates with constant maturities are used. The constant maturity rates used are the 3 and 10 year interest rate. Five rate indices are used:

Barclays US Aggregate Treasury Tracking all US treasury securities.

Barclays US Intermediate Treasury Tracking all US treasury securities with a maturity of more than 1 and less than 10 years.

Barclays US 5-7 year Treasury Tracking all US treasury securities, with a maturity of more than 5 and less than 7 years.

Barclays US 7-10 year Treasury Tracking all US treasury securities, with a maturity of more than 7 and less than 10 years.

Barclays US 4-10 year Treasury Tracking all US treasury securities, with a maturity of more than 4 and less than 10 years.

All interest rate indices has a corresponding modified adjusted duration to each value.

3.1.3 Stock index data

In order to analyze dependence structure between asset classes, stock index data is needed. In this thesis the S&P 500 index is used, which is an index tracking 500 large companies that are publically traded in the US, on either the NYSE or NASDAQ stock exchange.

3.2 Data analysis

In order to find a suitable model for credit spreads in the Investment grade and High yield segments, it is important to understand the characteristics of the historical data. By analyzing some statistical features the choice of model will be anchored in theory. Since this thesis is using an autoregressive approach, autocorrelation plays a big part in the evolution of the process. The goal is to capture as much of the structure in the data as possible, so that the only thing remaining is some random white noise, which can be simulated.

3.2.1 Stationarity

A stochastic process is weakly stationary if

$$\begin{aligned} E[X_t] &= \mu, \quad \forall t, \\ E[(X_t - \mu)(X_{t-j} - \mu)] &= \phi(j), \quad \forall t, j, \end{aligned}$$

where $\phi(j)$ is the autocovariance function. This means that the mean function of a wide sense stationary process is a constant, and that the autocovariance only depends on the lag difference, and not the time. An AR process is stationary if and only if the roots of the $A(z)$ polynomial are inside the unit circle, as discussed in Section 2.3.1. Historical data of credit spreads seems to be stationary, since the processes look like they have a constant mean function. However, when modeling the spreads as autoregressive processes, it turns out that they have roots relatively close to the unit circle. This can affect the parameter estimation, which is good to have in mind when modeling these processes.

3.2.2 Data transformation and trends

In order to fit a model to historical data, denoted y_t throughout this section, transformation of said data might be necessary for better results. Transformations can make the data better meet assumptions needed for statistical inference. One usual transformation is the logarithmic transformation given by equation 3.1.

$$z_t = \log(y_t). \quad (3.1)$$

The logarithmic transformation is popular since it can make the process more stationary, make the data more visually interpretable, make the model residuals less skewed, linearize variable relationships and make multiplicative noise additive, to name a few.

Another usual transformation when it comes to financial time series is differentiation of the data, making it stationary. Together with a logarithmic transformation this is commonly known as the log-returns of the time series. The log returns are calculated according to equation 3.2.

$$z_t = \log(y_t) - \log(y_{t-1}). \quad (3.2)$$

Another group of transformations in popular demand with time series modelers are transformations to get rid of linear trends in data. If a linear trend is present in the data, this trend can be removed by simply regressing a straight line on the data, and subtracting the process with this line. The regression line will look like $\alpha_0 + \alpha_1 t$, and the detrended process will be $z_t = y_t - \alpha_0 - \alpha_1 t$, where y_t is some time series in need of detrending. Removing the drift of the data can fix problems due to spurious relationships.

A transformation, often used by statisticians, is the standardization of random variables. It simply means transforming a random process such that it has mean zero and unit variance. The transformation is seen in the equation below.

$$\mathbf{Z} = \frac{\mathbf{X} - \mu}{\sigma}. \quad (3.3)$$

3.2.3 Serial correlation

Autoregressive and moving average processes both exhibit serial correlation, that is, dependence with itself over time. In estimating the model order of AR(p), MA(q) and ARMA(p, q) processes, serial correlation plays a central role.

Autocorrelation function Autocorrelation refers to the correlation of a process with itself, lagged some time-steps. It is defined as

$$\rho(s, t) = \frac{E[(X_t - \mu_t)(X_s - \mu_s)]}{\sigma_t \sigma_s}. \quad (3.4)$$

To estimate the autocorrelation of a process, the sample autocovariance is used. This is defined as $\hat{\rho}(k) = \frac{c_k}{c_0}$, where c_j is the autocovariance estimated as described by equation 3.5 below.

$$c_j = \frac{1}{T-1} \sum_{t=1}^{t-k} (y_t - \bar{y})(y_{t+j} - \bar{y}), \quad (3.5)$$

where \bar{y} is the sample mean.

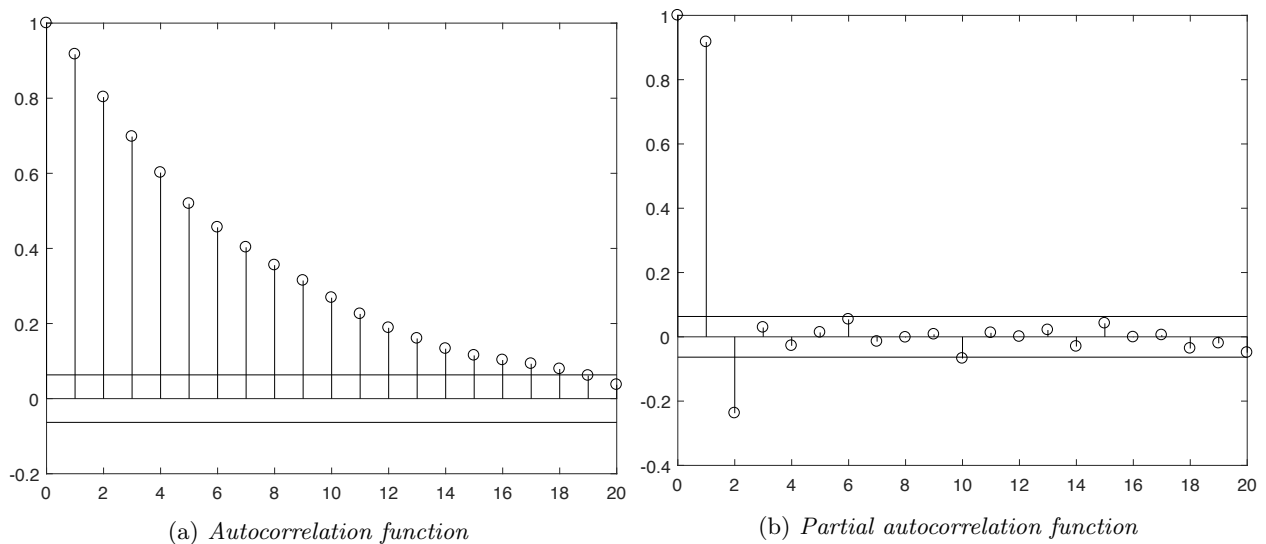
Looking at the autocorrelation of an AR(1) process, it can be viewed as an infinite MA process, since

$$y_t = a_1 y_{t-1} + \varepsilon_t = a_1(a_1 y_{t-2} + \varepsilon_{t-1}) + \varepsilon_t = \dots = a_1^n y_{t-n} + \sum_{k=0}^{n-1} a_1^k \varepsilon_{t-k}. \quad (3.6)$$

This is true for any AR(p) process. Since ε_t is white noise, an MA(q) process will have an autocorrelation structure that is zero after q lags. AR processes thus show an exponentially damped behavior, and/or sine functions, in the autocorrelation function. An example of the ACF for an AR(2) process can be seen below, in Figure 3.1a. The dotted lines denotes the confidence intervals for the function, where it is not statistically different from zero.

Partial autocorrelation function For autoregressive processes, the order can not be determined by the autocorrelation function. Instead the partial autocorrelation function can be used. It measures the correlation between y_t and y_{t+k} after adjusting for linear effects in $y_{t+1}, \dots, y_{t+k-1}$. This is done by fitting AR processes of order $1, 2, \dots, N$ to the process using an OLS approach. After p lags, the PACF will theoretically be zero. The PACF of a simulated AR(2) process can be seen in Figure 3.1b below.

Figure 3.1: *Serial correlations of simulated AR(2) process*



The autocorrelation function looks like a sine wave, and the partial ACF is statistically zero beyond lag 2. By looking at the patterns in both the ACF and the PACF, a reasonable model order for an ARMA process can be chosen.

3.2.4 Conditional heteroscedasticity

Financial time series often exhibit volatility clustering. One way of handling this is by introducing an autoregressive structure in the volatility. This is discussed in Section 2.3.2.

Engle's ARCH test To test for ARCH structures in the volatility, one can use Engle's ARCH test, described in Engle (1988). The test uses the residuals after fitting the data to the best AR(p) model. After this, a regression on a constant and lagged values are made, as such

$$\hat{r}_t^2 = \alpha_0 + \sum_{k=1}^q \alpha_k r_{t-k}^2. \quad (3.7)$$

If any of the parameters α_i is statistically different from zero, there are ARCH-effects in the data.

Visual assessment Another way to look for volatility clusters is by simply plotting the model residuals. Optimally, the residuals should show no autocorrelation, and no clusters of high or low values. Just looking at the plots of the residuals can reveal a lot of information about its structure.

3.3 Model building

3.3.1 Choice of model

One of the main problems with credit modeling is the lack of reliable data. Defaults are rare events and on firm specific level defaults are non existing. A good model is therefore one that is simple enough for accurate calibration, but at the same time captures as much of the structure as possible.

One way to model defaults could be to use credit ratings, or more specifically, a credit transition matrix based on historical defaults. Credit ratings are robust and could be used for predictions over long time horizons. However, since the transition matrix are entirely based on empirical analysis of historical events, it contains no information about the future. Therefore, it is far from certain that the constituents in a given credit segment has the same probability of default today as the historical defaults of the segment. The price given by the market, assuming efficiency, ought to be based on all information available including the credit ratings but also the common opinion about future events. When predicting future credit index levels it is therefore a better choice to build a model based on market prices than using credit ratings. Popular models that can be calibrated to market data are Structural, Hazard rate, and Credit spread models.

There is little literature on autoregressive credit models, and few valuable comparisons to Structural or Hazard rate models. Therefore, the choice of model will be based on theoretical arguments.

The Structural models are in theory the most intuitive. They are based on the assumption that a company's assets develop like a geometric Brownian motion and that default occurs if the assets are worth less than the liabilities at maturity. The recovery rate is then determined as the total value of the assets at maturity. In practice there are several problems with the structural models. The first being that the development of the price often is more volatile than what can be explained by the Brownian motion driving the asset value (Manzoni, 2002). Another problem is that Structural models assume that the assets are observable. For a specific company it is of course possible, but for an entire credit index, observing the asset values is an unmanageable task. Furthermore, in the case of default it is unlikely that the asset value represent a credible recovery rate (McNeil et al., 2005).

Hazard rate models handle defaults as jumps in a Poisson process. The recovery rate is either modeled as a stochastic variable or chosen as an historical average of recovery rates. The recovery rate has a great impact on the price and is very difficult to model. However, for an entire index the historical average should be a valid estimate. Default intensities are modeled using stochastic differential equations with Wiener increments. Just as in the case with Structural models these increments can not empirically explain the variability of the price (Manzoni, 2002).

For both Structural and Hazard rate models there are ways of handling all previously mentioned problems, but most of them result in an increased complexity. One approach however, using an autoregressive credit spread model, handles all the problems whilst keeping the model simple.

The market price is made up of risk neutral expectations of default probabilities and recovery rates, but tells nothing about what portion of it that comes from default probability and what comes from recovery rates. Instead of trying to figure this out, like when using a Hazard rate model, a credit spread model will simply model the price as described in equation 3.8 and not separate the two.

$$P(t, T) = e^{-(r+y)(T-t)}. \quad (3.8)$$

If one disregards the recovery rate a credit spread model can be viewed as a type of default intensity model in discrete time. Although such a model lacks immediate theoretical underpinning it should in theory be able to capture all structure in the price. There is no extensive literature on modeling of credit spreads, see Section 1.2. There is, however, a lot of literature on modeling log returns of stock prices. By adding an AR term to those models the theory can be directly translated to modeling of credit spreads. Such an autoregressive credit spread model handles all problems that the Structural and Hazard rate models can

not, whilst keeping the model simple, easy to calibrate and robust. Therefore, it is chosen as the model of interest in this thesis.

3.4 Model evaluation

A number of models will be evaluated in this thesis. It is not uncommon that more than one model passes the residual tests described in Section 2.8.1. There are a number of ways of comparing these models to see which one is the preferred model for our purposes. The methods used are described in this section.

Likelihood ratio test If a model is nested in some larger model, as in the case of AR and AR-GARCH models, the likelihood ratio test can be used to assess goodness of fit. Given a nested statistical model, the null and alternative hypothesis are given by the parameter θ being in a given parameter space, as such

$$\begin{aligned} H_0 &: \theta \in \Theta_0, \\ H_1 &: \theta \in \Theta_0^c. \end{aligned}$$

The likelihood ratio is defined as

$$\Lambda(x) = \frac{\sup\{\mathcal{L}(\theta|x) : \theta \in \Theta_0\}}{\sup\{\mathcal{L}(\theta|x) : \theta \in \Theta\}}, \quad (3.9)$$

where \mathcal{L} is the likelihood function, $\mathcal{L}(\theta|x) = P(x|\theta)$. Following the results from Wilks (1938), the test statistic $-2 \log(\Lambda)$ under H_0 asymptotically follows a χ^2 distribution with degrees of freedom corresponding to the extra parameters in the unrestricted model. Following this methodology, a statistic can be calculated with a corresponding p -value if the null hypothesis should be rejected or not.

Akaike Information Criterion The Akaike Information Criterion, AIC, is a measure of how a general model compares to other models, for a given data set. The measure will not tell whether a model is good or not, just if it is better or worse than another model. It uses the likelihood function and the number of parameters in the given model. AIC is defined as

$$AIC = 2k - 2 \log(\hat{L}), \quad (3.10)$$

where k is the number of parameters in the model, and \hat{L} is the maximized value of the likelihood function.

Bayesian Information Criterion BIC is closely related to AIC, and is defined as

$$BIC = \log(n)k - 2 \log(\hat{L}), \quad (3.11)$$

where k is the number of free parameters, n is the number of data points, and \hat{L} is maximized value of the likelihood function. This measure penalizes a large number of parameters more than the AIC, and is thus less prone to overfitting models.

3.5 Index model

The credit index will be recreated from the rate and the credit spread as described by equation 3.12.

$$P_t = P_{t-1}(1 + D_t(\rho - 1)) \frac{e^{-(r_t + y_t)\tau}}{e^{-(r_{t-1} + y_{t-1})\tau}} e^{(r_t + y_t + r_{t-1} + y_{t-1})\frac{1}{24}}, \quad D_t \sim Bin(q_t, n), \quad (3.12)$$

where q_t is the P -probability for default during one month, n is the number of index constituents, ρ is the historical average recovery rate and P_t denotes the price of the index at time t .

The P -default intensity should theoretically be calculated as equation 2.35, $\lambda^P = \lambda^Q - rp$. However, since some simulated credit spread paths might generate \mathbb{Q} -hazard rates smaller than the historical average

risk premium, using equation 2.35 could result in negative default probabilities. Therefore the P -default intensities will be approximated in another way. The approach is not theoretically completely correct, but will compared to historical data, yield the same results as equation 2.35 and never generate negative default intensities for simulated spreads. This approach assumes that

$$\lambda^P = \lambda^Q - rp = \lambda^Q - \alpha\lambda^Q = \gamma\lambda^Q, \quad (3.13)$$

where α is a constant and $\gamma = 1 - \alpha$. γ is estimated estimated from historical data as

$$\gamma = 1 - \frac{\log\left(\frac{P_T^{Credit} P_0^{Treasury}}{P_0^{Credit} P_T^{Treasury}}\right)}{\frac{1}{12} \sum_{i=1}^{12T} \lambda_i^Q}. \quad (3.14)$$

For a derivation of the above equation see appendix C.1.

3.5.1 Approximation errors

If the historical spreads and rates are used to recreate the historical credit index the results will be similar but not identical. The difference is a result of several factors:

In the historical index, defaults have occurred that will affect the index value. Since the constituents of the index is not available in the data, no default data is available either. One factor affecting the difference is thus unknown historical defaults. Another is the use of monthly rate and spread data, instead of daily. This can cause approximation errors. A third factor is that the duration of the index does not completely describe the time to maturity of the underlying bonds, which can cause approximation errors in the discounting of the price. One additional factor has to do with interest rates. The price depends on both the spread and the risk-free rate. However, finding the risk-free rate matching each spread value can be complicated, since the duration does not completely represent the maturity of the bond. This is further discussed in Section 3.5.2 below. The effects of these approximations are hard to examine, but an effort will be made to minimize approximation errors as much as possible, as well as locate the factors that contribute the most to the discrepancies between recreated and historical indices. The results can be seen in Section 4.7.

