Emma Skarnes

Investigating the Effect of Norwegian Policy Rate Changes on Interest Rate Risk

Master's thesis in Applied Physics and Mathematics
Supervisor: Jacob Kooter Laading
June 2022

Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Norwegian University of
Science and Technology

Emma Skarnes

Investigating the Effect of Norwegian Policy Rate Changes on Interest Rate Risk

Master's thesis in Applied Physics and Mathematics
Supervisor: Jacob Kooter Laading
June 2022

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Norwegian University of
Science and Technology

Preface

This thesis is the final assignment of my Master of Science degree in Applied Physics and Mathematics at the Norwegian University of Science and Technology. In the final three years of my study program, I have specialised in Industrial Mathematics with a heavy focus on statistics courses. When I was first introduced to quantitative finance, I found it very interesting how mathematics and statistics were applied to the subject; later, I have found my statistical and mathematical background to be very helpful for learning and understanding the mechanisms behind financial modelling and thus also of finance itself.

The thesis was carried out at the Department of Mathematical Sciences during the spring of 2022. I would like to thank my supervisor Jacob Kooter Laading for his valuable guidance, suggestions and feedback. Finally, I would like to thank my family for their support, and my friends for five wonderful years of studies.

Emma Skarnes
Trondheim, June 2022

Abstract

The main objective of this thesis is to investigate the risk of interest rate derivatives. We test the effect on risk from changes in the Norwegian policy rate using INLA, integrated nested Laplace approximation. INLA is an established method within statistics but has not been used much in quantitative finance. Studying and evaluating the properties of this framework is a central part of the thesis. Two INLA models are fitted to the entire period from January 2005 to May 2022, and the best performing one is fitted to four subperiods with lengths of approximately two years each. HJM models are fitted to the same four subperiods, giving a qualitative and quantitative comparison of INLA and a more established method for measuring financial risk.

The INLA models give sensible and precise results, with risk measures for every day of the data set. The high persistence of the models indicates that the global risk is incorporated in the local risk measurements, which is a desired property of our time series. Our model does not indicate that the policy rate changes influence the risk measures, as the information from the changes seems to be compromised in the derivative prices. We conclude that INLA is an appropriate and effective framework for modelling financial time series.

Sammendrag

Denne masteroppgavens hovedmål er å utforske risiko tilknyttet rentederivater. Vi tester effekten endringer i den norske styringsrenten har på risikoen ved å bruke INLA (Integrated Nested Laplace Approximation). INLA har blitt en etablert metode innen statistikk, men er lite brukt innen kvantitativ finans. Utforsking og evaluering av rammeverkets egenskaper utgjør en sentral del av oppgaven. To INLA-modeller er benyttet på perioden fra januar 2005 til mai 2022, og den beste modellen av disse er benyttet på fire underperioder med lengder på omtrent to år hver. HJM-modeller er benyttet på de samme fire underperiodene for å gi en kvalitativ og kvantitativ sammenligning av INLA og en mer brukt metode for å måle finansiell risiko.

INLA-modellene gir fornuftige og presise resultater, med risikomål for hver dag i datasettet. Den høye persistensen av modellene indikerer at den globale risikoen er bakt inn i de lokale risikomålene, som er en ønsket egenskap av tidsrekken vi modellerer. Modellen vår indikerer at endringene i styringsrenten ikke påvirker risikomålene, som kan være en konsekvens av at informasjonen fra endringene allerede er inkorporert i derivatene. Vi konkluderer at INLA er et nyttig og effektivt rammeverk for modellering av finansielle tidsrekker.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Mathematical models play an important role in a wide variety of subjects because of their ability to precisely describe the behaviour of different systems. The field of economics and finance is one of the subjects where mathematics come in handy - nowadays investors should have some basic mathematical knowledge to keep up with their colleagues and competitors. In fact, multiple mathematical disciplines are used in finance, with some of the most used ones being statistics, numerical methods and differential equations. These disciplines are often combined and applied to a dynamical system aiming to describe some financial quantity.

Applied mathematics associated with finance is better known under the term quantitative finance. The roots of the field go back to 1900, when Louis Bachelier published his PhD thesis *Théorie de la spéculation* on modelling stock options using stochastic processes [1]. 73 years later, Fischer Black and Myron Scholes published the well-known *The Pricing of Options and Corporate Liabilities* [2], which is now a cornerstone within quantitative finance.

A person working with quantitative finance typically wishes to do one of the following: price derivatives or measure and manage risk in some portfolio. In derivative pricing the most important tools are partial differential equations combined with stochastic calculus, as well as different simulation methods. On the other hand, risk modelling relies heavily on statistics. With new technology, quantitative methods such as Markov chain Monte Carlo (MCMC) and bootstrapping have replaced the earlier qualitative methods, yielding much faster and easier analysis. Risk management is highly relevant for both commercial and savings banks, as well as for insurance companies. For that reason, it is internationally regulated through the Basel frameworks of The Bank of International Settlements.

The outbreak of the Covid-19 pandemic caused abrupt declines in worldwide markets in March 2020. Norges Bank, the central bank of Norway, reduced its policy rate to 0% for the first time in history to try to absorb the worst economic shock as a response to the outbreak. Although the halt in the economy negatively impacted both people and businesses, it recovered relatively fast because of the expansionary monetary policy applied by many central banks all around the world. A complex situation with increasing prices, supply chain issues and lately also Russia's invasion and the following war in Ukraine has resulted in volatile and uncertain market conditions in the Spring of 2022. Alarming economic events have happened, and many are asking if we will soon be facing a new recession.

Because of the late situation, risk measurement and management are highly relevant today. Knowing how to identify and quantify risk is important for all investors to prevent extreme and fatal losses. Thus, this thesis will be focusing on financial risk and how to calculate two loss-related risk measures from financial time series. To do so, we will study a relatively new method within quantitative finance and see how this method performs compared to more used ones.

Although MCMC methods are much faster than the old, non-computational methods, the algorithm is relatively slow. INLA, integrated nested Laplace approximations, was introduced in 2009, as a computationally efficient alternative to MCMC. As opposed to the simulation-based methods of MCMC, INLA performs approximate Bayesian inference by making numerical approximations and solving integrals [3]. Since it was released in 2009, INLA has grown in popularity among statisticians. Despite the fast algorithm and very similar performance to corresponding MCMC methods [4], it is not yet established within quantitative finance. Only a few articles are published, most of them focusing on the results of stochastic volatility models.

The goal of this thesis is to investigate the risk of interest rate derivatives, and to study how changes in the policy rate influences this risk. For both purposes, we use the INLA framework to model the stochastic volatility of a bond portfolio. This volatility is then used to calculate the risk measures value at risk and expected shortfall of the models, and model selection criteria are used to decide which model is the most optimal - the one with or the one without the policy rate variable. We also compare INLA to HJM, a well-known and used method to measure financial risk.

A possible drawback of a stochastic volatility model in INLA is the distribution. This model must be either Gaussian, Student-t or normal inverse gamma distributed, neither of which behave like the typical financial return series. These returns series are more likely to have skewed distributions, e.g., the lognormal distribution. However, we know that established methods such as the Heath, Jarrow & Morton framework also implement Gaussian distributions for their increments in the return series. Comparing INLA to such a method could result in proposing a new, more efficient way to perform risk management.

The risk management methods used today yield the desired results but are computationally demanding. To run the necessary computations, servers must continuously work on the problems. This requires power and capacity. Thus, reducing the required capacity by using a more efficient method will also reduce the operating costs. If INLA proves to be a sufficient alternative to the existing risk measurement methods, banks and other investors could reduce their time and resources spent on risk management.

This thesis is a continuation of the author's project thesis, and hence parts of the theory presented here coincides with that of the project thesis. In particular, the HJM and HJM related topics are built upon the project thesis work. The new element of this master thesis is INLA, which will be our main focus.

In Chapter 2, we present basic financial concepts as well as some relevant theory for this thesis. The basic concepts include time valuation, the efficient market hypothesis and three models related to the latter. The next part, Section 2.3, introduces and discusses monetary policy thoroughly. Topics such as the policy rate, recessions and today's monetary situation are reviewed. This part focuses on Norway, but also discusses the international situation. The chapter finishes by Section 2.4 about interest rates and financial instruments.

Chapter 3 presents relevant background theory from mathematics and the more specific models used in this thesis. Understanding the fundamental theory of Brownian motion, Itô's lemma and Bayesian hierarchical models, the reader will have a good basis for understanding the more specific theory described later. INLA, the main method of this thesis, is presented in Section 3.5, before we go through the HJM framework and model selection criteria.

Financial risk, risk management and regulations are the topics of Chapter 4. The chapter introduces five risk classes and presents the two risk measures we later use in our analysis; value at risk and expected shortfall. The Basel II and Basel III frameworks of the Bank of International Settlements are brought up in the final section of this chapter.

The data, as well as some preliminary analysis, is presented in Chapter 5. We present the original and transformed interest rate data, descriptive statistics and introduce the four subperiods we later return to in the analysis. The specific implementation of the INLA framework relevant for this thesis is described in Chapter 6.

Chapter 7, Results and Discussion, presents and discusses the analysis of the thesis. The section is opened by some fundamental historically measured risk as an indication of how we expect the results to behave, before we apply the model from Chapter 6 to our data. Furthermore, we compare the methods and strategies of INLA and HJM in Section 7.2. Finally, the thesis is concluded in Chapter 8. Here we also suggest some future work within the topic of this thesis.

Chapter 2

Finance

The concept of finance is built up by elements from science as well as from social sciences. Unlike the natural sciences, such as physics and chemistry, there are no underlying natural laws in finance; we can easily calculate when a falling feather will hit the ground, but we cannot know what the stock price of Apple will be tomorrow. Movements in the financial markets are affected by human emotions and expectations, as well as financial and non-financial events. Such events and human behaviour are hard to predict, and thus it is hard to perform controlled studies to provide new knowledge.

In the following sections, important financial concepts are described. First some basic concepts are introduced, before monetary policy and relevant financial instruments and derivatives are discussed.

2.1 Time Valuation

The concept of time value of money can be described by the statement that one dollar today is worth more than one dollar tomorrow. There are two main drivers behind this concept; investment opportunities and time preference, the latter also called human impatience by the economist Fisher [5]. Time preference refers to the general preference for present income instead of future income. Some expenses cannot be postponed until later, and buying less necessary expenses in the future instead of in the present may be associated with a risk. Even though the investor has a certain income in ten years, the desired expenses to be invested in may not be available anymore. For the investment opportunity, assume you could get $1000 today or in one year. If you choose to receive the money today, you could invest the amount and earn, say, 9.5%, or 95 dollars on the investment in one year. Choosing to receive the money in one year from now instead, you lose the opportunity to make such a profit while waiting to receive the amount.