To eliminate errors due to unknown historical defaults treasury indices will be examined. These will be recreated using equation 3.15.

$$P_T^{Treasury} = P_{T-1}^{Treasury} \frac{e^{-r_t\tau}}{e^{-r_{t-1}\tau}} e^{(r_t+r_{t-1})\frac{1}{24}}. \quad (3.15)$$


The impact of daily versus monthly data will be examined by recreating an index using daily data as well as monthly data and comparing the results. Thereafter, the impact of the unknown maturities of the bonds within the index will be examined by recreating the index using monthly yield curve rates for different maturities and comparing the results.

One approach to examine the impact of defaults on the credit indices, is to recreate the indices without using historical rates as in equation 3.12, but instead using treasury indices, as in equation 3.16. In this way errors with respect to the historical rates will be minimized and only errors resulting from the credit spread and the defaults will be seen.

$$\frac{P_T^{Credit}}{P_{T-1}^{Credit}} = (1 + D_t(\rho - 1)) \frac{P_T^{Treasury}}{P_{T-1}^{Treasury}} \frac{e^{-y_t\tau}}{e^{-y_{t-1}\tau}} e^{(y_t+y_{t-1})\frac{1}{24}}, \quad D_t \sim Bin(q_t, n) \quad (3.16)$$

Calculating $1 + D_t(\rho - 1)$ from the equation above should give a value between 0 and 1, indicating the fraction of the index which has not defaulted. As seen in Figure 3.2 below, this is not the case. Our interpretation is that approximation errors makes the data too noisy to be able to say anything about the historical unknown defaults. Instead, one has to estimate the defaults by testing the simulated indices against historical data. In Section 3.5.3, the approach to model defaults within indices is explained further.

Figure 3.2: The figures show the fraction of the HY and IG indices that remains after credit events have been accounted for during each time step. The fractions calculated using equation 3.16. In theory the fraction should never be higher than one, but as can be seen the value oscillates around one. Our interpretation is that our approximation yields so much noise that it is impossible to extract the credit event data this way.



3.5.2 Finding matching rates

In order to simulate credit index values, the risk-free rate corresponding to each spread value is needed. The data used in this thesis consists of spread values with corresponding modified adjusted duration. The risk-free rates are taken from the US government's Daily Treasury Yield Curve Rates. Rate indices from US Treasury rates with corresponding modified adjusted duration are also available, as seen in Section 3.1.2. Since the modified adjusted duration does not entirely capture the maturity of the bonds, extracting the right rate becomes a more complicated matter. The methodology consists of two steps. Firstly, creating a risk-free rate index corresponding to each duration value for the spread. Secondly, creating a rate that, when reconstructing an index, gets as close as possible to the index from the first step. A more detailed explanation for each step is provided below.

Creating a risk-free rate index First off, a rate index is constructed, such that the modified adjusted duration of the rate index corresponds to the spread. Since we have access to rate indices with higher modified adjusted duration than the spreads, as well as indices with lower duration, a linear combination of the existing rate indices are calculated such that the modified adjusted duration matches that of the spread.

$$\text{index}_t = c_t \cdot \text{index}_{1,t} + (1 - c_t) \cdot \text{index}_{2,t}, \quad c \in (0, 1). \quad (3.17)$$

Creating the rate The rate is constructed in a similar way as the rate index. A linear combination of the 3 and 10 year constant maturity treasury rate is used, as such

$$r = c \cdot r_{10} + (1 - c) \cdot r_3. \quad (3.18)$$

c is calculated as the constant that minimizes the squared error between the rate indices constructed from equation 3.17 and 3.19 below. The index is constructed as

$$\text{index}_t = \text{index}_{t-1} e^{-(r_t - r_{t-1}) \text{mod}_{t-1} + \frac{r_t + r_{t-1}}{24}}, \quad (3.19)$$

where mod is the modified adjusted duration. The 3 and 10 year rates are simulated using an AR(1) model, and the rates corresponding to IG and HY spreads are created using c above.

3.5.3 Modeling historical defaults within index

The defaults within an index will be modeled using a binomial distribution, where the probability of success, representing the probability of default within one time step for a specific company, will be calculated using the credit spread approach described in Section 2.5.2, instead of the credit ratings method. The ratings method will result in a constant probability for future defaults, which is unrealistic. It also demands that the proportion of different rating classes within the index is known and compensated for at all times. Therefore the spread approach is both a better and simpler choice.

The main problem with using the credit spread approach is that the recovery rate is crucial and usually unknown. To solve this the recovery rate is assumed to be constant and chosen such that historical predictions will be as accurate as possible.

The binomial distribution used for simulations is defined as a probability for default as well as a number of trials representing the number of companies within the index. Initially this number will be chosen as the actual number of companies, but if it turns out that this will result in incorrect volatility of the index the number of trials can either be increased to lower the variance or decreased to increase the variance of the simulated index.

3.6 Asset class dependence

As discussed in Section 2.6 it is possible to introduce dependence between simulation of different assets by correlating the noise in the asset models. In this thesis, the dependence structure between credit spreads for the IG and HY segments, two riskfree rates and the S&P500 index is analyzed.

The risk free rates are a crucial part of modeling the credit indices and the correlation between rates and spreads is therefore needed in order to simulate credit indices. The S&P500 index is not needed in order to simulate credit indices, but since the dependence between stocks and bonds should theoretically be high, the correlation with the S&P500 index is modeled to test whether the coupled innovation approach is sufficient to model this dependence.

An important aspect when choosing what copula to use for the dependence structure is the amount of tail dependence in each asset class. To investigate this one can simply look at some chosen quantile level q , and calculate

$$P(\mathbf{X} > F_{\mathbf{X}}^{-1}(q) | \mathbf{Y} > F_{\mathbf{Y}}^{-1}(q)), \quad (3.20)$$

$$P(\mathbf{X} \leq F_{\mathbf{X}}^{-1}(q) | \mathbf{Y} \leq F_{\mathbf{Y}}^{-1}(q)), \quad (3.21)$$

for the upper and lower tail dependence respectively. As it turns out, the historical data shows tail dependence, which is not present in the Gaussian copula dependence structure. Therefore, the t -copula or the grouped t -copula is a better choice. Fitting a Student's t -copula to the different groups individually suggests that the numbers of freedom for each group is not that different and that also the Student's t -copula is a plausible choice.

To investigate whether the grouped t -copula or the standard t -copula is a better fit to the data AIC and BIC will be calculated and compared. Lastly, to see if the copula could actually represent the true dependence structure parametric bootstrap will be used to calculate a confidence interval for the log-likelihood of the sample. Also, parametric bootstrap will be used to see if the upper and lower tail dependence of the log-returns of the simulated indices match that of the historical data.

3.6.1 Estimation

After having chosen a copula, the parameters need to be estimated. In the case of a grouped t -copula this is done in two steps. First, estimation of the correlation matrix is done. As shown in Daul et al. (2003) an approximation of the elements in this matrix can be constructed as shown in equation 3.22. Let $X_i = \sqrt{W_1}Z_i$ and $X_j = \sqrt{W_1}Z_j$ just as in Section 2.6.1.

$$\tau_{\text{Kendall}}(X_1, X_2) \approx \frac{2}{\pi} \arcsin(\rho) \quad (3.22)$$

where ρ is the correlation between Z_i and Z_j , and τ_{Kendall} is Kendall's τ coefficient, defined below.

Definition - Kendall's tau Kendall's τ is defined as the difference between the probability of concordance and the probability for discordance between two two-dimensional random vectors, Nelsen (2006).

$$\begin{aligned}\tau_{\text{Kendall}} &= P(\text{concordance}) - P(\text{discordance}) \\ &= P((X_1 - X_2)(Y_1 - Y_2) > 0) - P((X_1 - X_2)(Y_1 - Y_2) < 0).\end{aligned}\tag{3.23}$$

For a random sample an unbiased estimator of Kendall's tau is given by:

$$\tau_{\text{Kendall}} = \frac{c - d}{c + d} = \frac{c - d}{\binom{n}{2}},\tag{3.24}$$

where c is the number of concordant pairs and d is the number of discordant pairs.

After this, maximum likelihood estimation is used to estimate the parameters ν_k for each subgroup. In order to do so, the pdf of the grouped t -copula is required. The pdf is derived by first finding the cdf of the grouped t -distribution and using this to derive the pdf of the copula, see appendix B.1 for details. If all subgroups within the grouped t -copula have at least two members, then the usual t -copula can be used to find the ν corresponding to each subgroup and there is no need for deriving an expression for the pdf of the grouped t -copula. The results will however be suboptimal. A third approach is to find the degrees of freedom for each group using the t -copula approach and use the values as starting values for the maximum likelihood estimation using the grouped t -copula.

3.6.2 Simulation

The procedure of generating random variables with a dependence structure given by a Student's t -copula is described below. SIGMA denotes the correlation matrix, ν_k denotes the degrees of freedom corresponding to group k and F_d^{-1} is the inverse CDF corresponding to the marginal distribution of the d :th variable.

1. **Generate a uniform random variable, U , from the unit set.**
2. **Generate multivariate Gaussian random variables, $Z = [Z_1, \dots, Z_d]$, with zero mean and covariance matrix SIGMA.**
3. **For each subgroup k , set ν_k**
4. **Generate Inverse-gamma distributed random variables, $W = [W_1, \dots, W_d]$.** Do this by using the inversion method, such that $W_i = G_{\nu_i}^{-1}(U)$, where $G_{\nu_i}^{-1}$ denotes the inverse cdf of the Inverse-gamma distribution with ν_i degrees of freedom.
5. **Generate random variables, C , from the copula.** Generate the random variables by setting $C = [C_1, \dots, C_d] = [Z_1 W_1, \dots, Z_d W_d]$.
6. **Transform the copula values using the marginal distributions.** The transformation to the marginal distributions F_i for each asset is performed in the following way:

$$[\mathbf{X}_1, \dots, \mathbf{X}_d] = [F_1^{-1}(\mathbf{U}_1), \dots, F_d^{-1}(\mathbf{U}_d)]\tag{3.25}$$

This is expressed in Matlab code below.

```

1  U = rand(n,1)
2  MU = zeros(k,1)
3  Z = mvnrnd(MU,SIGMA,n)
4  V = [v1, v1, ..., v2, v2, ..., vk ];
5  for i = 1:length(V)
6      W(:,i) = V(i)./chi2inv(U,V(i));
7      Y(:,i) = sqrt(W(:,i)).*Z(:,i)
8      copula(:,k) = tcdf(Y(:,i),V(i));
9      simulation(:,i) = finv_i(copula(:,i))
10     % finv_i denotes the marginal distribution quantile function for
11     marginal i.
end

```

3.6.3 Calibration

To calibrate the credit spread models the maximum likelihood method will be used. The starting values will be found by first estimating the parameters of the AR model without the volatility model using a Kalman filter. These estimates will be set as the starting values for the maximum likelihood estimation of the entire model. Parameters that were not estimated in the pure AR model will be set to zero, yielding valid starting values. For updating the parameters when new information is available it is suggested that maximum likelihood is used and that the starting values are chosen as the previously estimated ones.

3.7 Model validation

A first step to validate the model for credit indices is to validate the credit spread models on historical credit spreads. If the credit spread models passes the tests the entire index models will be evaluated on historical credit indices. To use all data to calibrate the models, and not waste any on testing, the parametric bootstrap approach described in Section 2.8.2 will be used to validate the models. Apart from passing the statistical tests the simulated paths should never explode, get stuck at low values or display any other unlikely behaviours. To examine this a number of simulated paths will be plotted and visually inspected.

The simulation method requires that some statistic is calculated from the historical data. The interesting statistics differ depending on if the credit spread model or the index model is to be tested. In the sections below the methods of validating the credit spread model as well as the index model are listed.

3.7.1 Credit spread model validation

When finding the best model the residuals for different models are analyzed with respect to distribution and independence. If several models have white residuals the model which performs the best with respect to the number of parameters used and log-likelihood of the residuals is chosen. The problem with choosing the model that maximizes the log-likelihood is that the likely residuals will have much more influence than the unlikely ones on the estimation. This leads to a good fit for most values but not necessarily for the tails of the distribution. Therefore when validating the model on simulated data the effect of badly fitted tails is interesting to examine.

Results of badly calibrated tails could be that the simulated credit spread displays unrealistically big negative or positive shocks, not seen in the historical data. Since the volatility process is a function of these shocks, the result could either be that the credit spread gets stuck at low values or reaches really high values with high volatility. To put these behaviours to the test, relevant statistics to choose for the bootstrap are average values, as well as variances for some different time intervals. The average value test of the process makes sure that the process does not get stuck or explode too often. It is possible that the process would sometimes get stuck at a low value after an explosion yielding good average values. To test that this is not the case, the variance of the average values will also be examined.

It could be interesting to also examine the distribution of historical spreads compared to simulated ones.

All distributional tests described in Section 2.8.1 rely heavily on independence between the observations. Since the autocovariance between spreads is really high this means that consecutive spreads can not be used to perform the tests. As can be seen in Figure 4.1a and 4.1c, the autocovariance is significantly different from zero for about 40 lags for Investment grade and 20 lags for High yield. This means that every 20th or 40th observation can be used which will result in too little data to achieve any significance in any of the distributional tests.

3.7.2 Index model validation

The model for the index is a combination of models for the credit spread, the interest rate and credit events within the index. For one and each of these models the marginal distribution have been examined. The models does of course not explain their respective processes fully and will result in individual errors. The entire model for the index will thus combine these errors. The index is not a stationary time series and therefore all tests will be done on the log returns of the index.

The log returns of the index are functions of the credit spread and the interest rate, see equation 3.26. For different simulated paths of the credit spread and the interest rate their sum might be different from what is seen in the historical data. Therefore, if performing a test whether the log returns from the simulation and those of the historical data have been generated by the same distribution, the mean value of the log returns from the simulation will be different from that of the historical data. The same argument can be used to argue why the hypothesis of the same generating distribution will be rejected based on different variances. Therefore to test the distribution of the log returns, first the mean value and standard deviation will be examined using parametric bootstrap. Then the returns for each simulated path will be normalized by subtracting the mean and dividing with the standard deviation.


The test of generating distribution will be performed by creating an empirical distribution function using all normalized simulated log returns. Since the number of simulated log returns is so big, this distribution will be very close to the actual generating distribution. To test whether this distribution could have generated the normalized historical log returns the *Brownian bridge test for distribution* and the *Kolmogorov-Smirnov test* will be performed. Since the credit events within the index will have more impact on the log returns for longer return periods, see equation 3.26, a longer return series will be examined as well.

$$\log \left(\frac{P_t^{Credit}}{P_{t-k}^{Credit}} \right) = \sum_{i=0}^{k-1} \log (1 - D_{t-i}(1 - \omega)) + \tau(r_{t-1} + y_{t-1} - r_t - y_t) + \frac{1}{12} \sum_{i=1}^{k-1} r_{t-i} + y_{t-i} \quad (3.26)$$

3.7.3 Parametric bootstrap

The model which is being validated will be used to generate 10 000 independent paths. Each of these paths will represent a 248 months real world simulation, which is the same length as the historical data. For statistics calculated with a time series of independent observations a rolling window will be used as well as a separate window. Using a rolling window means that the samples evaluated are created using overlapping segments of the time series. A comparison between separate windows and rolling windows can be seen in Figure 3.3 below. Validation based on rolling windows is, for instance, used in Committee of European Insurance and Occupational Pension Supervisors (2010), and as can be seen in Appendix D.1 the estimates using rolling windows will have about the same variance as if not using them, but making sure that all data samples are used. For dependent time series statistics, like the mean and variance of the credit spread paths, there is a chance that the rolling window method will impose too much dependence between the calculated statistics for each path. Therefore these statistics will be evaluated using separate windows. Also, all distributional tests will be performed without using a rolling window since these tests are sensitive to dependence.

Figure 3.3: Difference between separate and rolling window approach when calculating returns.



(a) Example of rolling windows, 3 month returns

(b) Example of separate windows, 3 month returns

Chapter 4

Results

In this chapter, results are derived using the methods listed in the previous chapter. The chapter initially evaluates the data and uses the results for model selection for the credit spread models. A number of models are calibrated and compared. After selecting the best models the results of the models will be presented. Credit indices are recreated using simulated spreads and rates. The approach of going from credit spreads and rates to a credit index is finally evaluated and the validation results are presented in the end of this chapter. The dependence between the credit spreads, the rate and the S&P500 index will then be modeled and tested both on spread and index level.

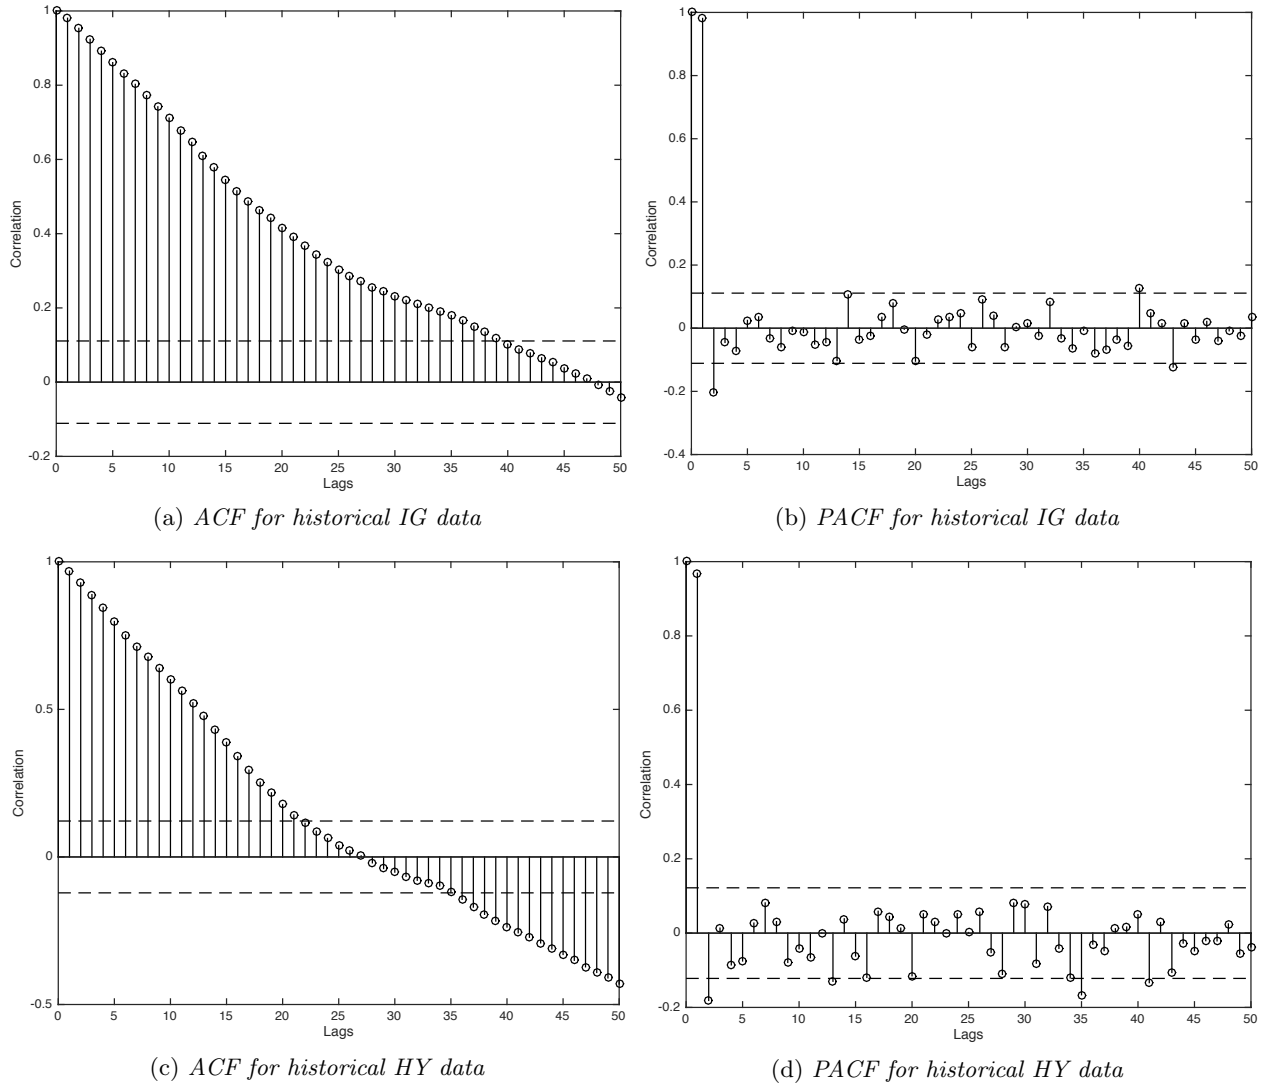
4.1 Data analysis

Analyzing the data is a vital part of creating any model. In this section the data will be analyzed according to the methods described in the previous chapter.

Data transformation In order to make the spread data more stationary and ensure positiveness of simulations, a logarithmic transformation was made. After this, the process was demeaned in order to skip the estimation of the constant in the AR process. Differentiating the data did not improve results for the models tested, and this transformation was therefore not used.

Serial correlation The ACF and PACF of the transformed spread data for Investment grade and High yield segments respectively can be seen in Figure 4.1. As can be seen the PACF for both Investment grade and High yield drops to zero after lag 2, indicating AR(2) processes.


Figure 4.1: Autocorrelation and partial autocorrelation for historical data for both Investment grade and High yield segments. The exponential decay in the ACF and the sudden drop to zero in the PACF indicates an autoregressive structure



4.1.1 Conditional heteroskedasticity

Since the AR(2) model was suggested by the initial data analysis, this model was fitted to the historical data. A normal probability plot of the residuals suggested a distribution with fatter tails but with the same bell curved shape as a normal distribution. Therefore the t -distribution was used. This significantly increased the likelihood. The residual plots for both the Investment grade and High yield segments can be seen in Figure 4.2.

Figure 4.2: *Residuals after fitting an AR(2) model with t -distributed noise to the historical Investment grade and High yield data.*



For neither one of the Investment grade and High yield segments ARCH effects were found, after conducting Engle's ARCH test on the residuals from the AR(2) model with t -distributed noise. However, by looking at the residuals from the pure AR model some volatility clustering was seen, leading to the conjecture of some other volatility structure. Since the pure AR model is a nested model of an AR model with a time-dependent noise structure, one can use the likelihood ratio test to investigate whether some volatility model will be yield a better fit to the data. Additionally, if the pure AR model is wrong, the parameters might not be very good estimates, resulting in residuals not showing any ARCH structure, even if the true model have such residuals. To further test for volatility clustering, a number of EGARCH models were tested and the goodness of fit was evaluated with the likelihood ratio test. The models considered are described in Section 4.2.

4.2 Model comparison

All credit spread models evaluated have an autoregressive structure. It turns out that the Investment grade spread was best modeled using an AR(2) model and the High yield spread using an AR(1), this contradicts Figure 4.1d saying that an AR(2) would be the best fit for both segments, but after adding a time-dependent volatility the AR(2) term for High yield was reduced to insignificance. All models tested from here and on will be AR(2) and AR(1) models with time-dependent volatility structure for IG and HY respectively. The volatility models will belong to the EGARCH-family. The reason for the choice of EGARCH structure in favour of regular GARCH is that it ensures positiveness of the standard deviation and that it is easy to impose skewness. In the following sections the tested models are presented. To start off with the results using a pure AR model will be presented.

4.2.1 AR model

For pure autoregressive models the AR(2) model is the best fit for both the Investment grade and the High yield segment. Comparing Gaussian and t -distributed noise, the model with t -distributed noise had a significantly higher likelihood, see Table 4.1 for the log likelihood using t -distributed noise. The model structure can thus be described by the following equation.

$$\log(y_t) = \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \sigma \varepsilon_t \quad (4.1)$$

$$\varepsilon \sim t(\nu) \quad (4.2)$$

Model	AR(2)
Investment grade	367.2291
High yield	258.5000

Table 4.1: Log-likelihood for a pure AR(2) model for Investment grade and High yield.

4.2.2 Leverage models

In financial time series, the impact of process shocks on the volatility are typically larger in one direction than the other, as discussed in Section 2.3. Credit spreads for both Investment grade and High yield seem to have this behavior, with positive shocks having a larger impact than negative. There are a number of ways to model this behavior, and in this thesis three different approaches are proposed.

Model 1

$$\log(\sigma_t^2) = \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 \varepsilon_{t-1} \quad (4.3)$$

In this model, the leverage term comes from $\alpha_2 \varepsilon$, where a positive sign of ε will increase the volatility whereas a negative will decrease it.

Model 2

$$\log(\sigma_t^2) = \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 (|\varepsilon_{t-1}| + \varepsilon_{t-1}) + \alpha_3 (|\varepsilon_{t-1}| - \varepsilon_{t-1}) \quad (4.4)$$

This model separates the impact of negative and positive shocks and is thus more flexible than **Model 1**.

Model 3

$$\log(\sigma_t^2) = \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 (|\varepsilon_{t-1}| - E[|\varepsilon_{t-1}|]) \quad (4.5)$$

The last term of this model gives big shocks more impact on the volatility than small regardless of the sign of the shock. $E[|\varepsilon|]$ can be calculated for a t -distribution as $\sqrt{\frac{\nu-2}{\pi} \frac{\Gamma(\frac{\nu-1}{2})}{\Gamma(\frac{\nu}{2})}}$.

Log-Likelihood of the models The log-likelihood for the different models are shown in Table 4.2 below.

Model	1	2	3
Investment grade	380.1344	380.9250	373.1279
High yield	263.4793	264.0897	259.6112

Table 4.2: Log-likelihood for leverage models.

Looking at BIC values, **Model 1** is the best choice for both Investment grade and High yield. The AIC values are better for **Model 1** as well for Investment grade, but not for High yield. However, the overall assessment of both AIC and BIC values for the High yield segment makes the choice lean towards **Model 1** as well.

4.2.3 Process value models

It is reasonable to assume that the volatility is somehow proportional to the level of the spread. To adjust for this it is possible to include a process term in the volatility structure. The following models were tested.

Model 4

$$\log(\sigma_{t-1}^2) = \omega + \alpha y_{t-1} \quad (4.6)$$

Model 5

$$\log(\sigma_{t-1}^2) = \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 y_{t-1} \quad (4.7)$$

The only difference between **Model 4** and **Model 5** is the GARCH-term, which makes the volatility depend on the volatility level in the previous time step. In Table 4.3 below, the log likelihoods for the two models are listed. The α_1 parameter of **Model 5** is insignificant for both Investment grade and High yield.

Model	4	5
Investment grade	381.6756	382.1676
High yield	260.4694	261.1439

Table 4.3: *Log-likelihood for process input models.*

Since these two models are nested it is possible to use the likelihood ratio test to evaluate if the improvement of adding additional parameters is significant. The test concludes that the improvement of adding an additional to **Model 4**, creating **Model 5**, is insignificant for both High yield and Investment grade.

4.2.4 Combination models

By selecting the volatility models with the highest likelihood, based on the model performance in the two sections above, two combination models containing both leverage and process terms are tested. These are presented below.

Model 6

$$\log(\sigma_t^2) = \omega + \alpha_1 \varepsilon_{t-1} + \alpha_2 y_{t-1}. \quad (4.8)$$

Model 7

$$\log(\sigma_t^2) = \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 \varepsilon_{t-1} + \alpha_3 y_{t-1}. \quad (4.9)$$

Model	6	7
Investment grade	383.2053	388.6637
High yield	262.8008	264.0360

Table 4.4: *Log-likelihood for combination models.*

Model 6 is a nested in **Model 7**. The likelihood ratio test concludes that **Model 6** is better for the High yield segment, and that **Model 7** is preferred when it comes to the Investment grade segment.

4.2.5 Choice of model

For a comparison of log likelihood, AIC and BIC values for all models tested this far see Table 4.5. The models containing process values perform better than the pure leverage models for the Investment grade segment. However, as will be seen in Section 4.3, the models containing process values, including the combination models, will not produce stable simulations. For High yield, the pure leverage models are the better ones, looking at the likelihood. All in all, also comparing to a pure AR model, **Model 1** is the model giving the best results.

Model	AR(2)	1	2	3	4	5	6	7
# Parameters(IG/HY)	3/3	5/4	6/5	5/4	4/3	5/4	5/4	6/5
Investment grade								
Log Likelihood	367.23	380.13	380.93	373.13	381.68	382.17	383.21	388.66
AIC	-728	-750	-750	-736	-755	-754	-756	-765
BIC	-717	-731	-727	-717	-740	-735	-738	-743
High yield								
Log Likelihood	258.5	263.48	264.09	259.61	260.47	261.14	262.8	264.04
AIC	-511	-517	-516	-509	-513	-512	-516	-516
BIC	-500	-499	-495	-491	-499	-494	-498	-495

Table 4.5: The table shows the log likelihood along with AIC and BIC values for all models tested this far. The best AIC and BIC values for each model are highlighted with bold letters. It can be seen that **Model 7** performs the best for IG while either **Model 1** or the pure AR 2 model works the best for HY. A combined assessment of AIC and BIC favours **Model 1**.

4.3 Model evaluation

Residuals from all models were tested and the choice of model was based on the AIC, BIC and likelihood ratio when possible, as described in Section 4.2.5. Since the models will be used for simulation purposes a number of simulation tests were performed, see Section 3.7.1 for details. It turns out that all estimated models resulted in exploding simulations. By performing simulations where the degrees of freedom of the t -distribution was increased it could be seen that this significantly decreased the explosion probabilities. However, whilst keeping the degrees of freedoms as low as required in order for the residuals to pass the distributional tests no models with the log credit spread as input to the volatility process stopped exploding for the Investment grade models. Therefore all models with the credit spread as input were dismissed in the following for Investment grade.

The model that performed the best out of the remaining models was kept and tests to choose the best possible numbers of freedom were performed. For both High yield and Investment grade the model that performed the best when the degrees of freedom was kept fixed was **Model 1**.

4.3.1 Finding the degrees of freedom

When finding the degrees of freedom, ν , corresponding to the noise distribution, one desired property is that the spread does not explode. The variance of this parameter is relatively high, pointing to a high uncertainty in the maximum likelihood estimation. In order to estimate a good number for the ν parameter, which generates non-exploding spreads, the model was calibrated using a number of different ν 's. For every ν a p -value, for the hypothesis that the assumed distribution had generated the residuals, was calculated. Finally, simulations were examined by simulating 10 000 paths over 30 years and counting the number of exceedances over some explosion threshold. This threshold is assumed to be three times higher than the maximum historical spread value. That the spread should reach this value in 30 years seems unlikely. The threshold values was calculated to 20 for the IG segment, and 60 for HY.

As can be seen in Figure 4.4a and 4.4b. The best choice of ν for the High yield segment seems to be somewhere between 12 and 18, and somewhere between 10 and 12 for the Investment grade segment, since these are the degrees of freedom that maximizes the p -values from the tests. Figure 4.3 shows that these degrees of freedom can be used without the simulations exploding very often. For High Yield the degrees of freedom was chosen to 15 and for Investment Grade the degrees of freedom was chosen to 11. As it turns out, the degrees of freedom has a relatively small impact on the statistics of the spread, unless the spread explodes to very high values. This is fortunate, since the estimate of ν is relatively uncertain. By selecting a higher number for ν , especially for the Investment grade segment which is more prone to explosions, the spread is more likely to stay on a reasonable level.

Figure 4.3: Percentage of paths with a maximum value larger than an explosion threshold, which is 20 for IG and 60 for HY. The number of simulated paths are 10 000, and the length of each path is 325 and 270, for IG and HY respectively.