Because of the time value of money, we need a way to compare past, present and future values. Comparing forward in time is called compounding and comparing backward in time is called discounting. Interest is compounded according to some predefined time interval, for example continuously, monthly or annually. Compounded interest is interest on the principal amount plus the accumulated interest from previous time intervals, known as interest on interest. If interest is compounded annually according to a rate r, the value at a time T in the future is given by

$$FV_T = PV(1+r)^T,$$

where FV and PV denote the future and present values, respectively. Oppositely, discounting a future amount, the present value of a future investment is given by

$$PV = FV_T(1+r)^{-T}.$$

If interest is compounded n times per year, the future value is given by

$$FV_T = PV(1+r/n)^{nT} = PV\big((1+r/n)^{n/r}\big)^{rT}.$$

As $n \to \infty$, $(1+r/n)^{n/r} \to e$, which gives

$$FV_T = PV \cdot e^{rT} \quad \Longleftrightarrow \quad PV = FV_T \cdot e^{-rT}. \tag{2.1}$$

Moving the terms and taking the logarithm of Equation (2.1), we find that the logarithmic returns are given by

$$rT = \ln \frac{FV_T}{PV}$$

from which it can easily be seen that the compound return simplifies to the sum of the interest rates. For this reason, the logarithmic returns are widely used when modelling financial time series.

2.2 Efficient Markets and the Efficient Market Hypothesis

Fama [6] defines an efficient market as a market where all available information is fully reflected in the prices. Such prices can be defined mathematically as the deviation of the expected return from the realised return, that is

$$\epsilon_{j,t+1} = \frac{p_{j,t+1} - p_{j,t}}{p_{j,t}} - \frac{\mathrm{E}[p_{j,t+1}|\Phi_t] - p_{j,t}}{p_{j,t}} = \frac{p_{j,t+1} - \mathrm{E}[p_{j,t+1}|\Phi_t]}{p_{j,t}}.$$

Here ϵ denotes the excess returns, p is the price, Φ_t is the information set at time t and j is the asset index. Using this definition, Fama introduced three models for market efficiency; the fair game, submartingale and random walk models.

The **Fair Game model** states that the information set Φ_t cannot be used to generate positive excess returns. Equivalently,

$$\mathrm{E}[\epsilon_{j,t+1}|\Phi_t] = 0.$$

This model relies on the assumption that the market equilibrium can in fact be modelled by expected returns where the equilibrium is fully reflected in the information set.

In the **Submartingale model**, the price follows a submartingale. That is, at time $t+1$, the conditional expectation of the price depends only on the history Φ_t, and is equal to or larger than the price at time t. For all times t and information sets Φ_t,

$$\mathrm{E}[p_{j,t+1}|\Phi_t] \geq p_{j,t}. \tag{2.2}$$

If Equation (2.2) is an equality, the process is called a martingale. However, the prices are expected to increase over time with the expected return, giving a higher expectation than the previously observed price. Thus, the price will follow a submartingale, while the properly discounted expected price will be a martingale [7].

Finally, the **Random Walk model** is a more thorough model, specifying all moments of the distribution of ϵ. The former two models only specify the expected value, which is the first moment. In the random walk model, the returns are assumed to be independently and identically distributed, and the process has the Markov property. Formally, these assumptions imply that the marginal and conditional distributions of the returns $r_{j,t+1}$ are equal,

$$f(r_{j,t+1}|\Phi_t) = f(r_{j,t+1}),$$

where $f(\cdot)$ denotes the probability distribution function.

Market efficiency may be categorised in three groups, depending on the type of information reflected in the current prices. The three categories of the efficient market hypothesis are:

- **Weak form:** All historical prices are fully reflected.
- **Semi-strong form:** All historical prices, including publicly available information, are fully reflected.
- **Strong form:** All information, including public and private/inside information, is reflected.

An important concept related to efficient markets, is the no-arbitrage condition. Arbitrage is simultaneously buying and selling financial instruments in two (or more) different markets to make a profit [8]. Since the action is instantaneous, there is no associated risk, and thus the no-arbitrage condition states that making an instantaneous risk-free profit is impossible. Later, we will see that this condition is frequently used as an argument when modelling prices of financial instruments.

2.3 Monetary Policy

Monetary policy refers to the available tools of central banks to regulate the economy, and more specifically the amount of money, of its country. The implemented monetary policy may vary from country to country, depending on the economic foundation and situation. Examples of such tools and instruments are policy rates, (open) market operations and reserve requirements, among others. This thesis will mainly focus on the monetary policy of Norway and the Norwegian central bank.

Within monetary policy, the main objective of Norges Bank is to maintain financial stability through a low and stable inflation, as well as a high and stable production and employment rate. Multiple other central banks share this objective. In Norway, the purpose of the central bank is enshrined in paragraph 1-2 in the Central Bank Act [9] and in the Regulation on Monetary Policy. By paragraph 3 in the latter, the operational target is an inflation close to 2% over time [10], where the inflation is measured by the consumer price index (CPI). This target is shared by the majority of Norway's main trading partners [11].

A low and stable inflation is desired as a too high inflation progressively will make the monetary value unpredictable [12]. In addition, the probability of occurrence of significantly serious economic costs caused by price variations will be much higher in an instable economic growth. On the other hand, zero inflation will typically cause a halt in the demand, resulting in a low or even declining wage growth. In the long term, this brings the economy itself to a halt, similar to what happens during a deflation. Deflation is associated with falling prices, resulting in a lower economic activity, lower wages and a higher unemployment rate. As opposed to during an inflation, where central banks can increase the policy rate to simulate the economy, deflations are harder to influence for the central banks. Lowering the policy rate far beyond zero is not desired during a deflation, as the tendency among consumers is then to save rather than invest, hoping that the monetary value will soon increase. The combination of low interest rates and high savings rates yielding ineffectiveness in the monetary policy is known as a liquidity trap [13].

In order to steer the inflation towards the desired level, Norges Bank sets the policy rate eight times yearly. Depending on the current situation and expected events, the policy rate may be raised, reduced or remain unchanged. For almost 20 years, the policy rate has never been increased by more than 25 basis points at once. On the other hand, reductions in the rate have been of the same size in calmer market situations, and of larger amplitudes in more turbulent financial periods [14]. Increasing the policy rate is an instrument used to slow down the economy, making it less attractive to spend and more attractive to save money, known as a contractionary monetary policy. On the other hand, reducing the rates, a tool of expansionary monetary policy, is used to increase spending and boost the economy.

2.3.1 Channels

The policy rate does not directly affect the population; it does so through different channels. In particular, the policy rate is the rate on overnight deposits in Norges Bank for Norwegian banks up to a given amount [15], and how the banks choose to react to the policy rate changes is of higher importance for the population than the actual changes. The banks' reaction does not solely depend on the level of the policy rate; in general, market conditions are typically just as important. To understand how the policy rate impacts the economy, some knowledge regarding the mechanism of two other interest rates is required.

If a bank's overnight deposit in Norges Bank exceeds the given amount of the policy rate, the bank receives the **reserve rate** [16] on this exceeding amount. This rate is typically 100 basis points lower than the policy rate, and thus the bank will aim to avoid exceeding the policy rate quota. If the bank instead has a negative balance in Norges Bank at the end of the day, it has to pay the **D-loan rate** so that the overnight balance is zero. The D-loan rate is typically 100 basis points above the policy rate.

Consequently, the banks will try to ensure that their overnight deposits in Norges Bank is within the specified amount, as well as above zero. If this is not the case, they typically borrow or lend among other banks at the overnight rate, a rate close to the policy rate. In this way, a bank exceeding the quota will receive a higher rate on the exceeding amount, while a bank with a negative deposit at the end of the day will have to pay a lower rate to obtain the desired zero balance. The reserve rate has been negative since March 2020, emphasising why the banks wish to participate in the overnight redistribution.

As mentioned, the policy rate affects the population and the general economy through different channels. In Norway, there are three main channels [17]:

- **The demand channel:** When the policy rate is changed, the demand levels will often move in the opposite direction as soon as the market has adjusted to the change. Thus, the effect of changes through the demand channel will be a lower or higher inflation for an increased or decreased policy rate, respectively.

- **The foreign exchange channel:** Consumers, both individuals and companies, are influenced by the exchange rate when buying or selling abroad. If the Norwegian policy rate is changed while the rates abroad stay fixed, the value of the Norwegian krone may change in the same direction as the policy rate change. An increased policy rate simultaneously lowers import prices and decreases the export, the latter resulting in a lower demand in the export sector. The foreign exchange channel thus works by changing the import prices and the demand from abroad trading partners.

- **The expectations channel:** The final channel is closely related to wages and prices on goods and services. Wages must follow the general inflation for the employees to maintain their habits - which they typically wish to do - and so the expectation of higher inflation may lead to higher inflation.

2.3.2 Recessions

For many years, the international macroeconomic trend has been positive and increasing. However, larger economic declines have occurred along the way, affecting entire economies. Such a decline is known as a recession, a significant negative fluctuation with a duration of multiple months as defined by the National Bureau of Economic Research [18]. There are many theories of what causes a recession, where economic, financial and psychological factors are some of the recurring keywords. The global Financial Crisis of 2007-2009 is an example of a recession caused by financial factors. A series of events related to financial regulations created a perfect storm in the US housing market, causing severe subprime mortgage problems in banks and hedge funds [19]. On the other hand, global events such as the outbreak of the Covid-19 pandemic in 2020 may also cause recessions. The lockdowns across numerous countries worldwide caused a halt in the economic activity, which combined with the uncertainty about the future caused a recession.

Predicting when a recession will hit is difficult. Recessions have only occurred a certain number of times in our available historical data, and the general market conditions in the time leading up to these recessions have varied. This makes it hard to build a precise predictive model. In addition, the market itself has changed over the years. However, certain variables seem to be able to indicate an upcoming recession - for example yield curve inversion, typically represented by a negative spread between a longer and a shorter yield. From 1970 to 2018, a negative term spread of the U.S. ten-year minus three-month yield always occurred before a recession [20]. D. S. Miller of FED concludes in a FED Note that there is no single best predictor of the different term spreads [21]. The best spread will vary from one situation to another, but Miller states that in general, the principal components model of Johansson and Meldrum [20] is the most accurate one.

Low interest rate environments can cause investors and other market participants to take a higher risk, as borrowing money is cheaper. When more people can afford to buy goods, services and housing, prices are pushed up. Certain prices may experience an abrupt growth, for instance housing or electricity prices. Combined with a low rate, this causes high debt in households, which again causes a higher sensitivity to lost income. Once interest rates rise as the central bank increases the policy rate to steer towards the inflation target, households with a high debt-to-income ratio are at especially high risk. Besides the risk of the term amount of the mortgage increasing with a higher interest rate, lower income or decreasing housing prices may be critical for high debt households. Because of this, a low-rate environment is not desirable over time. The propagation of an issue in a smaller part of the market can cause the market in total to collapse, as seen in the global Financial Crisis.

Market declines have their natural place in all business cycles; a market cannot grow uninterrupted forever. The actions of the central bank can play an essential role in navigating out of a recession without causing too much harm to businesses and households. Learning from previous declines and recessions, central banks and monetary organisations work both independently and internationally to constantly develop new tools and frameworks to attenuate the consequences of future recessions. For instance, the global Financial Crisis revealed a need for better risk management in banks. An extension to the existing framework was published and agreed to already by 2010. An example of how a central bank can dampen the impact and duration of a recession besides lowering the policy rate, is to purchase Treasury securities. This was done by the FED in the US during the Covid-19 pandemic, to ensure that the cash flow would not be too low for the markets to function [22]. Risk management is discussed in closer detail in Section 4.2.

2.3.3 Today's Monetary Situation

In the wake of the financial crisis of 2008-2009, policy rates in most western countries have been relatively low compared to historical levels. For instance, the Norwegian policy rate has not exceeded 2% since 2011 [14], and as of May 2022, the policy rate is 0.75%. Like the majority of other countries, the outbreak of the coronavirus in the first quarter of 2020 had a great impact on the Norwegian economy. Interest rates fell abruptly, unemployment rates rose, and companies of all sizes had to get governmental funding to avoid bankruptcy. From March 2020 until the lifting of the final restrictions in the first quarter of 2022, the trend of the economic activity has been correlated to the degree of national restrictions [23].

As a consequence of the pandemic, multiple industries experienced supply and shipping challenges leading to increased commodity prices. Combined with expensive electricity, this resulted in a higher inflation. In April 2022, the CPI was 5.4% while the adjusted CPI, CPI-ATE, was 2.6% [24], a little bit above the inflation target of 2%. Especially transportation, electricity and consumer goods contributed to the high CPI. To slow the economy, the policy rate is thus expected to increase steadily in the coming years, and the central bank has implemented a contractionary monetary policy. The Committee expects a policy rate of 2.5% by the end of 2023 [25].

In Norway, the low-rate economy has been particularly visible in the housing market, as it has enabled more people to apply for, and get, a mortgage. Housing prices have steadily increased as a result of the increasing demand, and in 2021 the housing price index for existing housing of Norway was 7.1% [26], compared to the CPI that increased by 5.3% in the same period [24]. The result of the growing debt combined with rising interest rates may be a noticeable, and even problematic, decrease in the personal finance for those with large mortgages relative to their income and assets. A research conducted by Finanstilsynet in 2021 found that the total debt-to-income ratio was above 400% for 47% of all new loans [27]. This number is close to the maximal debt-to-income ratio of 500%, and the level is high compared to historical numbers. If the wages do not increase when the policy rate increases, many young homeowners may face tighter economic times ahead.

A high inflation is not exclusive for Norway - in fact, multiple countries experience an even higher inflation than 5.4%. For instance, in March 2022, the inflation rate in the USA was 8.5%, a level it has not reached since 1981 [28]. As a result of this, the Federal Reserve increased the policy rate by 0.5 percentage points for the first time since 2000 in May 2022 [29]. It was questioned whether the increase would be as large as 75 basis points, but the Federal Reserve did not find it necessary; however, more 50 basis points increases are expected in the upcoming time.

Russia's invasion of Ukraine and the following war has exaggerated the effect of rising prices; usually, both Russia and Ukraine are important international suppliers of essential goods, such as oil and wheat. The International Monetary Fund states in their World Economic Output report of April 2022 that certain supply chains will experience issues until 2023, yielding a higher expected inflation for the next years [30]. The same report concludes that central banks should closely follow the economic situation and take the necessary actions in time. Focusing on the long-term economy in addition to absorbing the impact of the war in Ukraine will be important for the countries as well as their residents in the upcoming months.

2.4 Interest Rates and Financial Instruments

In the modern society, we all get in touch with interest rates at some point in our lives. Our bank may pay us interest on our bank deposits, and when we repay our mortgage loan, we usually pay more than the originally borrowed amount. The interest rate is the additional fraction of the principal amount the borrower must pay the lender per time unit - essentially it is a charge for using the lender's assets. Oppositely, for the lender, the interest rate equals the rate of return. The lender's interest is of course to regain his or her assets from the borrower, so a borrower that is considered riskier will have to pay more interest than a less risky borrower. Furthermore, loans with longer maturities usually have a higher interest rate because of the potential risks taken. In addition to the baseline risk, potential losses can occur from not investing the lent amount in a possibly higher profit yielding investment.

Unlike stocks, we cannot buy interest rates; this makes modelling interest rates and pricing instruments technically harder as we have no underlying asset to hedge with. In addition, the relationship between interest rates of different maturities must be considered when modelling the rates. This must be done in a sensible manner, such that the fundamental properties of interest rates are maintained. For instance, when we model forward rates, compounded one-year rates up to five years in the future should not be lower than the five-year rate; if that was the case, nobody would have chosen the five-year interest rate on their loan.

Most interest rates are compounded according to some predefined time interval, for example continuously, monthly or annually. The compound interest is interest on the principal amount plus the accumulated interest from previous time intervals - or simply just interest on interest. If interest is compounded annually, the growth factor of an investment is

$$(1 + r_d)^{T-t}, \tag{2.3}$$

where r_d is a discretely compounded interest rate and $T - t$ is the time interval for the growth, here given in years.

Financial instruments are contracts holding a monetary value. Instruments can be based on both equity and debt, where stock options are equity-based, and interest rate swaps are debt-based examples. Furthermore, financial instruments are typically divided into two groups; cash and derivative instruments [31]. In this thesis, we will focus on derivatives, whose value is dependent on some underlying quantity. Specifically, we will model a portfolio of bonds.

A bond is a contract yielding a known amount, the par value, to the holder at maturity. Furthermore, it may also pay coupons at fixed times during the lifetime of the contract. It is a security typically issued by governments, municipalities or big companies to finance projects, and can be thought of as a loan where the investor is the lender, and the issuer is the borrower[8]. The value to be paid back to the lender at maturity is called the par or face value. Coupon payments can be paid at fixed times according to a fixed or a floating rate, or not paid at all - if so, the bond is said to be a zero-coupon bond.

For simplicity, we use zero-coupon bonds when we model the interest rates in this thesis. The price of an annually compounded zero-coupon bond $Z(t,T)$ with face value 1 is the inverse of the growth factor given in Equation (2.3),

$$Z(t,T) = (1+r_d)^{-(T-t)}, \tag{2.4}$$

where t is the time at which the zero-coupon bond is priced and T is the maturity.

Finally, we present two central concepts of the thesis:

Definition 1 (Yield curve)**.** *The yield curve $Y(t)$ shows the yield, i.e., the interest rate, for a financial instrument as a function of different available maturities at time t.*

Different shapes of the yield curve, such as increasing, flat or inverted, have different macroeconomic implications. Some examples of different yield curves are presented and discussed in Chapter 5.

Definition 2 (Forward rate)**.** *The forward rate $f(t,T)$ is the future interest rate on a bond investment starting at some specified time in the future, calculated from today's yield curve.*

The zero-coupon bond price in Equation (2.4) can also be expressed in terms of the forward rate,

$$Z(t,T) = \exp\big\{-\int_t^T f(t,s)ds\big\}, \tag{2.5}$$

from which it directly follows that the forward rate can be expressed by the zero-coupon bond price by the equation

$$f(t,T) = -\frac{\partial}{\partial T}\log Z(t,T). \tag{2.6}$$

Chapter 3

Theory and Models

Quantitative finance makes use of multiple mathematical concepts, such as stochastic differential equations, numerical mathematics and statistics. This chapter starts by presenting Brownian motion and Itô's lemma, two cornerstones in financial mathematics. Then we discuss the related Bayesian hierarchical models and Monte Carlo methods, as well as INLA, a special case of a Bayesian model and an alternative to Monte Carlo methods. Finally, we present a more widely used method for modelling interest rates and some preliminary theory for this model; the Heath, Jarrow & Morton framework and principal component analysis, respectively.

3.1 Brownian Motion

Brownian motion, better known as the Wiener process in mathematics, is a continuous stochastic process that in finance is used for its prominent representation of random processes in modelling of both stocks and stochastic interest rates. Mathematically, the following properties describe a Wiener process [32]:

Definition 3 (Wiener Process). *Let W be a Wiener process. Then*

1. $W_0 = 0$.
2. *For $t > 0$, the increments $W_{t+1} - W_t$ are independent of all past values.*
3. *The increments of W are Gaussian, i.e., for $t, \tau \geq 0$, $W_{t+\tau} - W_t$ is normally distributed with mean 0 and variance τ.*
4. *The paths of W are continuous, but nowhere differentiable.*

In practice, Definition 3 yields some important properties for the modelled quantity [33]. The Brownian motion in quantitative models will be Markovian, that is W_{t+1} will be independent of all past values except for the present value W_t, and the martingale property holds, i.e., $E[W_t|\tau < t] = W_\tau$. Furthermore, since the Brownian motion is a random walk, we have almost surely quadratic variation,

$$\sum_{i=1}^{n}(W_{t_i} - W_{t_{i-1}})^2 \to t,$$

when the full time interval $[0, t]$ is partitioned into $n+1$ parts $t_i = \frac{it}{n}$, and $n \to \infty$. The chosen scaling of the increment by the square root of the time increment ensures that the model is finite; any other choice of the scaling factor would either blow up the model or make it converge to a limit yielding no motion [8].

3.2 Itô's Lemma

Itô's Lemma is the stochastic variable equivalent to Taylor's theorem in classical calculus, and it plays a highly important role in quantitative finance as we need a way to study the evolution of stochastic processes. Examples of use include derivation of the famous Black-Scholes model, pricing of derivatives and market risk management.

Theorem 1 (Itô's Lemma)**.** *Assume that $f(S, t)$ is a smooth function of the stochastic process S and time t, and that S satisfies*

$$dS = \mu S dt + \sigma S dX,$$

where dX is a Wiener process and μ and σ are the mean and volatility, respectively, of the process. Then, if $dX^2 \to dt$ as $dt \to 0$ with probability 1, and $(dt)^2 = 0$,

$$\begin{aligned} df &= \frac{\partial f}{\partial S}(\sigma S dX + \mu S dt) + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} dt \\ &= \sigma S \frac{\partial f}{\partial S} dX + \big(\mu S \frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} + \frac{\partial f}{\partial t}\big) dt \end{aligned}$$

For a derivation of Itô's lemma in a quantitative financial setting the reader is referred to Wilmott [8]; or, for a more general derivation and proof, to the original by Itô [34].

3.3 Bayesian Hierarchical models

Hierarchical modelling is used for multilevel statistical models where the observations are clustered in groups sharing some properties. In this thesis, the Bayesian method is preferred over the frequentist method because of its ability to account for uncertainty, as well as providing interesting information about the parameters and their intervals [35].

As opposed to the controlled experiments of frequentists, the Bayesian method makes an assumption about the likelihood and combines it with a prior distribution to find the posterior distribution by using Bayes' theorem:

Theorem 2 (Bayes' theorem)**.** *Let A and B be two events with $P(B) \neq 0$. Then*

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{3.1}$$

Here $P(\cdot|\cdot)$ denotes a conditional probability, where the first event is conditional on the second.