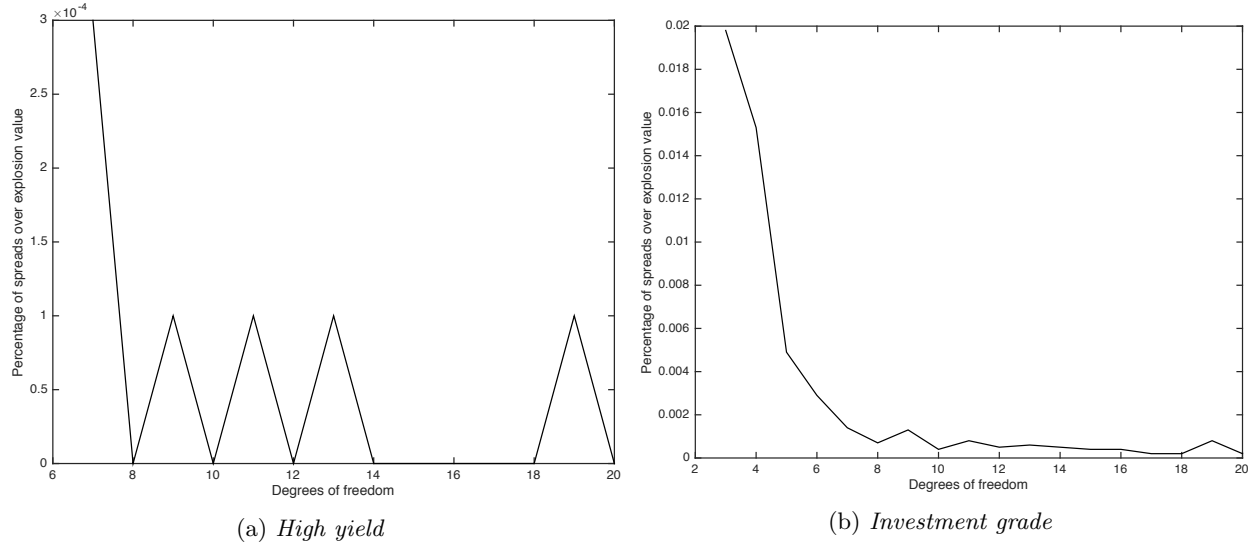
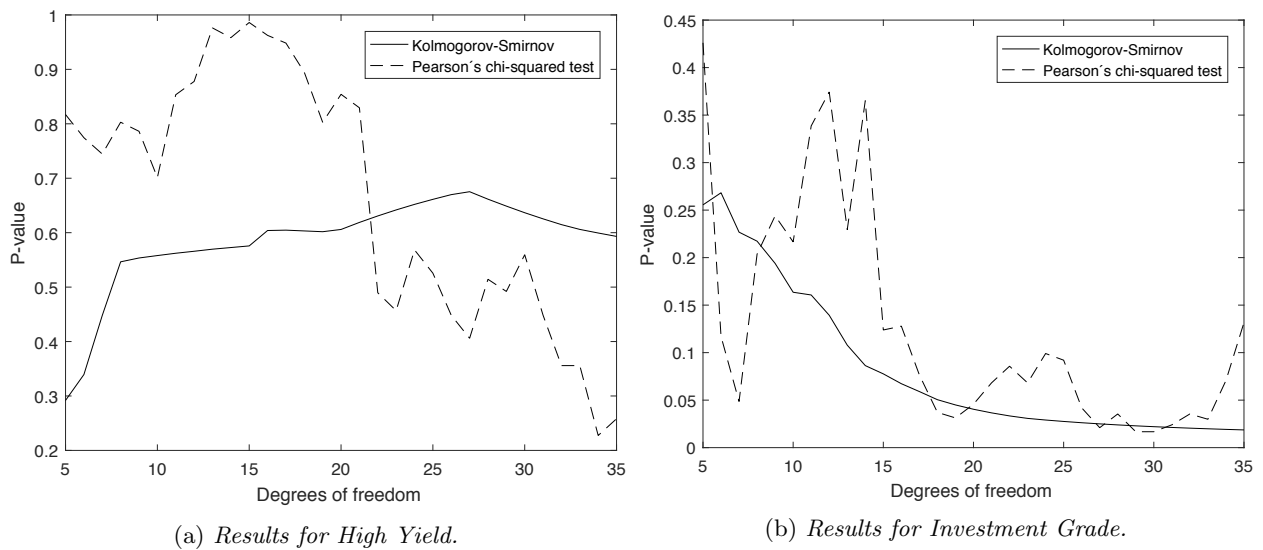


Figure 4.4: The figure shows the P-value of the hypothesis that a t -distribution with ν degrees of freedom have generated the residuals, where ν is displayed on the x-axis. Pearson's chi squared test was used to test the overall fit of the ECDF to the assumed CDF and the Kolmogorov-Smirnov test was used as well to assure that there was no trend in the ECDF's deviation from the CDF. It can be seen that a relative large number of degrees of freedom can not be rejected for neither HY nor IG, but a value of somewhere between 10 and 12 has the highest p-value with respect to both tests for IG and a value between 12 and 18 for HY.



4.3.2 Residual analysis and tests

The residuals are assumed to be uncorrelated and follow t -distributions with 15 and 11 degrees of freedom for High yield and Investment grade respectively. An analysis with respect to these assumptions was performed using the tests described in Section 2.8.1. The results for the High yield and Investment grade segments are summarized below.

Results for High Yield The results of the correlation analysis can be seen in Table 4.6. As can be seen the correlation assumption can not be rejected on a 95 % significance level. Figure A.1 shows the residuals, upper plot, divided by the estimated standard deviation, seen in the lower plot. The clustering previously seen when not using a time-dependent standard deviation seems to have been removed by the model introduced. This is also somewhat confirmed by the positive *McLeod-Li test* which tests for second moment dependence of the residuals. All in all, the residuals look white and non-clustered, which is good.

The assumption of the innovations being distributed as a t -distribution with 15 degrees of freedom can not be rejected at a 95 % significance level. The P -values for *Pearson's chi-squared test* and the *Kolmogorov-Smirnov test* can be seen in Figure 4.4a. Although the entire distribution can not be rejected there is strong evidence that the upper tail of the assumed distribution is insufficient for explaining the distribution of the residuals. This can be seen in Appendix A.1.1 in Figure A.2 and A.3, where the upper tail exceeds the 99 % confidence interval.

Test	Test Quantity	95% confidence interval cut-off
Ljung-Box-Pierce test	25.78	[0,36.42]
McLeod-Li test	10.12	[0,36.42]
Monti-test	24.64	[0,36.42]
Sign-change test	0.44	[0.44,0.56]

Table 4.6: *Residual tests for no autocorrelation of the residuals using Model 1 for High Yield data. As can be seen the residuals pass all the tests at a 95-% significance level.*

These results are good, but as it turns out, the model becomes even better when using the Investment grade spread as input to the High yield spread, which can be seen in Section 4.4.

Results for Investment Grade The results of the correlation analysis can be seen in Table 4.7. As can be seen the correlation assumption can not be rejected at a 95 % significance level for first moment correlation. The negative *McLeod-Li test* suggests that there still is correlation for the second moment of the residuals. These suspicions are further reinforced by the clustering of high volatility areas seen in Figure A.4, which shows the residuals, upper plot, divided by the estimated standard deviation, seen in the lower plot.

The assumption of the innovations being t -distributed with 11 degrees of freedom can not be rejected at a 95 % significance level. The P -values for *Pearson's chi-squared test* and the *Kolmogorov-Smirnov test* can be seen in Figure 4.4b. Although the tail fit is not perfect, see Figure A.5 and A.6, no percentiles can be rejected as generated by the assumed distribution on a 99 % significance level.

Test	Test Quantity	95% confidence interval cut-off
Ljung-Box-Pierce test	25.29	[0,36.42]
McLeod-Li test	48.44	[0,36.42]
Monti-test	26.02	[0,36.42]
Sign-change test	0.47	[0.45,0.55]

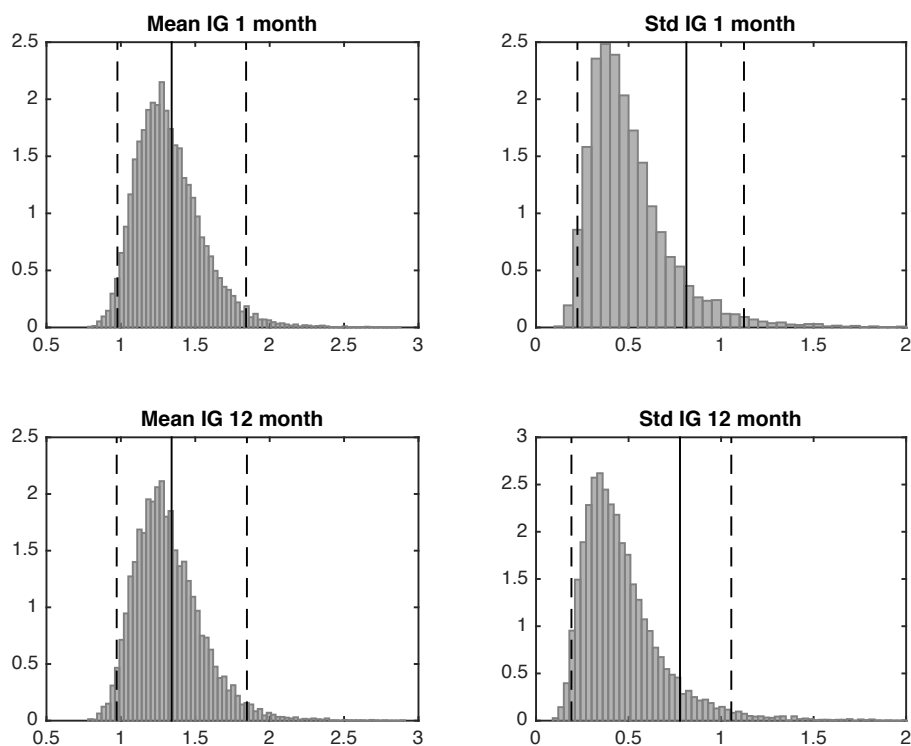
Table 4.7: *Residual tests for no autocorrelation of the residuals using Model 1 for Investment Grade data. As can be seen the residuals pass all the tests except the McLeod-Li test at a 95 % significance level.*

4.3.3 Bootstrap analysis

To further validate the models their simulation performance will be evaluated. The methods used are explained in Section 3.7.3. First off, the first two moments of the credit spread simulations will be analyzed, using both rolling and separate windows. The moments are calculated for 1 month and 12 months windows. Only results for the Investment grade segment will be shown here, and the results for the High yield segment can be seen in Appendix A.2. This is because the model with the IG spread as input is better for HY, and the bootstrap analysis results for this model will be shown in Section 4.4.

The bootstrap analysis for the IG segment can be seen in Figure 4.5. Simulated data seems to replicate the statistics of the historical data reasonably well, and the historical data is inside the 95 % confidence interval spanned by the simulation statistics, for both mean and standard deviation. Using the rolling window approach the results are almost identical.

Figure 4.5: Mean and standard deviation bootstrap analysis using separate windows, for the Investment grade segment using **Model 1**. The dotted lines are 95 % confidence intervals for the simulated data. The solid line indicates the statistic for the historical data. The histogram shows the distribution of the statistics for 10 000 simulated paths.



4.4 Models with input

The correlation between the log transformed and demeaned High yield and Investment grade credit spreads is 0.93 and the correlation between their lagged log-returns is 0.18. This suggests that it might be reasonable to use one as input variable to the other. In this section the performance of such models will be evaluated. The following models are considered:

Process input model - Model 8 Using the models seen in equation 4.10 and 4.11 below it will be tested whether either the High Yield credit spread can significantly explain changes in the Investment Grade

segment and vice versa. The time varying volatility process is initially chosen as **Model 1** but it will also be investigated whether **Model 9** significantly improves the performance.

$$y_t^{HY} = \alpha y_{t-1}^{HY} + \beta y_{t-1}^{IG} + \sigma_t z_t, \quad z_t \sim t(15), \quad (4.10)$$

$$y_t^{IG} = \alpha_1 y_{t-1}^{IG} + \alpha_2 y_{t-2}^{IG} + \beta y_{t-1}^{HY} + \sigma_t z_t, \quad z_t \sim t(11). \quad (4.11)$$

Note that **Model 8** is a process model and not a volatility model, as all other models listed here. **Model 8** can thus be used in combination with the volatility models, if the noise is to be time-dependent.

Model 9 The model described in equation 4.12 is used to test whether the performance of **Model 1** for High yield or Investment grade is improved by adding the volatility predicted for the other segment using **Model 1**.

$$\log(\sigma_{t,a}^2) = \omega + \alpha_1 \log(\sigma_{t-1,a}^2) + \alpha_2 \varepsilon_{t-1} + \alpha_3 \log(\sigma_{t,b}^2) \quad (4.12)$$

4.4.1 Comparison

The usual AR(1) and AR(2) models are nested in **Model 8** and **Model 9**, given that their respective volatility processes are governed by the same model. Since this is the case the likelihood ratio test can therefore be used to evaluate their respective performance. In Table 4.8 it can be seen that no significant improvements of the Investment grade model is achieved by using High yield values as input.

For High yield the log-likelihood improvements are not statistically significant with respect to the likelihood ratio test. However, when using a combination of **Model 8 and 9** for High Yield the α_1 and α_2 parameters in equation 4.12 are insignificantly different from zero. Removing these parameters and performing the analysis again with a combination of **Model 8 and 10** described by equation 4.13 yields a log-likelihood of 266.81. Since the models compared are no longer nested the likelihood ratio test can not be performed. Hence, the higher likelihood is enough to conclude that the combination of **Model 8 and 10** is preferred to **Model 1**.

Model 9

$$\log(\sigma_{t,a}^2) = \omega + \alpha_3 \log(\sigma_{t,b}^2). \quad (4.13)$$

Model	8	9	8 and 9	8 and 10	1
# Parameters(IG/HY)	6/5	6/5	7/6	-/4	5/4
Investment grade					
Log Likelihood	295	294	295	—	294
AIC	-578	-576	-576	—	-578
BIC	-556	-554	-551	—	-560
High yield					
Log Likelihood	264	262	267	267	262
AIC	-518	-514	-522	-526	-516
BIC	-500	-496	-500	-511	-501

Table 4.8: The table shows the log-likelihood as well as AIC and BIC values for the evaluated input models and the previous best model, Model 1. The models with the best, lowest, AIC and BIC values are highlighted with bold text.

4.4.2 Residual analysis

The only model evaluated will be the combination of **Model 8 and 10** for High yield with Investment grade as input since it was the only input model significantly improving the results compared to **Model 1**. The residuals between the one step predictions using a combination of **Model 8 and 10** and the historical values can be seen in Figure A.7 along with the estimated standard deviation.

If the correlation tests seen in Table 4.9 and 4.6 are compared all test-quantities are slightly improved for **Model 8 and 10** compared to **Model 1** except for the sign change test. Overall the results are in concordance with each other suggesting white residuals.

To test for the assumed distribution the *Kolmogorov-Smirnov test* and *Pearson's chi-squared test* were used resulting in P -values of 0.51 and 0.17. For **Model 1** the corresponding P -values were 0.58 and 0.91. The Kolmogorov-Smirnov results were about the same as for **Model 1**, meaning that the maximum distance between the ECDF and the assumed CDF are the same for both models. The P -value from *Pearson's chi-squared test* was much lower meaning that the overall shape of the empirical distribution was worse than for **Model 1**. A P -value of 0.17 still means that the assumed t -distribution with 15 degrees of freedom can not be rejected. The problem with a heavier upper tail seen for all previous models remain, see Figure A.8.

Test	Test Quantity	95% confidence interval cut-off
Ljung-Box-Pierce test	23.41	[0,36.42]
McLeod-Li test	8.65	[0,36.42]
Monti-test	23.08	[0,36.42]
Sign-change test	0.43	[0.44,0.56]


Table 4.9: Residual tests for no autocorrelation of the residuals using **Model 8 and 10** for High Yield with Investment Grade input. As can be seen the residuals pass all the tests except the *Sign-Change test* at a 95 % significance level.

4.4.3 Bootstrap analysis

The simulations using the combination of **Model 8 and 10** was analyzed with the same methods as in the case of no external input to the High yield spread model. The mean and standard deviation of the simulations are compared to their historical value. The results can be seen in Figure 4.6.

Both the mean and standard deviation lies inside the 95 % confidence intervals of the simulated distributions. The standard deviation seems to be slightly higher for the historical data than for the bulk of simulated paths, but not unreasonably so.

Figure 4.6: Statistics for the HY spread with the IG spread as input, using separate windows. The dotted lines are 95 % confidence intervals for the simulated data. The solid line indicates the statistic for the historical data. The histogram shows the distribution of the statistics for 10 000 simulated paths.



4.5 Parameter estimates and stability for final models

In this section the parameter estimates will be presented for the final models. The stability of the parameters with respect to time will also be investigated. In Table 4.10 and 4.11 the parameter estimates are presented along with their corresponding 95 % confidence intervals. In Figure 4.7 and 4.8 the parameters estimated, using a window length of 100 months that is rolled over the data set, can be seen along with the final parameter estimate.