In the Bayesian setting, we assume that $\boldsymbol{\theta}$ is a vector of unknown parameters and $\mathbf{y}$ is the vector of observations. Note that elements of $\boldsymbol{\theta}$ are not only parameters, but also stochastic variables. Then $P(\mathbf{y})$ and $P(\boldsymbol{\theta})$ are the priors and

$$P(\boldsymbol{\theta}|\mathbf{y}) = \frac{P(\mathbf{y}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathbf{y})}$$

is the posterior we wish to find. As the denominator does not depend on $\boldsymbol{\theta}$, our quantity of interest, we simplify calculations by finding the proportional instead;

$$P(\boldsymbol{\theta}|\mathbf{y}) \propto P(\mathbf{y}|\boldsymbol{\theta}) \times P(\boldsymbol{\theta})$$
$$\text{Posterior} \propto \text{Likelihood} \times \text{Prior}$$

3.4 Monte Carlo Integration and Markov Chain Monte Carlo

In finance, Monte Carlo methods are used in two different settings; to explore properties in order to find quantities of interest, and to find the value of derivatives. Modelling financial models as random walks and simulating a large number of scenarios allow us to estimate interesting quantities such as value at risk from the simulated data [36].

The idea behind Monte Carlo integration is to simulate n random draws $\boldsymbol{X_1}$, $\boldsymbol{X_2}$, ..., $\boldsymbol{X_n}$ from the target distribution f to estimate the value of an integral. From basic statistical theory, the expectation of a function h of a random variable $\boldsymbol{X}$, where f is the density of $\boldsymbol{X}$, is given by

$$\mu = E[h(\boldsymbol{X})] = \int h(\boldsymbol{X}) f(\boldsymbol{X}) d\boldsymbol{X}.$$

In many cases, this integral is not analytically tractable. As an alternative to numerical approximation, the mean μ may be estimated by simulation. Now, let $E[h(\boldsymbol{X})] = \mu$, and let $\boldsymbol{X_1}, \ldots, \boldsymbol{X_n}$ be independent and identically distributed random samples from our target distribution f. By the strong law of large numbers, the Monte Carlo estimate of μ is simply the average of the sample:

$$\hat{\mu}_{MC} = \frac{1}{n} \sum_{i=1}^{n} h(\boldsymbol{X_i}).$$

The variance of this estimator is given by

$$\hat{S}^2 = \frac{1}{n-1} \sum_{i=1}^{n} [h(\boldsymbol{X_i}) - \hat{\mu}_{MC}],$$

and by the central limit theorem, the distribution of $\hat{\mu}_{MC}$ converges to a normal distribution when $n \to \infty$. The convergence rate is relatively slow, $\mathcal{O}(n^{-1/2})$ [37].

When dealing with high-dimensional or complex distributions, Markov chain Monte Carlo (MCMC) is a popular choice. Introducing a Markov chain to the Monte Carlo integration, we get a process converging to the target distribution f so that we can perform inference for this distribution. The Markov chain is a stochastic process satisfying the Markov property, like the Brownian motion in Chapter 3.1 - which is, in fact, a special case of a Markov chain.

Using MCMC, we can perform inference for the distribution of interest by sampling non-uniform distributions. Compared to similar approaches, increased dimensionality in an MCMC algorithm does not change the speed of convergence. Multiple MCMC algorithms exist, with some of the most popular ones being the

Metropolis-Hastings algorithm and the Gibbs sampler. Before taking a closer look at Metropolis-Hastings to understand the mechanisms of a MCMC algorithm, we need to introduce two concepts [37].

Theorem 3 (Unique Limiting Distribution)**.** *If the Markov chain is irreducible, aperiodic and positive recurrent, it has a unique limiting distribution $f(x)$. For a state space S, the limiting distribution is given by*

$$f(y) = \sum_{x \in S} f(x)P(y|x) \quad \forall y \in S$$
$$\sum_{x \in S} f(x) = 1$$

Definition 4 (The detailed balance condition)**.** *The detailed balance condition is a sufficient condition for Theorem 3:*

$$f(x)P(y|x) = f(y)P(x|y) \quad \forall x, y \in S.$$

If this condition is satisfied, the Markov chain is time reversible.

The majority of all MCMC algorithms use reversible Markov chains, i.e., processes that satisfy the definition and theorem above. We now take a closer look at the Metropolis-Hastings algorithm, which in short can be described as follows:

1. Draw the initial state $\boldsymbol{x_0} = g(\boldsymbol{x_0})$ such that $f(\boldsymbol{x_0}) > 0$.
2. Generate a proposal $\boldsymbol{X}^*$ from the proposal distribution $g(\cdot|\boldsymbol{x_t})$.
3. Draw $u \sim$ Unif(0, 1).
4. Calculate the acceptance probability $\alpha = \min\left\{1, \frac{f(\boldsymbol{y})g(\boldsymbol{x_t}|\boldsymbol{y})}{f(\boldsymbol{x_t})g(\boldsymbol{y}|\boldsymbol{x_t})}\right\}$
5. If $u < \alpha$: $\boldsymbol{X_{t+1}} = \boldsymbol{X}^*$; otherwise $\boldsymbol{X_{t+1}} = \boldsymbol{x_t}$.
6. Increment t and repeat steps 2-6 until the desired amount of iterations, n, is performed.

The number of iterations n is chosen such that the Markov chain will have time to converge to the stationary distribution, and typical choices are $n > 1000$. Depending on the proposal distribution g, we usually discard the m first samples, called the burn-in period, to get rid of irrelevant effects of the initial proposal. If running a single simulation, the easiest way to find m is by plotting the Markov chain and locating where the process stabilises. For more details on the Metropolis-Hastings algorithm, the reader is referred to Givens & Hoeting [37].

3.5 Integrated Nested Laplace Approximation

Integrated nested Laplace approximation (INLA) is an alternative to Markov chain Monte Carlo for performing Bayesian inference. Instead of using the rather slow MCMC sampling methods to inference, INLA is based on Laplace approximations and numerical integration, yielding a computationally fast and flexible approach. Before the INLA approach was developed, Gaussian approximation-based proposals were popular to try to reduce the runtime of the MCMC algorithms. However, even the fastest samplers were slow; and so, the idea of inference using closed form Gaussian approximations arose, leading to the INLA framework.

INLA can be applied to a broad field of models, including almost all types of hierarchical Bayesian models. The common property of these models is the structure of the latent field. This field must be a Gaussian Markov random field (GMRF):

Definition 5 (Gaussian Markov random field). *A GMRF $\boldsymbol{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\sigma})$, $\boldsymbol{x} = (x_1, \ldots, x_n)$ is a Gaussian random variable that satisfies the Markov property, that is*

$$P(x_i | \{x_j : j \neq i\}) = P(x_i | \{x_j : j \in \mathcal{N}_i\})$$

where $\mathcal{N}_i$ are the neighbours of a point x_i. The Markov property is encoded in the precision matrix $\boldsymbol{Q} = \boldsymbol{\Sigma}^{-1}$ with entries $Q_{ij} = 0 \iff j \notin \{i, \mathcal{N}_i\}$ [3].

The first stage of an INLA model consists of the observations $\mathbf{y}$, assumed to be conditionally independent given the latent field $\mathbf{x}$ and possibly some parameters $\boldsymbol{\theta}_1$,

$$\boldsymbol{y} | \boldsymbol{x}, \boldsymbol{\theta}_1 \sim \prod_i \pi(y_i | x_i, \boldsymbol{\theta}_1),$$

where $i = 1, \ldots, n$ and n is the number of variables in the latent field.

The unobserved process makes up the second stage of the model, and this latent model is assumed to be a GMRF with a sparse precision matrix $\mathbf{Q} = \boldsymbol{\Sigma}^{-1}$. This field can be large, with up to 10^6 dimensions, and is on the form

$$\boldsymbol{x} | \boldsymbol{\theta}_2 \sim \mathcal{N}(\mathbf{0}, \boldsymbol{Q}(\boldsymbol{\theta}_2)^{-1}),$$

where $\mathcal{N}(\cdot, \cdot)$ denotes the Gaussian distribution. The sparsity of the inverse covariance matrix is an essential part of the INLA method, as it enables using efficient sparse-matrix calculations instead of standard dense-matrix calculations.

Finally, the third stage consists of the hyperparameters of the likelihood and the precision parameter of the GMRF. This corresponds to the prior distribution of $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$. The number of hyperparameters j should be small; $j \leq 9$ [38].

Using Bayes' theorem given in Equation (3.1), we find that the posterior distribution is given by

$$\pi(\boldsymbol{x}, \boldsymbol{\theta} | \boldsymbol{y}) \propto \pi(\boldsymbol{\theta}) \pi(\boldsymbol{x} | \boldsymbol{\theta}) \prod_i \pi(y_i | x_i, \boldsymbol{\theta}).$$

However, this is not the main quantity of interest when using INLA. Instead, we wish to approximate the posterior marginals of the hyperparameters and the latent Gaussian variables, i.e., $\pi(\theta_j|\mathbf{y})$ and $\pi(x_i|\mathbf{y})$ for $i = 1, \ldots, n$, respectively. Again, using Bayes' theorem, we find these to be given by

$$\pi(\theta_j|\boldsymbol{y}) = \int_{\boldsymbol{\theta}_{-j}} \int_{\boldsymbol{x}} \pi(\boldsymbol{x}, \boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{x} d\boldsymbol{\theta}_{-j} = \int_{\boldsymbol{\theta}_{-j}} \pi(\boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{\theta}_{-j} \tag{3.2}$$

and

$$\begin{aligned} \pi(x_i|\boldsymbol{y}) &= \int_{\boldsymbol{\theta}} \int_{\boldsymbol{x}_{-i}} \pi(\boldsymbol{x}, \boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{x}_{-\boldsymbol{i}} d\boldsymbol{\theta} = \int_{\boldsymbol{\theta}} \pi(x_i, \boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{\theta} \\ &= \int_{\boldsymbol{\theta}} \pi(x_i|\boldsymbol{\theta}, \boldsymbol{y}) \pi(\boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{\theta} \end{aligned} \tag{3.3}$$

The posterior marginals are found by approximating each of the components in the integrals, and then integrating out the remaining variables not of interest. In Equation (3.3), the approximation is

$$\hat{\pi}(\boldsymbol{\theta}|\boldsymbol{y}) \propto \left. \frac{\pi(\boldsymbol{x}, \boldsymbol{\theta}, \boldsymbol{y})}{\hat{\pi}_G(\boldsymbol{x}|\boldsymbol{\theta}, \boldsymbol{y})} \right|_{\boldsymbol{x}=\boldsymbol{x}^*(\boldsymbol{\theta})} \tag{3.4}$$

where the normalising constant is neglected, and $\hat{\pi}_G(\boldsymbol{x}|\boldsymbol{\theta}, \boldsymbol{y})$ is the Gaussian approximation of the full conditional of $\boldsymbol{x}$. $\boldsymbol{x}^*(\boldsymbol{\theta})$ is the mode of this full conditional, given a $\boldsymbol{\theta}$. The reader is referred to Chopin, Martino and Rue [3] for the full details of deriving the expression. We note that the numerator is simple to compute, while the denominator is more complex and expensive. For each value $\boldsymbol{\theta^k}$, a Cholesky decomposition is used to compute the precision matrix $\boldsymbol{Q}(\boldsymbol{\theta^k})$. It is possible to find the posterior marginals $\hat{\pi}(\theta_j|\boldsymbol{y})$ by numerically integrating Equation (3.4). However, doing so by using the quantities already found while deriving Equation (3.4), the process is much less computationally demanding.