Final models The final model for the High yield credit spread is seen in equation 4.14, and the final model for the Investment grade credit spread is seen in equation 4.15.

$$\begin{aligned}
 y_t^{HY} &= a_1 y_{t-1}^{HY} + a_2 y_{t-1}^{IG} + \sigma_{t,HY} \cdot z_t, & z_t &\sim t(15), \\
 \log(\sigma_{t,HY}^2) &= \omega + \alpha \log(\sigma_{t,IG}^2).
 \end{aligned}
 \tag{4.14}$$

$$\begin{aligned}
 y_t^{IG} &= a_1 y_{t-1}^{IG} + a_2 y_{t-2}^{IG} + \sigma_{t,IG} \cdot z_t, & z_t &\sim t(11), \\
 \log(\sigma_{t,IG}^2) &= \omega + \alpha_1 \log(\sigma_{t-1,IG}^2) + \alpha_2 z_{t-1}.
 \end{aligned}
 \tag{4.15}$$

Parameter	Estimated value	95 % confidence interval
a_1	0.8502	[0.78, 0.92]
a_2	0.1037	[0.05, 0.16]
ω	-1.8427	[-0.23, -3.46]
α	0.5946	[0.29, 0.90]

Table 4.10: *Parameter estimates for the High yield model with corresponding 95 % confidence intervals for the parameters.*

Parameter	Estimated value	95 % confidence interval
a_1	1.1904	[1.09, 1.29]
a_2	-0.2237	[-0.13, -0.32]
ω	-1.4400	[-0.85, -2.03]
α_1	0.7331	[0.62, 0.84]
α_2	0.3446	[0.22, 0.47]

Table 4.11: *Parameter estimates for the Investment grade model with corresponding 95 % confidence intervals for the parameters.*

Figure 4.7: High yield parameters: *The figures show the parameter estimates, blue, for a window of length 100 months rolled over the data set. The black solid line is the parameter estimate using the entire data set. The black dotted lines are approximated confidence intervals for the parameters using the 100 month window.*




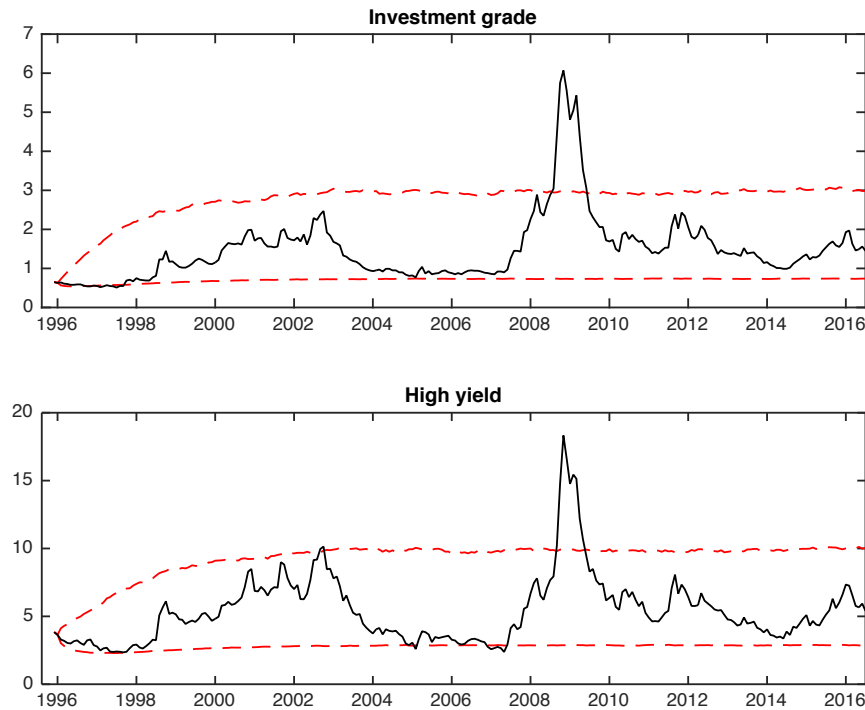
Figure 4.8: Investment grade parameters: The figures show the parameter estimates, blue, for a window of length 100 months rolled over the data set. The black solid line is the parameter estimate using the entire data set. The black dotted lines are approximated confidence intervals for the parameters using the 100 month window.



4.6 Historical spreads with confidence interval

A visual test for the spread simulations is to plot the historical spread together with 95 % confidence intervals spanned by the simulations. This is done in Figure 4.9. The confidence intervals stabilizes on a constant level after about 50 months for the upper confidence boundary, and almost immediately for the lower. There are few values falling short of the lower boundary, and only around the crisis of 2008 does the spread exceed the higher.

Figure 4.9: *Historical spreads for IG and HY segment, together with 95 % confidence intervals constructed by simulated spreads.*




4.7 Index model

The procedure that will be used to recreate a credit index using credit spreads and rates is described in Section 3.5. In this section the performance of the method and the impact of the different reasons for errors will be presented.

First off the impact of daily data as compared to the use of monthly data will be evaluated. Treasury indices were recreated using daily and monthly data for US treasury bonds with 3, 5, 7 and 10 years to maturity. Recreated treasury indices using maturities between 3 and 10 years and daily and monthly data can be seen in Figure 4.10. The difference between the final values of the indices using daily compared to monthly data was between 1.2 and 2.7 %. A reason for the difference might be that the number of trading days used to calculate the rate for daily data was set to 252 when the number of trading days actually differ from year to year. All in all the results suggests that monthly data can be used instead of daily data since the error for more than 20 years was less than three percent.

The impact of using different maturities was examined by recreating historical treasury indices with an average of 6.9 and 5.1 years duration using historical yield curve rates with 5 and 7 years to maturity. As can be seen in Figure 4.11, the results differ quite heavily between the recreated indices and the historical ones, suggesting that the duration of the historical indices might not be representative when deciding what maturities the interest rates should have, corresponding to the spread. The deviation between the recre-

Figure 4.10: *Recreated treasury indices using daily (grey) and monthly (black) data. The rate indices are matched to the duration of the IG and HY spread.*



ated indices and the historical indices increases when the difference between rates with different maturities increases, which is evidence further suggesting problems with duration. To deal with this the rates will be chosen as a combination of the 10 year and the 3 year yield curve rate, as discussed in Section 3.5.2. The rate used will be calculated according to equation 4.16, where the constant, c_i , is estimated to the one which minimizes the squared difference between historical indices and the recreated ones.


$$r_{\text{index}} = r_{3\text{-YCR}} + c_i(r_{10\text{-YCR}} - r_{3\text{-YCR}}), \quad (4.16)$$

where $r_{i\text{-YCR}}$ denotes the yield curve rate with i years to maturity. c_{IG} and c_{HY} were estimated to 0.26 and 0.14 respectively. There is still some difference between the indices, which will in turn affect the creation of the credit indices. However, after testing a few different interest rates, the effect of the rate on the performance of the simulated indices seems to be quite small.

Except for the credit spread and the rate, simulated credit events within the index are needed in order to retrieve index values. To simulate these events four parameters are needed, namely δ describing the historical recovery rate, γ needed to retrieve λ^{P} as $\lambda^{\text{P}} = \gamma\lambda^{\text{Q}}$, ω describing the actual recovery rate for credit events within the index and ' n ' representing the number of companies within the index.

The average recovery rate has been around 38 % for High yield and 40 % for Investment grade, between the years 1982-2010, as seen in Ou et al. (2011). Using this together with the historical credit spreads λ^{Q} was calculated using equation 2.34. Using the historical credit indices and treasury indices the risk premium for High yield and Investment grade was estimated to 0.0180 for HY and 0.0027 for IG. The risk premium along with the calculated λ^{Q} 's were used to calculate γ by using equation 3.14. Finally, by setting the number

Figure 4.11: *Historical treasury indices with an average duration of 5.1 and 6.9 years plotted together with indices recreated from the historical 5 and 7 year yield curve rate.*



of defaults during each time step to the expected number of defaults and recreating the index using the historical rates and spreads, ω was estimated to the value which gave the recreated indices the same final value as historical ones. All parameter estimates are summarized in Table 4.12.

To find the parameter n , representing the number of companies within the index, the obvious approach is to investigate the actual number of companies within the index. But, since defaults within credit indices are most likely not independent the number n will not represent the number of companies within the index but what proportion, $\frac{1}{n}$, of the index is most likely to experience credit events during the same month given that there are credit events during that month. Therefore n will be estimated by simulating a large number of credit indices using historical rates and spreads and comparing the distribution of the simulated indices to that of the historical index. The n which gives the best fit will be used.

	High Yield	Investment Grade
δ	0.38	0.40
γ	0.8124	0.8955
ω	0.70	0.84
n	50	100

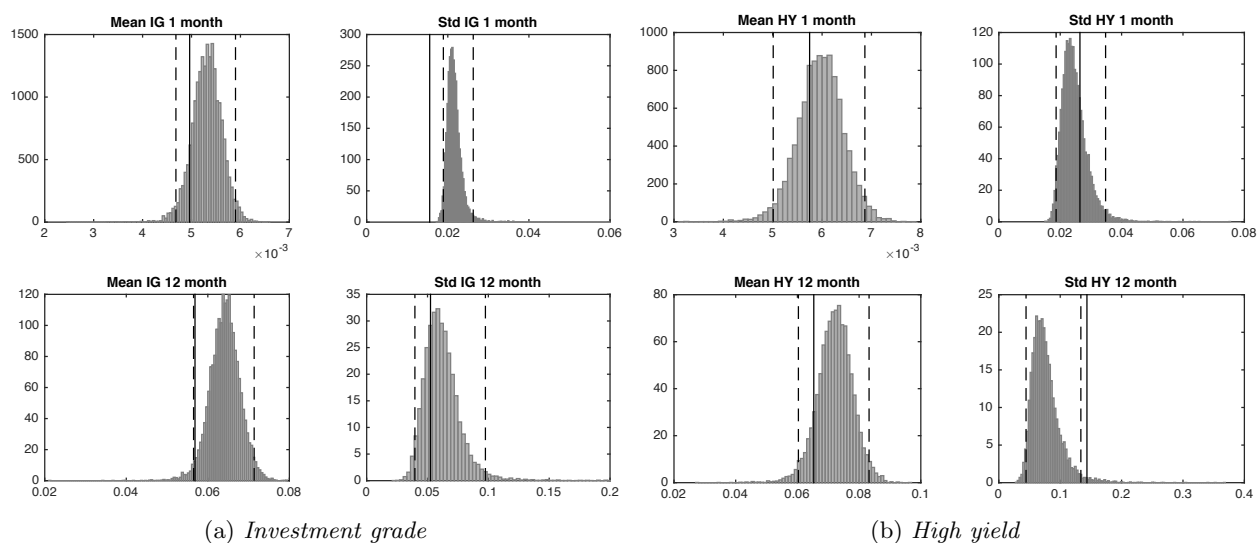
Table 4.12: Summarized parameters used to simulate index-within defaults for HY and IG respectively.

4.7.1 Credit Index validation

To validate the credit index models the same validation tools as in Section 4.3.3 are used. That is, the first two moments of the index log-returns, and the shape of the normalized distribution of the log-returns are evaluated. The lengths of the return periods considered are one and twelve months.

Mean and standard deviation Histograms of the mean and standard deviations of log-returns on indices created using simulated spreads and rates, together with historical mean and standard deviation are shown in Figure 4.12 below.

Figure 4.12: Mean and standard deviation of log-returns on historical and simulated credit indices, using separate windows.

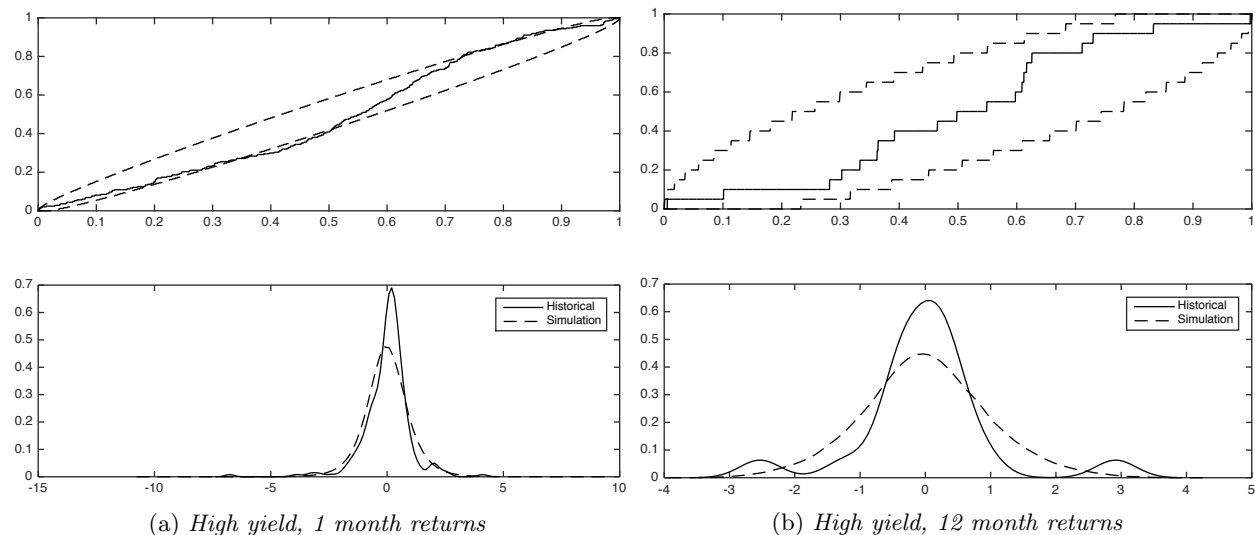


As seen in Figure 4.12a the standard deviation of the one month log-returns on simulated IG indices are higher than for the historical index. This might be due to problems with approximating maturity and face values when discounting the price. For longer returns this problem disappears. For the HY segment, the standard deviation of the 12 month log-returns are a little low. This is not the case for either 11 or 14 month log-returns, nor in the case when using the rolling window method, rather than separate. This could mean that the 12 month returns contain some outlier that makes the standard deviation high. The impact of the outlier will decrease when using rolling windows. For Investment grade, the statistics are not affected considerably using the rolling window approach.

Empirical PDF and Brownian bridge After testing the mean and standard deviation the fit of the normalized distribution is tested with the Brownian bridge test. Also plots of the empirical PDF of the simulated index together with the historical one will be displayed.

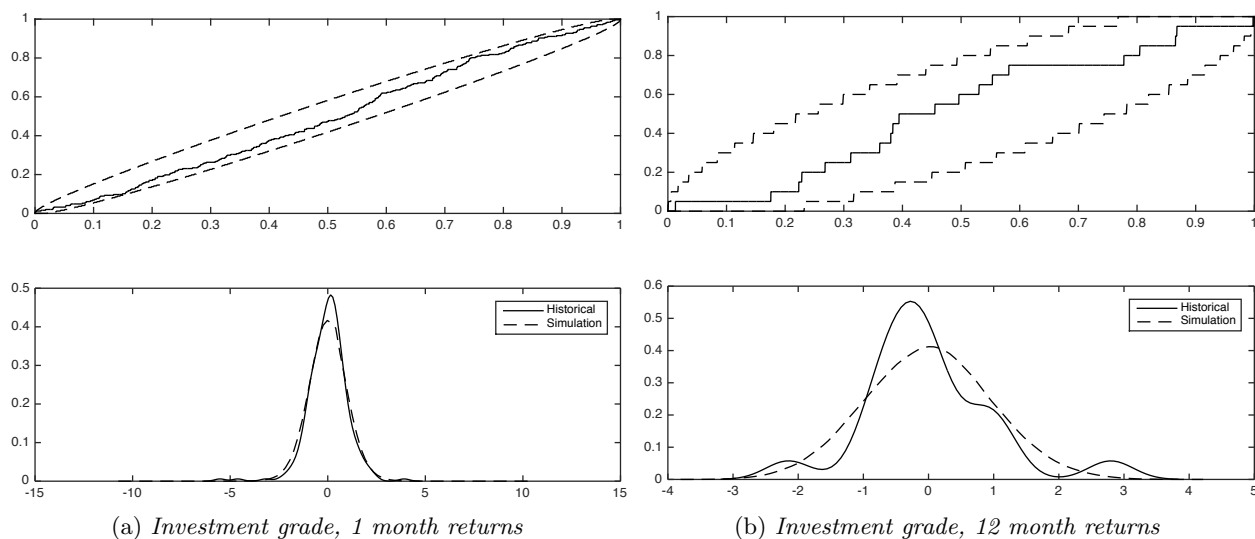
The 1 month returns are slightly off for High yield. The kurtosis is not high enough, which means that the simulated index have tails that are to large. The twelve month returns make a reasonable approximation

Figure 4.13: Brownian bridge test with 95 % confidence interval and density of simulated and historical index data for the High yield segment.



of the historical density. Due to the small sample, the confidence interval becomes rather large and caution should be exercised when drawing conclusions. The p -values from the χ^2 test and the Kolmogorov-Smirnov test can be seen in Table 4.13. The HY segment does not pass either of the tests, and this is probably due to the approximation errors made from the transition from spread to index.

Figure 4.14: Brownian bridge test with 95 % confidence interval and density of simulated and historical index data for the Investment grade segment.



Simulated 1 month returns on the index have an almost identical density as the historical data for Investment grade, and the Brownian bridge is inside the 95 % confidence interval at all times. The twelve month density for the historical data is not very smooth, but that is probably due to the relatively small number of samples, and should not be given to much attention. The p -values for the χ^2 and Kolmogorov-Smirnov tests can be seen in Table 4.13. The Investment grade segment passes both tests.