When the hyperparameters $\boldsymbol{\theta}$ are estimated, we can find $\hat{\pi}(x_i|\boldsymbol{\theta}, \boldsymbol{y})$. The approximation can be performed in multiple ways. The Laplace approximation is preferred for its accuracy, but with the high accuracy, an expensive computational cost follows. To significantly reduce the cost, while still being more accurate than the simplest methods, a simplified Laplace approximation is chosen for the analysis in this thesis. The simplified Laplace is also the default strategy used in the R-package of INLA [39], and the approximation is based on

$$\hat{\pi}_{LA}(x_i|\boldsymbol{\theta}, \boldsymbol{y}) \propto \left. \frac{\pi(\boldsymbol{x}, \boldsymbol{\theta}, \boldsymbol{y})}{\hat{\pi}_{GG}(\boldsymbol{x}_{-i}|x_i, \boldsymbol{\theta}, \boldsymbol{y})} \right|_{\boldsymbol{x}_{-i}=\boldsymbol{x^*}_{-i}(x_i, \boldsymbol{\theta})}$$

where $\hat{\pi}_{GG}$ is a different Gaussian approximation than the one used in Equation (3.4), and $\boldsymbol{x^*}_{-i}(x_i, \boldsymbol{\theta})$ is the corresponding modal configuration. Although $\boldsymbol{x}$ is a sparse matrix, the expression is very expensive to evaluate. Hence, the approximation is simplified by combing elements of the above equation and a plain Gaussian approximation. In particular, the simplified Laplace computes the Taylor series around the Laplace node, correcting the issues of the location and skew of the Gaussian approximation [39]. Finally, the integral in Equation (3.3) is solved by a numerical scheme.

3.6 Principal Component Analysis and Singular Value Decomposition

Principal component analysis (PCA) is a technique that transforms a set of possibly correlated variables to another set of linearly independent variables through an orthogonal transformation. The transformation is done in such a manner that the first principal components explain the most variance, which is why PCA is mainly used to reduce the number of dimensions in complex models.

Let $\boldsymbol{X}$ be a standardised and centralised $(m \times n)$ matrix, where m is the number of samples and n is the number of variables in our original data set, and denote the associated covariance matrix by $Cov[\boldsymbol{X}] = \boldsymbol{\Sigma}$. Then $\boldsymbol{\Sigma}$ is a square $(n \times n)$ and positive semi-definite matrix which can be diagonalized by

$$\boldsymbol{\Sigma} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^{-\mathbf{1}},$$

where the columns of $\boldsymbol{V}$ are the eigenvectors and $\boldsymbol{\Lambda} = \text{Diag}\{\lambda_1, \ldots, \lambda_n\}$ is a diagonal matrix which elements are the eigenvalues of $\boldsymbol{\Sigma}$. The eigenvectors are then the principal directions, and the corresponding eigenvalues are sorted in descending order where the highest eigenvalue explains the most variance.

The eigenvector corresponding to the highest eigenvalue will point in the direction explaining the most variance, i.e., the most important direction of the data. We can then reduce the dimension of the data from n dimensions to $p, p < n$, depending on how much explained variance we wish to keep in our model. The more correlated the original variables are, the fewer principal components are needed to explain a larger proportion of the variance. Multiple different methods exist for choosing the exact number of principal components to include in the model, but quite a few of these are ad hoc, without a formal justification. An example of such a method is visualising the principal components and choosing those who, from visual inspection, capture certain properties of the curve. Regardless of the missing formalities, this method has shown to be quite plausible and work well in practice [40].

The traditional PCA method where we decompose the covariance matrix is not optimal in all situations. In terms of numerical accuracy, singular value decomposition (SVD) is preferred because it extracts data in the highest variance directions compared to PCA which is quite simply a linear transformation of the data [40]. Let $\boldsymbol{X}$ be as above. Instead of finding the covariance matrix, we now find the eigenvectors and eigenvalues of $\boldsymbol{XX^T}$ and $\boldsymbol{X^TX}$, matrices which both carry some pleasant properties; independent of the structure of $\boldsymbol{X}$ they are symmetrical, square, of the same rank as $\boldsymbol{X}$, and either positive semi-definite or positive-definite. From these matrices, the SVD is given by

$$\boldsymbol{X} = \boldsymbol{USV^T}.$$

Here $\boldsymbol{U}$, with dimension $(m \times m)$, and $\boldsymbol{V}$, with dimension $(n \times n)$, are found from the orthonormal eigenvectors from $\boldsymbol{XX^T}$ and $\boldsymbol{X^TX}$, respectively. Both $\boldsymbol{U}$ and $\boldsymbol{V}$ are thus invertible and orthogonal matrices, the latter meaning that $\boldsymbol{U^TU} = \boldsymbol{I}_{m\times m}$ and $\boldsymbol{V^TV} = \boldsymbol{I}_{n\times n}$. $\boldsymbol{S}$ is a diagonal $(m \times n)$ matrix where the upper square is a diagonal matrix which elements are the singular values, i.e., $\text{Diag}\{\sigma_1, \ldots, \sigma_n\}$, where $\sigma_i = \lambda_i^2$.

Using the SVD results to perform PCA, we follow the procedure described for the traditional PCA. However, the columns of $\boldsymbol{V}$ are now the principal directions, while $\boldsymbol{XV} = \boldsymbol{USV^TV} = \boldsymbol{US}$ are the principal components and the eigenvalues are as mentioned above. When $m \gg n$, $\boldsymbol{U}$ and $\boldsymbol{S}$ are quite large - however, computer software usually deals with possible issues caused by this.

3.7 The Heath, Jarrow & Morton Framework

The Heath, Jarrow & Morton (HJM) framework is used for forward rate modelling. Instead of the classical derivation of the forward curves from short rate models, the concept of HJM is to model the evolution of the whole forward curve, denoted by $f(t,T)$ at time t where T is the maturity. Yield curve fitting is contained in the model as there is an explicit relation between the yield and forward curves.

Under the HJM framework, the modelled quantity may take negative values. Earlier this was regarded as a disadvantage; however, policy rates have now taken on negative values, such as at the European Central Bank where the rate was lowered from 0% to −0.10% in June 2014 [41].

As we wish to model both short and long forward rates, we follow the multi-factor HJM framework with d sources of randomness. From Equation (2.5) and Equation (2.6) in the previous chapter, we found the relationship between the zero-coupon bond price Z and the forward rate f to be given by

$$\begin{aligned} Z(t,T) &= \exp\Big\{-\int_t^T f(t,s)ds\Big\} \\ \iff f(t,T) &= -\frac{\partial}{\partial T}\log Z(t,T). \end{aligned}$$

Furthermore, we assume that changes in zero-coupon bonds are on the form

$$dZ(t,T) = \mu(t,T)Z(t,T)dt + \boldsymbol{\sigma}^T(t,T)Z(t,T)d\boldsymbol{W}, \tag{3.5}$$

where $\boldsymbol{\sigma}$ is a $(1\times d)$ vector of the volatility factors and $\boldsymbol{W}$ is a d-dimensional Wiener process. Combining Equation (2.6) and Equation (3.5), and differentiating with respect to t, the forward curve evolves according to the differential equation

$$df(t,T) = \frac{\partial}{\partial T}\big(\frac{1}{2}\boldsymbol{\sigma}^T(t,T)\boldsymbol{\sigma}(t,T) - \mu(t,T)\big)dt - \frac{\partial}{\partial T}\boldsymbol{\sigma}^T(t,T)d\boldsymbol{W}. \tag{3.6}$$

The two dt-terms, which are deterministic, express the drift in terms of volatility of the curves and in terms of the bond, respectively. We denote the factor associated with the random component $d\boldsymbol{W}$ by $\boldsymbol{\nu}(t,T)$,

$$\boldsymbol{\nu}(t,T) = -\frac{\partial}{\partial T}\boldsymbol{\sigma}(t,T), \tag{3.7}$$

which is the forward rate volatility function. The k'th volatility factor is given by

$$\boldsymbol{\nu}_k(t,T) = \sqrt{\lambda_k}\boldsymbol{u}_k, \tag{3.8}$$

where λ_k is the k'th eigenvalue and $\boldsymbol{u}_k$ is the k'th eigenvector, for $k = 1,\dots,d$.

An important assumption in the HJM framework is the no-arbitrage argument. Derivatives must be priced in risk-neutral settings to avoid arbitrage opportunities, which corresponds to $\mu(t,T) = r(t)$. The interested reader is referred to Heath, Jarrow and Morton [42] for the derivation and proof of this statement. In the limit the spot rate $r(t)$ equals $f(t,t)$, the forward rate with maturity today, and thus $\frac{\partial}{\partial T}\mu(t,T) = 0$. The forward rate drift is then related to the volatility by

$$m(t,T) = \frac{\partial}{\partial T}\big(\frac{1}{2}\boldsymbol{\sigma}^T(t,T)\boldsymbol{\sigma}(t,T) - \mu(t,T)\big) = \boldsymbol{\nu}^T(t,T)\int_t^T \boldsymbol{\nu}(t,s)ds. \tag{3.9}$$

Inserting Equations (3.7) and (3.9) into (3.6), we obtain the stochastic differential equation

$$df(t,T) = \left(\boldsymbol{\nu}^T(t,T)\int_t^T \boldsymbol{\nu}(t,s)ds\right)dt + \boldsymbol{\nu}^T(t,T)d\boldsymbol{W}, \tag{3.10}$$

which models the evolution of the forward rate curves in the multi-factor HJM framework.

Using HJM to price derivatives, we first model Equation (3.10) for a given amount of days, e.g., 252 days, which corresponds to one trading year. Then, the desired derivative is priced from the single simulation, and the process is repeated n times to obtain a sample from which the mean is the resulting price of the derivative. If we wish to investigate the price in closer detail, we repeat the whole process m times to obtain a distribution.

3.7.1 Discretisation of the HJM model

In most cases, exact simulation of the HJM model is impossible; only a few choices of $\sigma(t,T)$ allows exact simulation of Equation (3.10). In general, we must discretise the model in both arguments t and T of $f(t,T)$. The discretisation presented in this section follows the arguments of Glasserman [43].

We discretise the first argument t on the time grid $0 = t_0 < t_1 < \cdots < t_M$. In general, the complete forward curve $f(t_i,T)$ cannot be represented for a fixed time $t_i, t_i \leq T$, and thus the maturity grid also needs to be discretised. To simplify notation, which can here be done nearly without loss of generality, we let the discretisation of T be the same as the one for t.