In summary, the Investment grade segment passes the distribution tests with flying colors, while the High

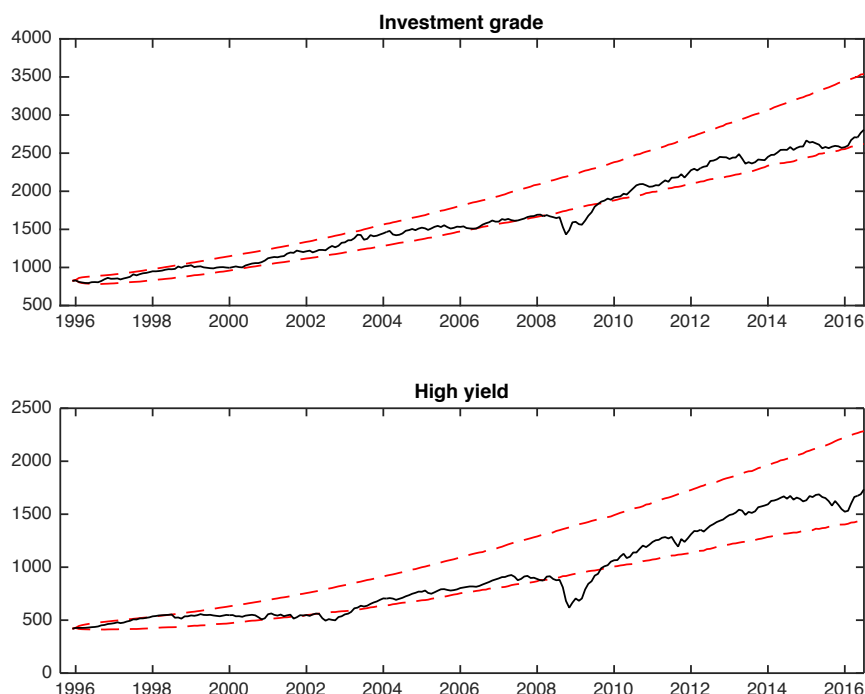
	High Yield	Investment Grade
Kolmogorov-Smirnov	0.0045	0.39
χ^2	0.00013	0.64

Table 4.13: p -values for the χ^2 and Kolmogorov-Smirnov tests. This tests the distributions of the historical and simulated log-returns on credit indices. A p -value above 0.05 means that the simulation distribution passes the test.

yield segment does not perform well at all. In order to make the tails less fat for HY, more degrees of freedom for the t -distribution could have been selected. This could perhaps make the distribution fit historical data better.

Historical indices with confidence intervals To see whether the simulated indices create a reasonable boundary for the historical data, the historical indices together with a 95 % confidence interval is plotted. The confidence intervals are constructed by the simulated indices, by taking the 95 % percentiles for each point in time. The result is seen in Figure 4.15 below.

Figure 4.15: *Historical indices for IG and HY segment, together with 95 % confidence intervals constructed by simulated indices.*



4.8 Dependence between asset classes

The assets considered were split into two groups, the first consisting of the 10 year- and 3 year Treasury rates and the second of the HY- and IG spreads and the S&P500 index. Fitting a t -copula to each group individually yielded 6.4 degrees of freedom for the first group and 8.8 for the second, whereas the degrees of freedom if a t -copula was calibrated to all assets in a single group yielded 32.6 degrees of freedom.

The method for calibrating the grouped t -copula described in Section 3.6.1, based on using the t -copula on each subgroup and method of moment estimates of the correlation, was used. This approach resulted in a

log likelihood of the grouped t -copula being lower than that of the t -copula meaning suboptimal parameter estimates.

Since the estimates of the degrees of freedom for the t -copula deviates that much from the estimates within each group the grouped t -copula will be used although the likelihood of the t -copula was higher. This is because the tail dependence within each group is more important than the tail dependence between groups, due to the correlation between groups being so low. Also, for practical usage purposes a stronger rather than a weaker tail dependence is desired since this will overestimate crashes rather than the opposite.

Bootstrapping a confidence interval for the log likelihood of the calibrated grouped t -copula resulted in a P -value of 0.0080, meaning that the dependence structure probably not can be represented by the calibrated grouped t -copula. However, looking at each subgroup the results are different. For the credit spread and stock index group the P -value was 0.63 and 0.22 for the rate group. This means that it can not be rejected on a 95 % significance level that the grouped t -copula can represent the dependence structure within each group.

4.8.1 Tail dependence

As discussed in Section 3.6, parametric bootstrap will be used to see whether the simulated indices has tail dependence that matches historical data. 10000 paths of each asset class were simulated using the grouped t -copula, resulting in 10000 tail dependence values. The span of these can be seen in Figure 4.16 below, together with the historical tail dependence value.

In almost all cases the historical tail dependence is inside the 95 % confidence interval. The exceptions are upper tail dependence between IG and the rate indices, as well as HY and the stock indices. This is also in the upper tail. Another important observation is that the historical tail dependence is different in the lower and upper tail, that is, it has a skewed tail dependence. The t -copula is symmetric and gives equal tail dependence for the upper and lower tail. To get better results one could use a skewed t -copula, as described in Demarta and McNeil (2005), or another copula with skewness. However, the estimation procedure becomes more involved with a skewed copula. In the figure, the tail dependence is defined as in equations 3.20 and 3.21. The negative tail dependence defined as $P(\mathbf{X} > F_{\mathbf{X}}^{-1}(q) | \mathbf{Y} \leq F_{\mathbf{Y}}^{-1}(q))$ is zero between all indices, both historically and for simulations; except between HY and the rate. Here, the negative tail dependence is about 0.25 historically, but zero for the simulations.

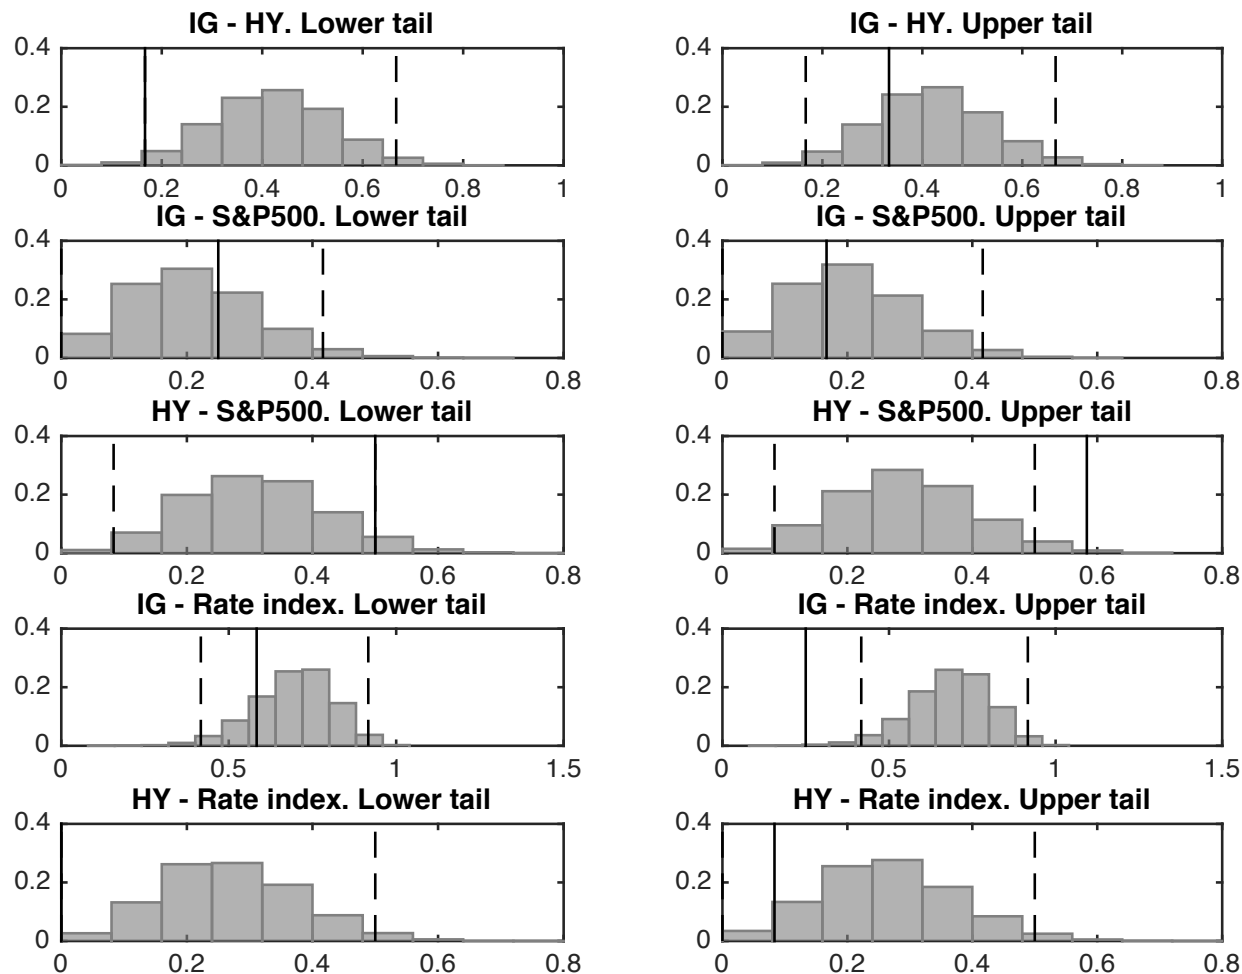


Figure 4.16: Upper and lower tail dependence between log-returns on indices. The dotted lines represent 95 % confidence intervals, and the solid black line represents the historical value of the tail dependence.

4.8.2 Correlation

Here, the correlation between asset classes are tested with the parametric bootstrap method. Both correlation between the spreads themselves, and log-returns for indices are tested. The results for the correlation between spreads, S&P500, and 3 year treasury rate can be seen in Figure 4.17.

The confidence intervals are very broad, but in most cases the mean of the simulation correlations seems to correspond to the historical correlations, which is good.

In the same way as above, the correlation between simulated and historical log-returns of indices can be tested with a parametric bootstrap approach. In Figure 4.18, a histogram of the correlations between 10 000 simulated indices paths are shown, as well as the the 95 % confidence interval spanned by the simulations, together with the historical correlation between asset classes.

Correlations between log-returns on indices does not perform well at all. The only simulation matching historical values are the Investment grade index and S&P500. This is probably due to the approximation errors discussed earlier, when moving from spread to index.

Figure 4.17: Histogram of correlations between simulations of IG and HY spreads, S%P500, and 3 year treasury rate, together with 95 % confidence intervals.

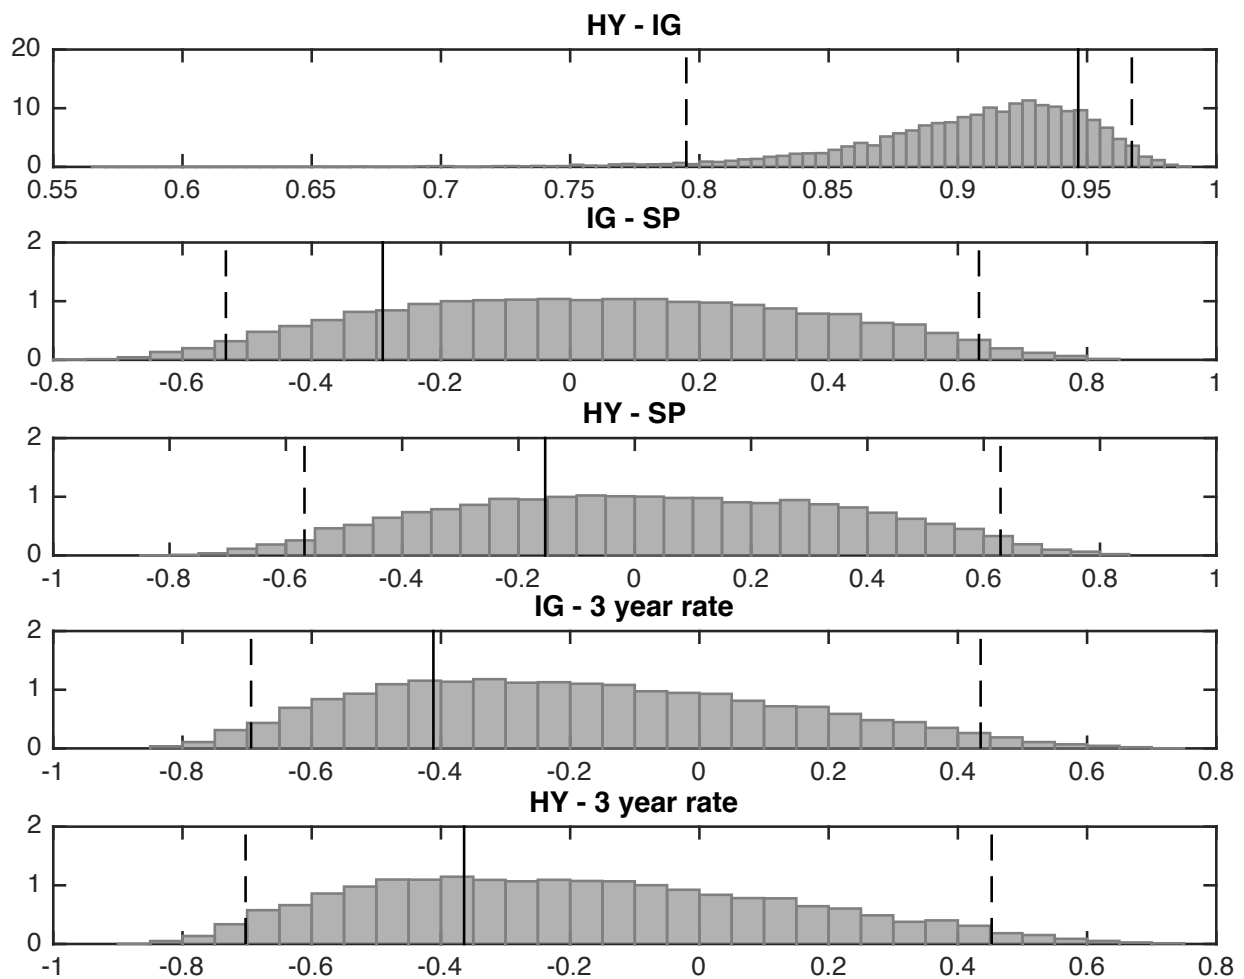
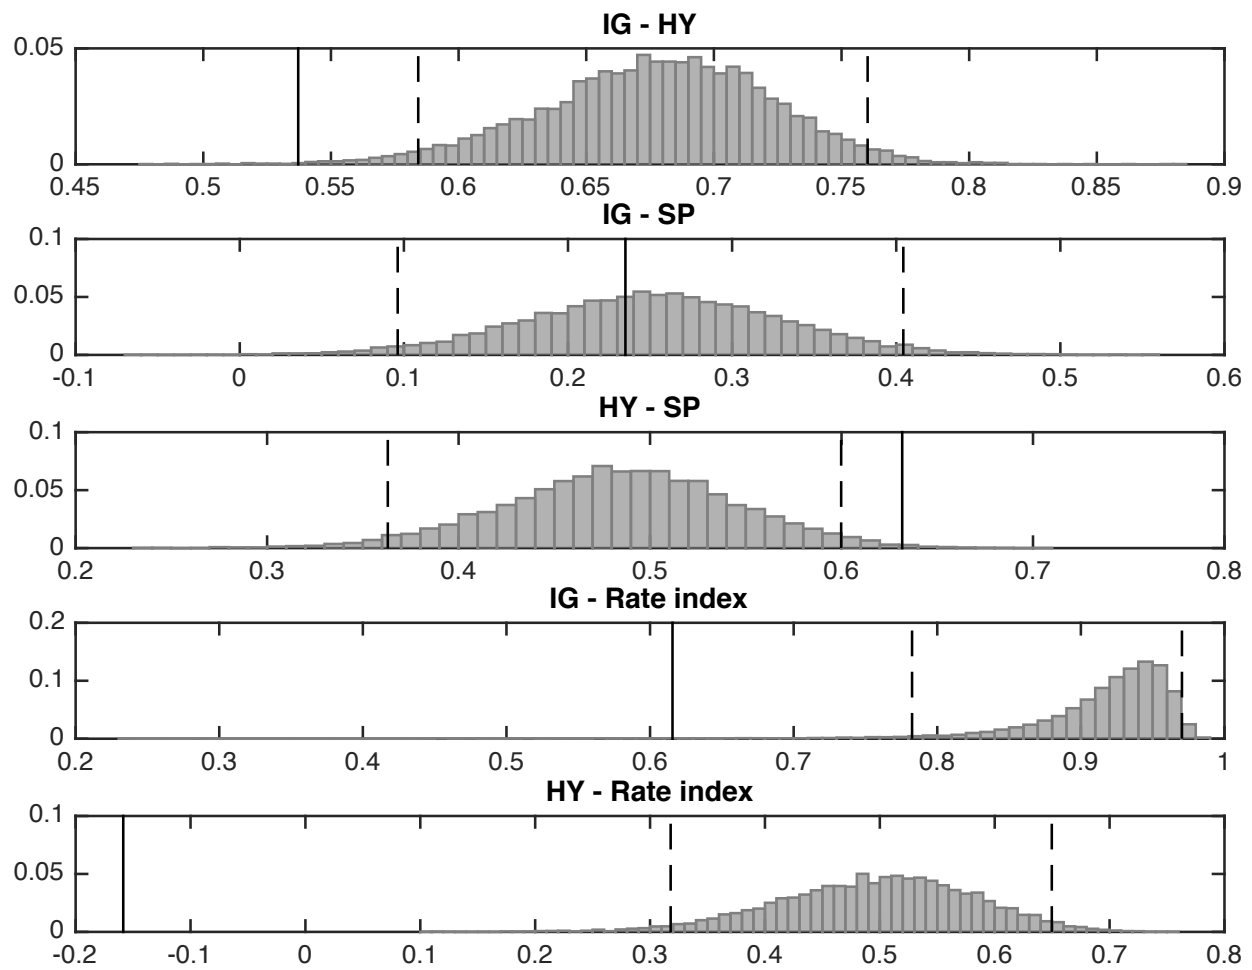


Figure 4.18: Histogram of correlations between simulations of IG and HY indices, S&P500, and a rate index, together with 95 % confidence intervals. The correlations are between log-returns on the indices.