In the further analysis, discrete variables are denoted with a hat to distinguish from their continuous equivalents; $\hat{f}(t_i,t_j)$ denotes the discretised forward rate at time t_i with maturity t_j for $i \leq j$. The discrete version of the zero-coupon bond price is then on the form

$$\hat{Z}(t_i,t_j) = \exp\left\{ -\sum_{l=1}^{j-1} \hat{f}(t_i,t_l)[t_{l+1}-t_l] \right\}. \tag{3.11}$$

Discretising a model yields a discretisation error. To avoid the introduced error to be larger than necessary, we let $\hat{Z}(0,t_j) = Z(0,t_j)$ for all maturities t_j. Combining Equations (2.5) and (3.11), we obtain

$$\sum_{l=0}^{j-1} \hat{f}(0,t_l)[t_{l+1}-t_l] = \int_0^{t_j} f(0,s)ds,$$

that is,

$$\hat{f}(0,t_l) = \frac{1}{t_{l+1}-t_l}\int_{t_l}^{t_{l+1}} f(0,u)du \quad \forall\ l = 0,\dots M-1$$

which corresponds to each $\hat{f}(0,t_l)$ being the average level of the continuous forward rate $f(0,T)$ on the interval $[t_l,t_{l+1}]$. The discretised evolution of the forward curve is thus

$$\begin{aligned}\hat{f}(t_i,t_j) =& \hat{f}(t_{i-1},t_j) + \hat{\mu}(t_{i-1},t_j)[t_i-t_{i-1}]+ \\ & \sqrt{t_i-t_{i-1}}\hat{\boldsymbol{\nu}}(t_{i-1},t_j)\boldsymbol{Z}_i\end{aligned} \tag{3.12}$$

for $i,j = 1,\dots,M$, where $\boldsymbol{Z}_1,\dots,\boldsymbol{Z}_M$ are independent standard normal d-variate variables.

The discretised risk $\hat{\mu}(t_{i-1},t_j)$ consists of the d contributions from each of the volatility factors,

$$\hat{\mu}(t_{i-1},t_j) = \sum_{k=1}^{d} \hat{\mu}_k(t_{i-1},t_j).$$

This drift can be derived in multiple ways; here, we derive it from the no-arbitrage condition, which corresponds to the discounted bond prices being martingales. With the discrete short rate $\hat{f}(t_k, t_k)$, this condition implies that

$$\hat{Z}(t_i, t_j) \exp\Big\{ - \sum_{k=0}^{i-1} \hat{f}(t_k, t_k)[t_{k+1} - t_k]\Big\}$$

is a martingale in i for all j. The martingale can be expressed by

$$\begin{aligned} &\mathrm{E}\big[\hat{Z}(t_i, t_j) \exp\big\{ - \sum_{k=0}^{i-1} \hat{f}(t_k, t_k)[t_{k+1} - t_k]\big\} | \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_{i-1}\big] \\ &= \hat{Z}(t_{i-1}, t_j) \exp\big\{ - \sum_{k=0}^{i-2} \hat{f}(t_k, t_k)[t_{k+1} - t_k]\big\}, \end{aligned}$$

which by Equation (3.11) and cancelling terms yields

$$\begin{aligned} &\mathrm{E}\big[\exp\big\{ - \sum_{l=i}^{j-1} \hat{f}(t_i, t_l)[t_{l+1} - t_l]\big\} | \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_{i-1}\big] \\ &= \exp\big\{ - \sum_{l=i}^{j-1} \hat{f}(t_{i-1}, t_l)[t_{l+1} - t_l]\big\}. \end{aligned}$$

Substituting in the evolution of the discretised forward rate given by Equation (3.12) gives

$$\begin{aligned} &\mathrm{E}\big[\exp\big\{ - \sum_{l=i}^{j-1} \sqrt{t_i - t_{i-1}} \hat{\boldsymbol{\nu}}^T(t_i, t_l) \boldsymbol{Z}_i [t_{l+1} - t_l]\big\} | \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_{i-1}\big] \\ &= \exp\big\{ - \sum_{l=i}^{j-1} \sum_{k=1}^{d} \hat{\mu}_k(t_{i-1}, t_l)[t_i - t_{i-1}][t_{l+1} - t_l]\big\}. \end{aligned}$$

Evaluating the expectation on the left-hand side and taking the logarithm, we obtain

$$\frac{1}{2} \sum_{k=1}^{d} \big(\sum_{l=i}^{j-1} \hat{\nu}_k(t_{i-1}, t_l)[t_{l+1} - t_l]\big)^2 = \sum_{k=1}^{d} \sum_{l=i}^{j-1} \hat{\mu}_k(t_{i-1}, t_l)[t_{l+1} - t_l],$$

from which we find the k'th contribution to the drift, $\hat{\mu}_k(t_{i-1}, t_j)$, that satisfies the martingale condition for any number d of volatility factors. This contribution is given by

$$\begin{aligned} \hat{\mu}_k(t_{i-1}, t_j) = \frac{1}{2(t_{j+1} - t_j)} \Bigg[&\Big(\sum_{l=i}^{j} \hat{\nu}_k(t_{i-1}, t_l)[t_{l+1} - t_l]\Big)^2 \\ &- \Big(\sum_{l=i}^{j-1} \hat{\nu}_k(t_{i-1}, t_l)[t_{l+1} - t_l]\Big)^2 \Bigg]. \end{aligned}$$

3.8 Model Selection Criteria

To decide which model is the best one out of a set of models, different criteria can be used. The most used one is perhaps the Akaike information criterion (AIC), which aims to find the best model in terms of the optimal balance of over- and underfitting [44]. The criterion can be used on Bayesian models, but it is mostly used by frequentists as it comes short in terms of some of the fundamental properties of Bayesian hierarchical models. In this section we discuss four different model selection criteria, of which all can easily be found when using the INLA R-package.

In this thesis, we want to test whether a variable should be included in a model or not. We thus use the same priors and initial values for the hyperparameters, and only differ the formulas of the models we want to compare. Then, the comparison is focused on the variable, and we avoid bias from other factors.

3.8.1 Deviance Information Criterion

The deviance information criterion (DIC) is a generalised version of AIC to enable better comparison of Bayesian models [45]. The AIC is given by

$$AIC = 2k - 2\ln \hat{L},$$

where k is the number of parameters and $\hat{L}$ is the maximum likelihood function of the model. The DIC is similar to AIC in many ways, with a main difference in the way the priors are considered. The DIC is defined as

$$DIC = 2p_D - 2\ln \pi(y|\hat{\theta}),$$

where $\pi(y|\hat{\theta})$ is the log-likelihood given the Bayesian estimate $\hat{\theta}$ and p_D is the effective number of parameters, given by

$$p_D = \ln \pi(y|\hat{\theta}) - \mathrm{E}[\ln \pi(y|\theta)].$$

Here, $\mathrm{E}[\cdot]$ denotes the posterior expected value. Using DIC to compare two or more models, the criterion favours the model with the lowest measure.

3.8.2 Watanabe-Akaike Information Criterion

The Watanabe-Akaike information criterion, also called the widely applicable information criterion (WAIC) is similar to the DIC, but the effective number of parameters is replaced by

$$p_{\text{WAIC}} = \sum_{i=1}^{n} \mathrm{Var}[\ln \pi(y_i|\theta)].$$

Summarising the variances from each observation gives a more stable expression than the one used in the DIC. The criterion is then given by

$$WAIC = 2p_{WAIC} - 2\sum_{i=1}^{n} \ln \int \pi(y_i|\theta)\pi(\theta|y)d\theta$$

where the latter term is the logarithmic pointwise predictive distribution [46].

3.8.3 Conditional Predictive Ordinal

Another Bayesian model selection criterion is the conditional predictive ordinal (CPO), which selects the best models in terms of how surprising new observations are. The measure is given by

$$CPO_i = \pi(y_i|\boldsymbol{y_{-i}})$$

where $\pi(\cdot|x)$ denotes the predictive distribution given some history x [47]. A smaller CPO corresponds to a larger effect of the new observation y_i, and thus the new observation can result in an underfitting model, or it can be an outlier. We note that a CPO can be unreliable, and if any of the CPO values are found to be so, they should be corrected.

The CPO is measured at each observation of the data set, and thus we compare models based on [48]

$$-\text{mean}\big(\ln CPO_i\big).$$

Comparing the negative mean of the logarithmic CPO's, we once again favour the model with the lowest measure.

3.8.4 Marginal Likelihood

The final measure used to select the best model is the marginal likelihood. Finding the marginal likelihood requires solving an analytically intractable integral, so a numerical approximation is used instead. The INLA approximation is similar to the expression in Equation (3.4), and is given by

$$\hat{\pi}(\boldsymbol{y}) = \int \frac{\pi(\boldsymbol{\theta}, \boldsymbol{x}, \boldsymbol{y})}{\hat{\pi}_G(\boldsymbol{x}|\boldsymbol{y}, \boldsymbol{\theta})}\bigg|_{\boldsymbol{x}=\boldsymbol{x}^*(\boldsymbol{\theta})} d\boldsymbol{\theta}$$

where $\hat{\pi}_G(\cdot)$ and the node $\boldsymbol{x}^*(\boldsymbol{\theta})$ is the same as in Equation (3.4). As for the other model selection criteria, we favour the model with the smallest marginal likelihood value.

Chapter 4

Financial Risk and Risk Management

Financial risk is a relevant topic for both businesses and governments, as well as for both professionals and individuals in the market. The term *financial risk* typically refers to the possibility of a loss of money, but it is also used in context with the possibility of default for companies. Many people associate the term with something negative, and even though it is not negative by definition, it is common to think that a risk will have significantly negative consequences. Regardless of ones understanding of financial risk, being able to detect and quantify it is important. Then the needed actions can be done in time, so that the negative outcome will not be fatal. Some forms of risk can be eliminated by diversification, while other forms are inevitable. Financial risk can be classified into the following groups [49]:

- **Market risk:** Also known as systematic risk, market risk originates from the financial markets. This type of risk cannot be eliminated, and examples of market risk sources from the last two decades are outbreaks of pandemics, market recessions and changes in policy rates.

- **Credit risk:** The risk of a counterparty not being able to carry out its obligations previously agreed on when entering a contract. The simplest example is a borrower not being able to repay his debt to the lender, which explains why your loan will be more expensive if you are considered as a riskier borrower.

- **Liquidity risk:** The risk of not being able to repay (short-term) debt obligations because of an insufficient liquidity. For instance, a real estate company may have valuable properties and carry most of its value in these. Properties are relatively illiquid, so there is a risk that the company will not be able to fulfil its obligations if the properties do not sell, of if they sell at a lower price than listed.

- **Operational risk:** The risk caused by organizational flaws and failures, either related to the management, the employees of the company or to other failures such as technical system ones. Human caused mistakes will almost always fall into this category of risk.

- **Legal risk:** Risks caused by legal constraints and uncertainty. For instance, misunderstanding and thus failing to meet the law can give extra costs through fines and fees. Legal risk was initially seen as a part of operational risk, but over the years the importance of the category has increased. Today this risk is thus typically distinguished from operational risk.