Chapter 5

Discussion

In this chapter the results and choices made throughout the thesis will be presented and discussed. Possible improvements will also be proposed. In the section *Summary* the results will be presented, in the sections *Credit spread model*, *Index model* and *Dependence* the results for each model will be discussed and in the section *Future research* improvements and future research topics will be proposed.

5.1 Objectives summary

In this section the resulting models of the thesis will be presented, as well as a summary and discussion over the results. This is done in order to answer the objectives of the thesis, described in Section 1.4.

Investment grade credit spread The final IG credit spread model is given by the following equation:

$$\begin{aligned} y_t^{IG} &= a_1 y_{t-1}^{IG} + a_2 y_{t-2}^{IG} + \sigma_{t,IG} z_t, & z_t &\sim t(11), \\ \log(\sigma_{t,IG}^2) &= \omega + \alpha_1 \log(\sigma_{t-1}^2) + \alpha_2 z_{t-1}, \end{aligned} \quad (5.1)$$

where y_t^{IG} denotes the log-transformed and demeaned Investment grade credit spread.

High yield credit spread The final HY credit spread model is given by the following equation:

$$\begin{aligned} y_t^{HY} &= a_1 y_{t-1}^{HY} + a_2 y_{t-1}^{IG} + \sigma_{t,HY} z_t, & z_t &\sim t(15), \\ \log(\sigma_{t,HY}^2) &= \omega + \alpha \log(\sigma_{t,IG}^2), \end{aligned} \quad (5.2)$$

where y_t^{HY} denotes the log-transformed and demeaned High yield credit spread.

Transition to index The transition to a credit index, assuming known rates and credit spreads, is described by the following equation:

$$\begin{aligned} \frac{P(t+1, T)}{P(t, T)} &= (1 - D_{t+1}(1 - \omega)) \frac{e^{-(r_{t+1} + y_{t+1})(T-t-1)}}{e^{-(r_t + y_t)(T-t)}}, & nD_{t+1} &\sim Bin(n, q_{t+1}), \\ q_{t+1} &= \left(\frac{1 - \delta}{e^{-y_{t+1}(T-t-1)} - \delta} \right)^{-\frac{\gamma}{12(T-t-1)}}, \end{aligned} \quad (5.3)$$

where δ denotes the historical recovery rate, γ is a constant converting the risk neutral probability of default to the objective probability of default and ω is a constant representing the impact of credit events on the price of a specific bond.

Dependence The dependence structure was represented by a grouped t -copula where the first group consisted of HY and IG credit spreads and the S&P500 index and the second group consisted of the 3 and 10 year US treasury rates.

Model validation The validation results differ between the rating segments and also between the spread models and credit index models. For the spread models the results are good, with white residuals that fit the assumed distributions of the models. The simulations perform well regarding the parametric bootstrap tests versus historical data. For the simulated indices, the results are not as good. The parametric bootstrap tests look reasonably well. An exception is the standard deviation for IG, where the simulations have a higher volatility than the historical data. The distributional tests does not perform well for the HY segment, but much better for the IG segment. The results, and possible reasons for poor results, are discussed further for each of the objectives below.

5.2 Credit spread model

5.2.1 Noise distribution

The driving noise of the credit spread model is t -distributed. The degrees of freedom estimated by maximum likelihood were so low that they made the simulations explode, especially for the Investment Grade segment, which is not a realistic spread behaviour. This led to the approach in Section 4.3.1, where the degrees of freedom were chosen such that the P -value from the Pearson's chi squared test was as large as possible. This approach can be defended as the variance of the maximum likelihood estimate was rather large and that the P -value was big enough not to reject the choice. Another way to avoid exploding simulations could have been to use some other distribution for the noise. Having tried the Normal-inverse-gamma distribution, which is more general than the t -distribution, the results did not improve.

Looking at the confidence intervals of the AR-parameters for the spreads a unit root can not be rejected. Although, a unit root is unrealistic for credit spreads the closeness to instability will increase the uncertainty of the estimated parameters. The uncertainty in the estimation of the parameters may be due to the fact that the process might actually change throughout time, see Section 4.5. The behavior of the process in times of turmoil, as around 2008, is different than it is in states of calm, pointing towards a regime switching model or time varying parameters. Then again, looking at the results from the single-model approach with adjusted degrees of freedom in this thesis, they are quiet good. The noise distributions match well for both High yield and Investment grade. Altogether, the results might improve with some other noise distribution, which incorporates skewness. However, the simplicity of the t -distribution makes it easier to calibrate and interpret.

5.2.2 Regime switching

Some of the parameters in the spread models varies a lot over time. This suggests that even better results could be achieved using some regime switching model, where parameters are different in different regimes, or even that the model has different specifications in different regimes. This is however a lot more difficult to implement, and the model will become less parsimonious. The models proposed in this thesis performs well on spread level, and the pain of making a regime switching model might exceed the gain.

5.3 Credit index

The results when going from credit spread to credit index are not perfect, even though the spread model performs well. This points to some other part of the transition being at fault, rather than the spread simulation. Most tests performed points towards the approximation of the time to maturity of the index as the index duration, as responsible for most of the error. However, there are more possible reasons for error and probably the bad results are a combination of several.

5.3.1 Risk-free rates

In Section 3.5.2 the method used to extract proper risk-free rates is described. This method contains some approximations. The first approximation is that the rate indices are created to fit the duration of the spread indices, by linearly combining rate indices with lower and higher duration than the spread index in each time point. This is a reasonable approximation, but an approximation nonetheless.

The second is that the rates are constructed as a linear combination between the 3 and 10 year US treasury rate. The real rate will diverge from this by some amount, since rate curves rarely are linear. Even so, it is a reasonable approximation since the real rate probably lies somewhere in between these two rates, and the movement of the rate produced will be somewhat realistic. The last step of creating the rates is fitting the constant such that the rate indices are as close to each other as possible. This can be seen as a third approximation, since the rate indices does not fit perfectly.

All in all, the rates might not be as close to the actual rates as need be to make the transformation from spread to index realistic. It could be the rate that introduces the volatility in the 1 month returns on the Investment grade index, seen in Figure 4.12a, making the simulated index volatility to high compared to the historical. However, testing with different rates the results differ very little, which points to the rate approximation not being at fault.

5.3.2 Duration

When moving from spread to index, the modified adjusted duration is used to discount the prices. Using this instead of the real maturity, coupons, and face values could generate discrepancies that causes the statistics of the spread index to differ from the historical. However, since data for individual bonds in the index is not available, this approximation is about the best one can achieve.

5.3.3 Credit events within index

The probability for a credit event to occur within the index is calculated based on an assumption that the credit events are directly proportional to the probability for default. Even if this assumption is correct, the way of calculating the \mathbb{P} -probability for default proposed in Section 2.5.1 is prone to generate errors. The biggest problem can be seen if a large number of paths are simulated - all of them resulting in different final outcomes yielding different risk premium estimations.

Say the \mathbb{P} -probability for default is known for every single company within the index. The problem of finding the correct correlation between the assets within the index remains. For the Investment Grade index there are about 8000 constituents, and for the High Yield index there are about 2000 constituents. Many of them are of course issued by the same company suggesting a really high correlation, whilst some are from different sectors suggesting a lower dependence. Since the task of finding a suitable dependence structure between the assets, and the fact that the assets are unknown, was so big the method of instead simulate the credit events using a binomial distribution was suggested. The correlation of the assets does not impact the expected number of defaults but the probability of more defaults to occur at the same time yielding a higher volatility. To increase the volatility of the binomial approach the number of assets was decreased. This approach will be able to capture the correct volatility, but since choosing a smaller n than the actual number of bonds will result in fewer possible outcomes, the method can not exactly generate all possible number of credit events that could possibly occur during a time step. This was not the only problem with the approach. Finding the number n which best represented the correlation structure was a difficult task due to the noisy observations. The impact on the index due to this was however not very big and the variance of the index due to credit events was small compared to changes in spread and rate.

5.4 Dependence

The chosen grouped t -copula was rejected when bootstrapping a confidence interval for the likelihood. The reason it was still used was the desire to keep tail dependence between assets within the proposed groups, although this imposed tail dependence between assets that empirically had uncorrelated tails. For each group the grouped t -copula was a good fit and the estimated degrees of freedom was low enough for significant tail

dependence, which was desired. The credit index model will be used to simulate future scenarios. For such a purpose it is more important not to generate too stable future scenarios than generating too volatile ones. The fact that the grouped t -copula imposed tail dependence between groups will result in a more volatile simulation of future events.

5.5 Future research

In this section improvements and future research topics will be presented. Each paragraph represents an area of this thesis.

Credit spread model Although the credit spread models yielded satisfactory results both with respect to residual analysis and bootstrapping tests there are a lot of research topics within this field. It would be interesting to use a stochastic volatility model in order to increase the speed of which a model adjusts to periods with different volatility. Residual analysis showed skewness in the residuals which suggests that some other distribution of the innovations than a t -distribution might be a better fit. Another interesting research topic is that of time varying parameters and hidden Markov models, none of which was handled in this thesis.

Index model The results of the index model was not as satisfactory as those of the credit model. Although this was the case the results did not necessarily depend on a bad model, but rather on the data. There is definitely work to be done on the topic of modeling credit events within an index. The standard model, depending on fitting a multivariate distribution function, is dependent on a large number of variables being estimated and that the choice of copula will be able to fit the structure of the noise. Due to the high number of constituents within many credit indices this approach will be too complex to be used in practice. Therefore, the need for further research within this area is big.

Dependence There are an infinite number of copulas. For high dimensional problems it is difficult to find an appropriate one due to different types of dependence between different variables. Therefore there is a chance that no standard copulas like the Frank, Gumbel, Clayton, Gaussian or t -copula will result in a good fit. An interesting research area would be to develop a non parametric, or quasi parametric, approach of simulating dependent noise.

Bibliography

- Arvanitis, A., Gregory, J., and Laurent, J.-P. (1999). Building models for credit spreads. *Journal of Derivatives*, Spring 99 issue.
- Bierens, H. J., Huang, J.-Z., and Kong, W. (2003). An econometric model of credit spreads with rebalancing, arch and jump effects. *EFA 2003 Annual Conference Paper No. 719*.
- Björk, T. (2009). *Arbitrage Theory in Continuous Time*. Oxford University Press, 3 edition.
- Bollerslev, T., Engle, R. F., and Wooldridge, J. M. (1988). A capital asset pricing model with time-varying covariances. *Journal of political Economy*, 96(1):116–131.
- Bolton, P., Freixas, X., and Shapiro, J. (2012). The credit ratings game. *The Journal of Finance*, 67(1):85–111.
- Committee of European Insurance and Occupational Pension Supervisors (2010). Ceiops’ advice for level 2 implementing measures on solvency ii: Scr standard formula article 111(d) correlations.
- Cox, J. C., Ingersoll Jr, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, pages 385–407.
- Damiano, E., Li, H., and Suen, W. (2008). Credible ratings. *Theoretical Economics*.
- Daul, S., De Giorgi, E. G., Lindskog, F., and McNeil, A. (2003). The grouped t-copula with an application to credit risk.
- Demarta, S. and McNeil, A. J. (2005). The t copula and related copulas. *International Statistical Review/Revue Internationale de Statistique*, pages 111–129.
- der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes*. Springer Series in Statistics.
- Engle, R. F. (1988). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 96:893–920.
- Engle, R. F. and Kroner, K. F. (1995). Multivariate simultaneous generalized arch. *Econometric Theory*.
- Engle, R. F. and Sheppard, K. (2001). *Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH*. National Bureau of Economic Research.
- Gabrielsen, A. (2010). *Modelling the Dynamics of Credit Spreads of European Corporate Bond Indices*. PhD thesis, City University Cass Business School.
- Hill, N., Auquier, R., Bauer, G., Foley, M., LaMonte, M., and Drevon, F. (2016). Rating methodology. Technical report, Moody’s Investors Service.
- Jacobs, K. and Li, X. (2008). Modeling the dynamics of credit spreads with stochastic volatility. *Management Science*, vol. 54, no. 6, 2008, pp. 1176–1188.
- Jakobsson, A. (2015). *An Introduction to Time Series Modeling*. Studentlitteratur AB.

- Krishnan, C., Ritchken, P., and Thomson, J. (2007). Predicting credit spreads. *Journal of Financial Intermediation*.
- Lando, D. (2004). *Credit Risk Models: Theory and Applications*. Princeton University Press.
- Levin, C. and Coburn, T. (2011). Wall street and the financial crisis: Anatomy of a financial collapse. Technical report, United States Senate, PERMANENT SUBCOMMITTEE ON INVESTIGATIONS, Committee on Homeland Security and Governmental Affairs.
- Lindström, E., Madsen, H., and Nielsen, J. N. (2015). *Statistics for Finance*. Chapman and Hall.
- Manzoni, K. (2002). Modeling credit spreads: An application to the sterling eurobond market. *International Review of Financial Analysis* 11.2:183-218.
- McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative Risk Management*. Princeton University Press.
- Merton, R. C. (1974). On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* 29:449-470.
- Nelsen, R. B. (2006). *An Introduction to Copulas*. Springer Series in Statistics, second edition.
- Ou, S., Chiu, D., and Metz, A. (2011). Corporate default and recovery rates, 1920-2010.
- Salmon, F. (2009). Recipe for disaster: The formula that killed wall street. *WIRED*.
- Wang, D., Rachev, S. T., and Fabozzi, F. J. (2009). Pricing tranches of a cdo and a cds index: Recent advances and future research. *Risk Assessment. Physica-Verlag HD, 2009. 263-286*.
- Wilks, S. S. (1938). The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9:60-62.

Appendix A

Plots

A.1 Residual and distributional analysis plots

A.1.1 Model 1

Figure A.1: **Residual analysis HY:** *The upper graph displays the normalized residuals between **Model 1** and the historical data. The lower figure shows the estimated standard deviation of the noise. There still seems to be some clustering in the residuals but overall they look white.*

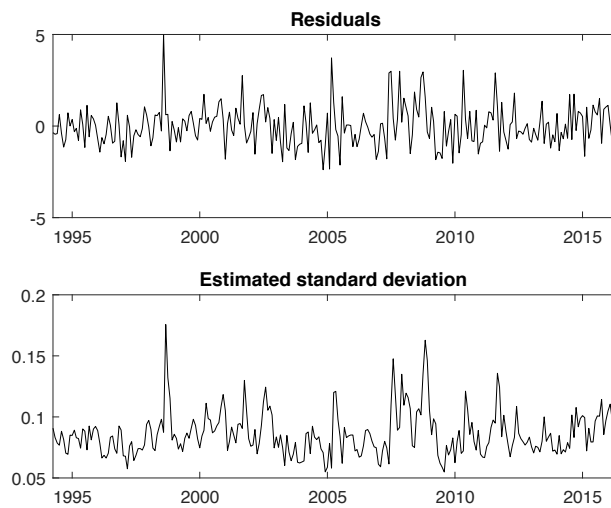


Figure A.2: **Distributional analysis HY:** The figure shows the transformed ECDF, which should lie on a straight line between 0 and 1. The surrounding ellipsoid is the corresponding 99 % confidence interval for the ECDF's percentiles. The transformed ECDF is inside the confidence interval everywhere except for the upper tail.




Figure A.3: **Distributional analysis HY:** Empirical density plot along with a density plot for a t -distribution with 15 degrees of freedom. As can be seen the densities corresponds well overall, but there seems to be some big empirical values which can not be explained by the t -distribution.




Figure A.4: **Residual analysis IG:** The upper graph displays the normalized residuals between one step predictions using *Model 1* and the historical data. The lower figure shows the estimated standard deviation of the noise. There still seems to be some clustering of high volatility areas in the residuals.

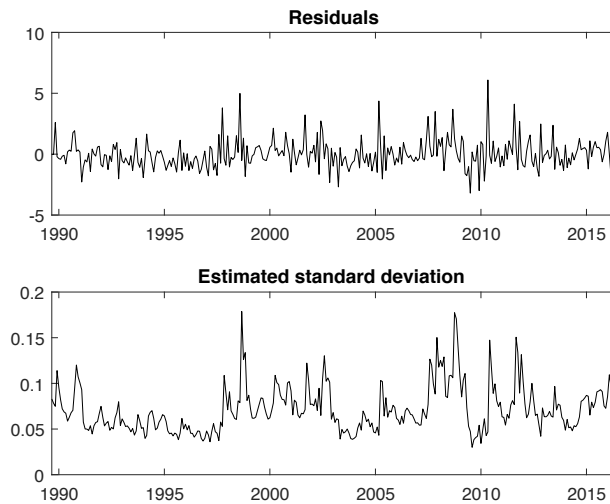


Figure A.5: **Distributional analysis IG:** The figure shows the transformed ECDF, which should lie on a straight line between 0 and 1. The surrounding ellipsoid is the corresponding 99-% confidence interval for the ECDF's percentiles. There are no exceedances of the confidence intervals.





Figure A.6: **Distributional analysis IG:** *Empirical density plot along with a density plot for a t -distribution with 11-degrees of freedom. As can be seen the densities corresponds well overall, but there seems to be some big empirical values which can not be explained by the t -distribution and to compensate for this the chosen hypothetical distribution is less peaked than the empirical distribution.*



A.1.2 Input model: Model 7 & 9

The only segment tested with this model is the High yield segment, since it uses Investment grade as input. Investment grade, however, was not governed by HY.

Figure A.7: **Residual analysis HY, using Model 7 and 9:** *The upper graph displays the normalized residuals between one step predictions using **Model 7 and 9** and the historical data. The lower figure shows the estimated standard deviation of the noise. No obvious clustering of high volatility areas can be seen.*

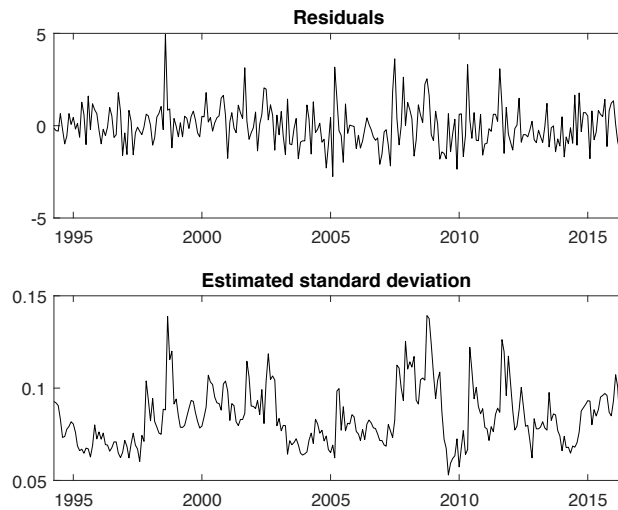


Figure A.8: **Distributional analysis HY, using Model 7 and 9:** The figure shows the transformed ECDF, which should lie on a straight line between 0 and 1. The surrounding ellipsoid is the corresponding 99-% confidence interval for the ECDF's percentiles. The upper tail of the ECDF is outside the confidence interval for the assumed distribution.



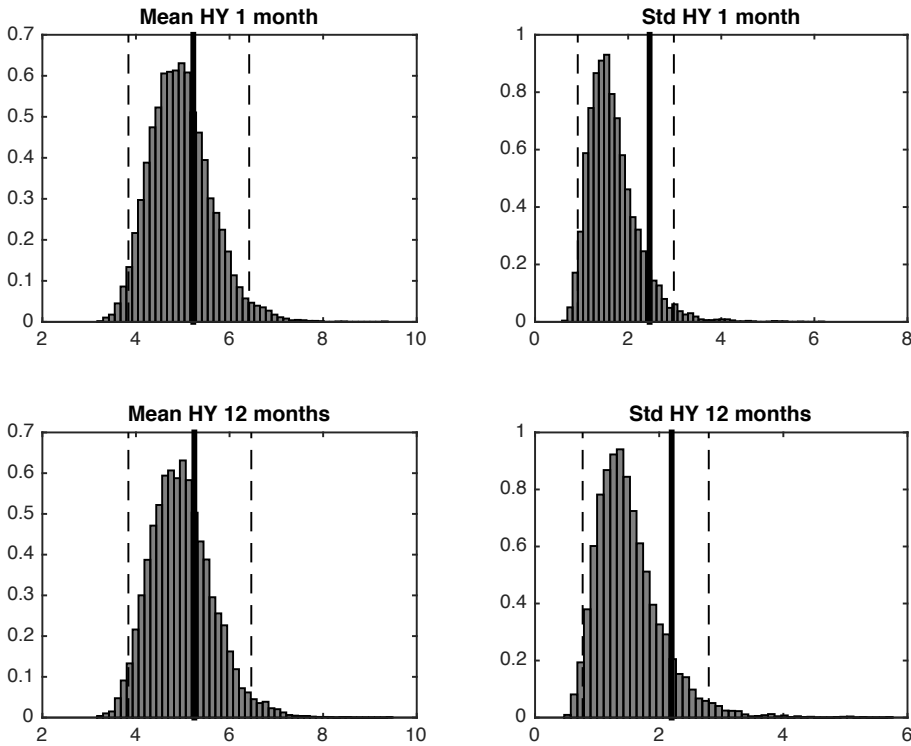


Figure A.9: **Distributional analysis HY, using Model 7 and 9:** Empirical density plot along with a density plot for a t -distribution with 15-degrees of freedom. As can be seen the densities corresponds well overall, but there seems to be some big empirical values in the upper tail.



A.2 Bootstrap analysis plots

Figure A.10: Mean and standard deviation bootstrap analysis using rolling windows, for High yield segment using **Model 1**. The dotted lines are 95 % confidence intervals for the simulated data. The solid line indicates the statistic for the historical data. The histogram shows the distribution of the statistics for 10 000 simulated paths. The separate window approach yields similar results for the 12 month windows.



Appendix B

Grouped t -copula

B.1 Derivation of PDF

$$\begin{aligned}
F(t_1, \dots, t_d) &= P(T_1 \leq t_1, \dots, T_d \leq t_d) \\
&= P\left(Z_1 \sqrt{\frac{\nu_1}{G_1^{-1}(U)}} \leq t_1, \dots, Z_d \sqrt{\frac{\nu_d}{G_d^{-1}(U)}} \leq t_d\right) \\
&= \int_0^1 P\left(Z_1 \leq t_1 \sqrt{\frac{G_1^{-1}(u)}{\nu_1}}, \dots, Z_d \leq t_d \sqrt{\frac{G_d^{-1}(u)}{\nu_d}} \mid U = u\right) f_U(u) du \\
&= \int_0^1 F_{Z, \Sigma}\left(t_1 \sqrt{\frac{G_1^{-1}(u)}{\nu_1}}, \dots, t_d \sqrt{\frac{G_d^{-1}(u)}{\nu_d}} \mid U = u\right) f_U(u) du
\end{aligned}$$

The PDF is then found by deriving this expression with respect to t_1, \dots, t_d yielding:

$$f(t_1, \dots, t_d) = \int_0^1 f_{Z, \Sigma}\left(t_1 \sqrt{\frac{G_1^{-1}(u)}{\nu_1}}, \dots, t_d \sqrt{\frac{G_d^{-1}(u)}{\nu_d}} \mid U = u\right) f_U(u) \prod_{i=1}^d \sqrt{\frac{G_i^{-1}(u)}{\nu_i}} du. \quad (\text{B.1})$$

Further the relationship between the pdf of the multivariate distribution function and the copula is found by the following relation.

$$\begin{aligned}
H(x_1, \dots, x_d) &= C(F_1(x_1), \dots, F_d(x_d)) \\
\Rightarrow h(x_1, \dots, x_d) &= c(F_1(x_1), \dots, F_d(x_d)) \prod_{i=1}^d f_i(x_i) \\
\Rightarrow c(u_1, \dots, u_d) &= h(F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)) \frac{1}{\prod_{i=1}^d f_i(F_1^{-1}(u_1))}
\end{aligned} \quad (\text{B.2})$$

Finally by using equation B.1 and B.2 it is found that the PDF of the grouped t -copula can be written as:

$$c(u_1, \dots, u_d) = \int_0^1 f_{Z, \Sigma}\left(F_1^{-1}(u_1) \sqrt{\frac{G_1^{-1}(u)}{\nu_1}}, \dots, F_d^{-1}(u_d) \sqrt{\frac{G_d^{-1}(u)}{\nu_d}} \mid U = u\right) f_U(u) \prod_{i=1}^d \sqrt{\frac{G_i^{-1}(u)}{\nu_i}} \frac{1}{f_i(F_1^{-1}(u_1))} du$$

Appendix C

Risk premium

C.1 Risk premium estimation

The drift of a risky asset should according to the capital asset pricing model be that of the risk free asset plus some compensation due to the risk. This compensation is known as the risk premium. The expected value of the price of any asset can be written in the following way:

$$E [P_t^{Credit}] = P_0^{Credit} \exp \left(\int_0^t r_f(s) + r_p(s) ds \right) \quad (C.1)$$

, where $r_f(s)$ denotes the risk free rate at time s and $r_p(s)$ denotes the risk premium at time s . Note that the risk premium for a risk free asset is zero. If $E [P_t^{Credit}]$ is estimated as:

$$\hat{E} [P_t^{Credit}] = \frac{1}{t} \frac{P_t^{Credit}}{P_0^{Credit}} \quad (C.2)$$

, then, assuming a constant risk premium, it is possible to construct an estimate the risk premium as:

$$\frac{1}{t} \log \left(\hat{E} \left[\frac{P_t^{Credit} P_0^{Treasury}}{P_0^{Credit} P_t^{Treasury}} \right] \right) = \frac{1}{t} \log (e^{r_p t}) = \hat{r}_p \quad (C.3)$$

An alternative way of estimating the risk premium is to assume that it is a proportion of the risk neutral hazard rate such that $r_p(s) = \alpha \lambda^Q(s)$. Then using equation C.1 and C.3 is is possible to estimate α as:

$$t \hat{r}_p = \hat{\alpha} \int_0^t \lambda^Q(s) ds \quad \Rightarrow \quad \hat{\alpha} = \frac{t \hat{r}_p}{\int_0^t \lambda^Q(s) ds} \quad (C.4)$$

For discretely estimated risk neutral hazard rates the integral in C.4 is substituted with a sum.

Appendix D

Rolling window

D.1 Rolling windows variance of estimates

Given a stationary time series $X = X_t : t = 1, 2, \dots, n$ the task is to create an estimate of the mean value as well as an estimate of the variance for a sum of consecutive values. The sum of the consecutive values is denoted $Y_j(k)$ and is defined as $Y_j(k) = \sum_{i=1}^k X_{j-1+i}$. Two different estimation methods will be considered and their corresponding variances will be evaluated. The different estimates will be denoted $\mu_{rolling}(k)$, $\mu_{usual}(k)$ for the mean values and $\sigma_{rolling}(k)$, $\sigma_{usual}(k)$ for the variances, where the k denotes the number of consecutive observations considered. It is assumed that the length of the time series, n , is divisible by the number k .

D.1.1 Mean value estimation

The mean value estimates are defined in the following way:

$$\begin{aligned}\mu_{usual} &= \frac{1}{n/k} \sum_{i=0}^{n/k-1} Y_{ik+1}(k) = \frac{k}{n} \sum_{i=0}^{n/k-1} \sum_{j=1}^k X_{ik+j} = \frac{k}{n} \sum_{i=1}^n X_i, \\ \mu_{rolling} &= \frac{1}{n-k+1} \sum_{i=1}^{n-k+1} Y_i(k) = \frac{1}{n-k+1} \sum_{i=1}^{n-k+1} \sum_{j=0}^{k-1} X_{i+j} \\ &= \frac{1}{n-k+1} \left(k \sum_{i=k}^{n-k+1} X_i + \sum_{i=1}^{k-1} iX_i + (k-i)X_{n-k+1+i} \right).\end{aligned}$$

The variance of the estimates are given by the following equations:

$$\begin{aligned}V[\mu_{usual}] &= V \left[\frac{k}{n} \sum_{i=1}^n X_i \right] = \frac{k^2}{n^2} (nV[X] + 2 \sum_{i=1}^{n-1} (n-i)C[X_{t+i}, X_t]), \\ V[\mu_{rolling}] &= V \left[\frac{1}{n-k+1} \left(k \sum_{i=k}^{n-k+1} X_i + \sum_{i=1}^{k-1} iX_i + (k-i)X_{n-k+1+i} \right) \right] \\ &= V[X] \left(\frac{(n-2l+2)k^2}{(n-k+1)^2} + 2 \sum_{i=1}^{k-1} i^2 \right) + \text{Sum of covariances},\end{aligned}$$

where it is assumed that the first and last $k-1$ variables are independent.

For an independent time series this means that the variances can be calculated as:

$$V[\mu_{usual}] = \frac{k^2}{n} V[X],$$

$$V[\mu_{rolling}] = \frac{(n - 2k + 2)k^2}{(n - k + 1)^2} V[X] + 2 \sum_{i=1}^{k-1} i^2 V[X].$$

Given these equations it is possible to calculate, for a given number k , when it is better to use rolling windows than not using them. Performing these calculations yields the results seen in equation D.1. It turns out that for all k 's the rolling window estimate will have a higher variance. Remember that k divides n and that n must be greater or equal to $2k$. But these calculations all assume that the length of the time series is divisible by the window length. The resulting mean value variances for different window lengths using overlapping- and non overlapping windows for a time series with 248 observations can be seen in Figure D.1. From this figure it can be seen that with regards to the variance it can sometimes be advantageous to use rolling windows for window lengths that does not divide the length of the time series.

$$V[\mu_{usual}] \leq V[\mu_{rolling}] \Rightarrow \frac{3k(k-1)^2}{2k^2 - 3k + 1} \leq n. \quad (\text{D.1})$$

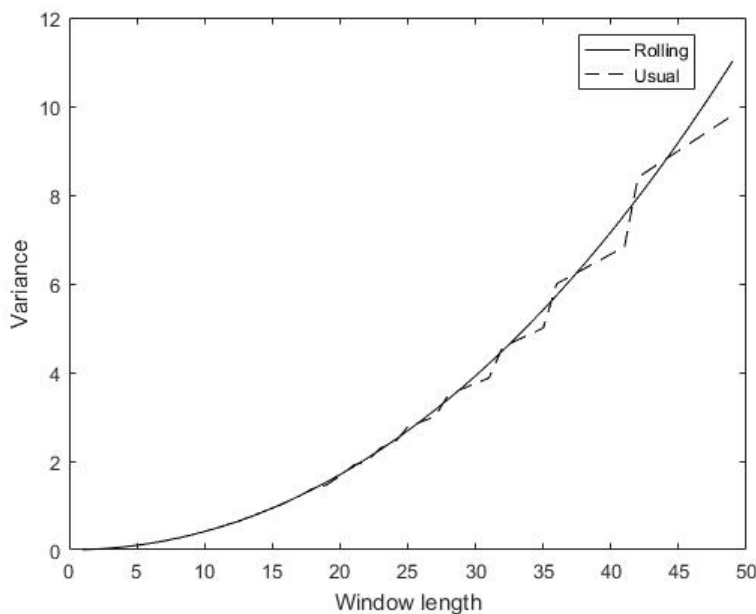


Figure D.1: The figure depicts the variance as a function of the window length using overlapping and non overlapping windows for a time series with 248 observations. It is assumed that $V[X]=1$.

D.1.2 Standard deviation estimation

As can be seen the expression for the variance of the mean value using a rolling window is rather complicated. That of the standard deviation will be even more so and therefore it will not be calculated.