The risk category of interest in this thesis is market risk, or systematic risk. Framework and regulations on systematic risk is presented in the final section of this chapter, Section 4.2. First, we present some risk measures and some desired properties of these.

4.1 Risk measures

Being able to identify, analyse and quantify risk is crucial for any investor or financial company. A high risk is often associated with a higher return but can also lead to a massive loss. A company or an investor will usually have a preferred tolerance for risk, which together with the risk measure is used to decide whether to make a certain investment or enter a contract. Risk measures can be relative to some benchmark, such as the measures alpha and beta, or they can tell something about how large expected losses might be. The latter is often used by investors to make decisions. To put it bluntly, being better than some benchmark does not help if the losses are so great that you go bankrupt. Measures related to the size of expected losses will thus be the focus of this thesis. Before we present two such measures, we present the concept of risk coherency.

4.1.1 Coherent Risk

Let Ω be a finite set of states equivalent to all outcomes of a portfolio, and let $\mathcal{G}$ be the set of all risks on Ω. Denote by $\rho(X)$ the risk measure for a risk $X \in \mathcal{G}$.

A risk measure is coherent if it satisfies the following axioms [50]:

1. Translation invariance: $\forall X \in \mathcal{G}$ and for all constant $\alpha \in \mathbb{R}$, $\rho(X + \alpha) = \rho(X) - \alpha$; that is, adding an amount α of cash or another riskless asset to the portfolio reduces the risk by the same amount. Conversely, subtracting a constant amount from the portfolio increases the risk by this amount.

2. Subadditivity: $\forall X_1, X_2 \in \mathcal{G}, \rho(X_1 + X_2) \leq \rho(X_1) + \rho(X_2)$; the risk cannot increase when we combine two portfolios. We will later see that a well-known risk measure fails to meet this property.

3. Positive homogeneity: $\forall \lambda \geq 0$ and $\forall X \in \mathcal{G}, \rho(\lambda X) = \lambda \rho(X)$. Increasing the portfolio size by a factor λ will increase the risk by the same factor.

4. Monotonicity: $\forall X, Y \in \mathcal{G}, X \leq Y, \rho(Y) \leq \rho(X)$. If a portfolio Y always performs better than another portfolio X, then the latter will have the highest risk.

The four coherency properties are logical and sensible considering the common understanding of risk as a concept. If, for example, increasing a portfolio size by a factor $\lambda \geq 0$ did not increase the risk by the same factor, but instead with a factor smaller than λ, investors would be able to take on much larger positions without adding too much risk. Then multiplying the size of the portfolio would result in less risk, and a portfolio size converging to infinity would have a risk converging towards zero. This is an obvious contradiction of any reasonable concept of risk and emphasises why we want the coherent risk axioms to be satisfied. However, not all risk measures satisfy all the axioms; one of the most popular ones does not satisfy the subadditivity axiom in general. This measure is value at risk, which is presented in the following section.

4.1.2 Value at Risk

Value at risk, or VaR, is a quantified measure of potential losses for investments or businesses in a specified period. It is used to measure and control the risk of portfolios, and the metric states what the potential loss is and what the probability is for this loss to happen. VaR can be applied to a long list of financial instruments and derivatives of those, e.g., stocks, bonds, options, swaps and many more. The concept is easy to understand and simple to calculate as long as certain properties of the underlying are known, which is why it is broadly used in the financial industry. VaR is usually calculated in one of three ways; by historical simulation, the variance-covariance method or by Monte Carlo simulation, where the first is the most used one.

More precisely, VaR is defined as the smallest value l such that losses greater than l have probability L of not exceeding $(1 - \alpha)$, where $\alpha \in (0, 1)$ is a chosen significance level, typically 5% or smaller. In mathematical notation,

$$\mathrm{VaR}_\alpha(X) = \inf\{l \in \mathbb{R} : P(L > l) \leq 1 - \alpha\} = \inf\{l \in \mathbb{R} : F_L(l) \geq \alpha\}$$

where $F_L(l) = P(L \leq l)$ is the cumulative distribution function for the losses.

In other words, VaR is simply the lower α quantile of the loss distribution. This leads to some serious disadvantages of using the method; VaR does not indicate anything about the size of the loss if it exceeds the confidence level. Loss distributions for investments and portfolios can often have heavy tails and even quite many outliers, which are not accounted for in VaR.

In general, VaR is not a coherent risk measure, since it fails to satisfy the subadditivity axiom. That is, the VaR may increase when combining two portfolios, contrary to the diversification principle. Because of this, an investor or a company basing their decisions on VaR can incorrectly appraise the maximal expected loss of a portfolio. This may cause them to not hedge an investment that should ideally have been hedged, or to make another suboptimal decision. To handle the issues of VaR, the measure expected shortfall (ES) was introduced by Artzner et al.[50] in the late 1990's.

As previously mentioned, VaR can be measured in three ways. The historical method and Monte Carlo simulation both fall under the non-parametric calculation. Once the data is available, either from simulation of from some historical data, the α level 1-day VaR is simply the lower α quantile of the profit-and-loss distribution.

The parametric calculation of VaR is also called the variance-covariance method. This method assumes that the data follows a Gaussian distribution, so that the standard normal distribution curve can be used to find the VaR. The 1-day VaR is given by

$$\mathrm{VaR}_\alpha = \mu + \sigma\Phi^{-1}(\alpha), \tag{4.1}$$

where μ is the mean and σ is the standard deviation of the return series, and $\Phi(\cdot)$ is the cumulative distribution function of $\mathcal{N}(0,1)$ [51]. The same method also works for a Student-t distributed profit-and-loss distribution. The VaR is then the same as shown in Equation (4.1), but with the standard normal cumulative distribution function replaced by the Student-t version, with $\nu > 2$ degrees of freedom. To obtain the n-day VaR, we simply multiply the 1-day VaR measurement by $\sqrt{n}$. If the data follows, or is close to following, a Gaussian distribution, the parametric method is suitable and easy to implement. However, financial time series are often more likely to follow a lognormal distribution. If so, the non-parametric method is preferred over the parametric method, as the latter will underestimate the risk to a potentially critical degree.

4.1.3 Expected Shortfall

ES, also known as conditional VaR (CVaR), quantifies the amount of risk in the tail by finding the weighted mean of the extreme losses in the returns, i.e., the return losses below the VaR cut-off. Mathematically, ES is defined as

$$\mathrm{ES}_\alpha = \frac{1}{1-\alpha}\int_\alpha^1 q_u(F_L)du = \frac{1}{1-\alpha}\int_\alpha^1 \mathrm{VaR}_u(L)du.$$

Unlike VaR, ES is a coherent risk measure as it satisfies all four axioms. The proof of this statement is not trivial, and a mathematical background is strongly recommended to understand the technicalities. For multiple ways to prove the subadditivity of ES, the reader is referred to Embrechts and Wang [52].

When we have calculated the VaR for some financial data, finding the ES is not too hard. For the non-parametric method, ES is the average of all observations exceeding the α level VaR. The parametric ES is similar to the parametric VaR given in Equation (4.1), with some modifications;

$$\mathrm{ES}_\alpha = \mu + \sigma\frac{\phi\big(\Phi^{-1}(\alpha)\big)}{1-\alpha},$$

where $\phi(\cdot)$ denotes the probability density function of $\mathcal{N}(0,1)$ and the other parameters are as in the VaR definition [51].

The concepts of VaR and ES are visualised in Figure 4.1. Here we observe how for the normal distribution, VaR would not outperform ES as there is not much tail risk present. However, in the mixed distribution displayed in Figure 4.1c, there is much more risk in the tail. Such distributions, or even more extreme ones, may occur in the real world. This shows the importance of measuring the ES.

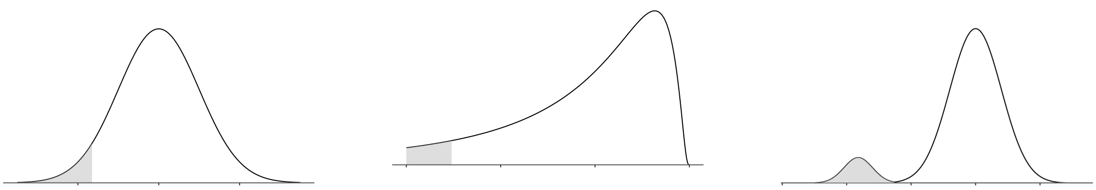

(a) Normal distribution (b) Lognormal distribution (c) Mixed distribution

Figure 4.1: A visualisation of VaR and ES for three different distributions. Note that the lognormal distribution is reversed to replicate a typical loss distribution. The shaded area shows the lower 5% quantile. VaR is the cut-off between the shaded and non-shaded areas, while ES is the expected value of the shaded area. The mixed distribution clearly depicts a situation where VaR would give an unrealistic indication of the risk.

4.2 Risk Management and Regulations

All banks, investment funds, financial institutions and even individuals participating in financial markets should be familiar with risk management. While banks and institutions are required to implement risk management by the law, investors and individuals should have a professional or personal interest of managing their risk to some degree. For instance, when deciding which investment to add to your portfolio, you would choose the one with the lowest risk for your desired expected return. The volatility of an asset is a simple, yet relatively informative measure of risk for an individual in the stock market. If you want a more informative risk measure quantifying the profits and losses of your investments, VaR and ES are good choices of risk metrics.

The Basel Committee on Banking Supervision (BCBS) under the Bank for International Settlements (BIS) provides the Basel frameworks for banks worldwide. BIS was established in 1930 as the first financial institution in history, to deal with the economic aftermath of the First World War; in particular regarding the payments Germany was required to do by the Treaty of Versailles. Half a century later, the global economy had seen a stagflation, and multiple countries requested a framework on how to ensure financial stability. The Basel Capital Accord, today known as Basel I, was issued in 1988 as a response to this request, and has been refined twice since; by Basel II in 2004 and Basel III in 2017 [53]. BIS currently has 63 members, of which all are either central banks or monetary authorities. Examples include the central banks of Norway, United States, Japan, EU, England, and many more [54].

The decisions of the BCBS are not legally required to follow for the members of BIS but are rather implemented by each central bank or authority for their countries. For instance, in the European Union and the European Economic Area, the Capital Requirements Directive IV (CRD) and the Capital Requirements Regulation (CRR) legislations implement the Basel III framework for banks and other financial institutions [55]. Since 2019, the implemented framework has been a revised version of Basel III. The current edition was BCBS' response to the financial crisis of 2008. During this crisis, banks suffered from great economic losses. The need for a more precise and comprehensive framework became apparent, as the pillars of the prevailing Basel II capital framework came short. In particular, topics such as liquidity, regulatory capital and risk management were not sufficiently well covered by the three pillars of Basel II [56].

Market risk measurement requirements are some of the most important differences of the current and previous versions. The Basel Committee highlights five main drivers of market risk; commodity prices, equity prices, interest rates, credit spreads and foreign exchange [57]. In terms of quantitative standards, the following were prevailing in Basel II:

1. VaR had to be computed every day.
2. The 10-day 99% VaR was to be used. The banks could either compute the 10-day VaR directly, or scale it up from a shorter holding period.
3. The historical data period used to calculate the VaR had to be at least one year.
4. All the three methods of measuring the VaR mentioned in Section 4.1.2 were allowed.

This list does not show the complete list of quantitative standards in the Basel II framework; the reader is referred to BCBS [58] for the full version.

The financial crisis revealed the insufficiency of VaR in terms of extreme losses, even though stress tests had to be carried out. Thus, ES was introduced in Basel III, replacing both VaR and stressed VaR. The full list of requirements of how to calculate ES in the new framework can be seen in BCBS [59]. A short summary of the most relevant requirements for this thesis are given here:

1. ES must be calculated on a daily basis.
2. The 97.5% quantile ES must be used.
3. ES must be calculated based on the worst-case scenario 12 months available in the data set. This data set must, at least, date back to 2007.
4. No specific method of how to measure the ES is determined - both parametric and non-parametric methods are fine.

The remaining of the Basel III capital requirements calculations concerns calibration of stress periods, model and risk factor standards and that the banks should be able to distinguish the different risk classifications as well as calculating the correct correlation between these.

The capital requirements discussed above only make up a small part of the huge Basel III framework; even within market risk, the framework is much more comprehensive than seen here. Besides the risk measurement, BCBS brings up three key characteristics of the updated framework [57]:

- On the boundary of the trading and banking books, regulatory requirements are specified to prevent banks and institutions from moving capital within the institution to meet the capital requirements.
- The internal models of banks are tested by improved criteria which are better at finding inappropriate models.
- The risk measurement requirements for options, foreign exchange and similar instruments are improved in terms of how the risk is weighted.

Chapter 5

Data

The data set used in this thesis is provided by DNB and represents NIBOR, Norwegian Interbank Offered Rate. The data includes maturities ranging from 3 months to 15 years, covering the time period from 3 January 2005 to 30 May 2022. In addition, the Norwegian policy rate is included to observe how the interest rates change compared to the policy rate. This data is downloaded from the open API of Norges Bank, the central bank of Norway [14].

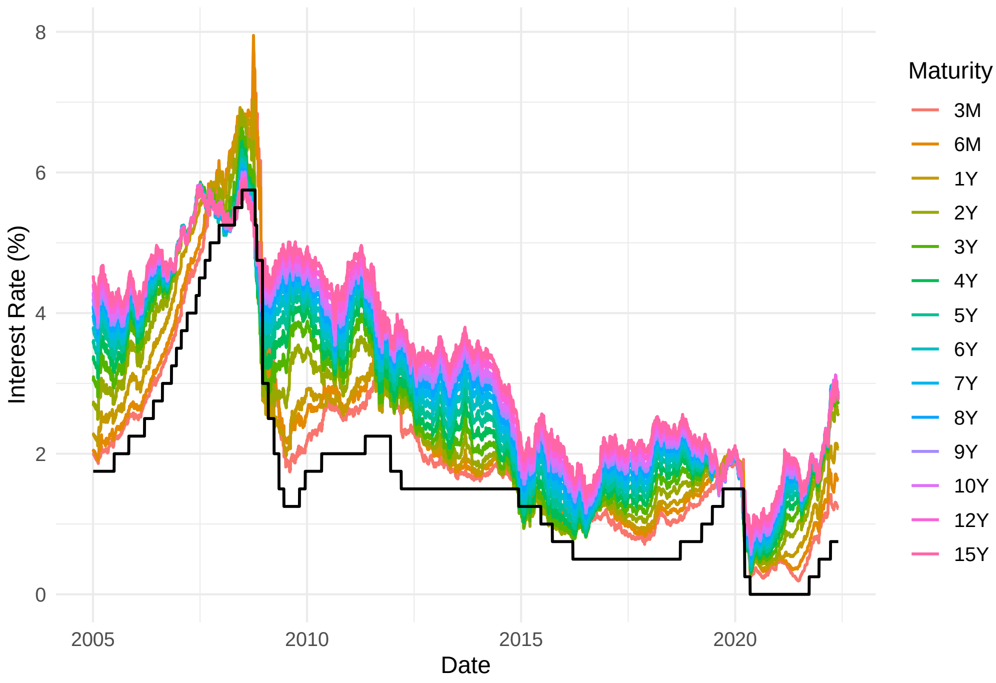

Figure 5.1: NIBOR rates for the period January 2005 to May 2022, with the Norwegian policy rate for the same period displayed by the black line.

The NIBOR rates are displayed in Figure 5.1, where the Norwegian policy rate is represented by the black line. There are certain patterns in the data set; during financially stable times, the yield curve takes an inclining slope, while the opposite is true during financial reclines. We recognise two main recessions in the data set. The first abrupt fall of the interest and policy rates is the financial crisis of 2007-2009, while the second large fall is the outbreak of Covid-19 in March 2020. From the financial crisis to 2020, there is a downward trend in the interest rate levels, as a result of the strategy of the central bank decisions both in Norway and on an international level. In 2021, and especially in 2022, this trend seems to take a turn, into a growing trend over multiple quarters. In particular, we note the difference in the NIBOR rates when the policy rate was increased in 2018-2019 versus the most recent series of policy rate increases. In 2018-2019 the yield curve flattened, while the most distinct movement of the yield curve during the most recent changes in the policy rate seems to be the increased level.

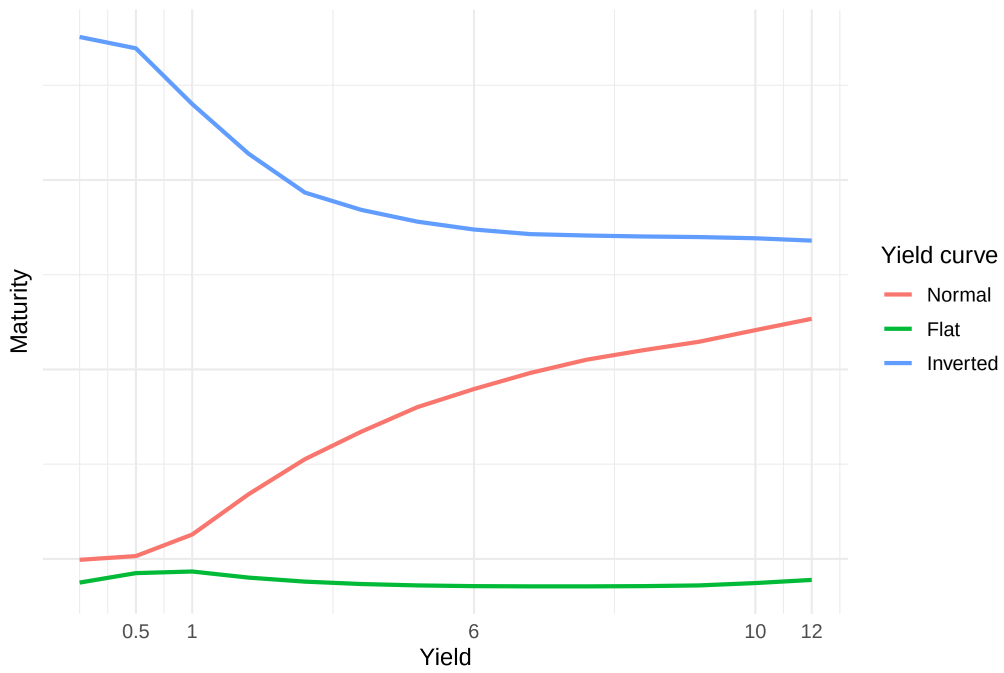

Figure 5.2: Three categories of yield curves observed in the data set; the normal, flat and inverted yield curves. Note that the x-axis is scaled to show changes between short maturities better.

Figure 5.2 displays the yield curves of three different dates found in our data set. The normal curve is from 3 January 2005, the flat curve is from 21 February 2020 and the inverted curve is from 26 September 2008; these dates are chosen to represent the characteristics clearly, as each category has its implications [60]. In the normal curve, represented by the red line, longer maturities have higher yields. This is the yield curve we observe most of the time, and it is an indicator of healthy market conditions. The steeper the normal yield curve, the higher is the expected future growth.

In uncertain economic times, a flat or humped yield curve is a common observation. The difference between a flat and humped yield curve is that in the latter, middle maturities - typically ranging from 6 months to 2 years - have a higher yield than the short and long rates, which we can see some hints of in the green line. This can for instance happen because of an increased volatility [60]. Finally, if the yield curve is inverted as displayed by the blue line, we are in one of two situations; either there is a recession going on, or there is a strong indication that a recession is coming up in the near future. In the case of this example, taken from September 2008, there was an ongoing recession.

5.1 Preprocessing

The interest rate data set is incomplete in the sense that not all series have the same length. While the majority of the series have dates equivalent to all business days in the given time period, the dates of the first three and latter two series do not match the rest. Certain problems arise from this situation; for example, we cannot perform the desired matrix calculations to our data set. To deal with these problems, we preprocess the data before performing the analysis.

The complete series consist of 4542 observation pairs on the form (Date, Value) for each maturity. However, not all the series are complete - for instance the shortest ones, the 3- and 6-month interest rates, have 160 less observations, i.e., about 3.5% of the data are missing compared to the longest series. The missing observations are distributed somewhat uniformly in their respective series, and there are no long periods where the data is missing. Since interest rates have a fair chance of remaining at the same level on consecutive days, expanding the shorter series by adding the initially missing dates seems like a fairly good solution to the problem. The series which after the expansion have missing values for the added dates are then completed by inserting the interest rate value from the previous date. All series initially have data for the first dates in the data set. Thus, to avoid problems if two or more consecutive dates miss data, the function is implemented in a forward loop so there is always a previous value available. Another method would be to cut down the length of the longest series to the length of the shortest series. This must be used for instance if the method described here results in unstable decomposed matrices.

The means and standard deviations of the raw and processed versions of the relevant NIBOR series are given in Table 5.1. There are no drastic changes in any of the parameters for any of the series. Except from the 0.011 percentage points difference in the 1-year mean of the raw and processed series, there are no differences bigger than 0.003 percentage points. Furthermore, the standard deviations of the raw series are relatively large, and so the processed data series are contained in the uncertainty intervals.

	Raw		Processed	
	Mean	Sd.	Mean	Sd.
3 months	2.158	1.498	2.156	1.494
6 months	2.294	1.506	2.292	1.502
1 year	2.316	1.501	2.327	1.486
12 years	3.248	1.355	3.251	1.356
15 years	3.313	1.348	3.316	1.349

Table 5.1: Means and standard deviations for selected raw and preprocessed NIBOR rates, given in percent.

5.2 Descriptive Statistics

The NIBOR swap rates are transformed to forward rates by Equation (2.6) in Section 2.4 to be suitable for modelling. The forward rates are plotted in Figure 5.3 together with the Norwegian policy rate, and a summary of the means and standard deviations for the interest rate series is given in Table 5.2. We recognise the same general patterns from the swap rates also in these series, but the spread between the different rates seems to be larger for the forward rates. This results in somewhat different yield curves for the forward rates than for the NIBOR rates.

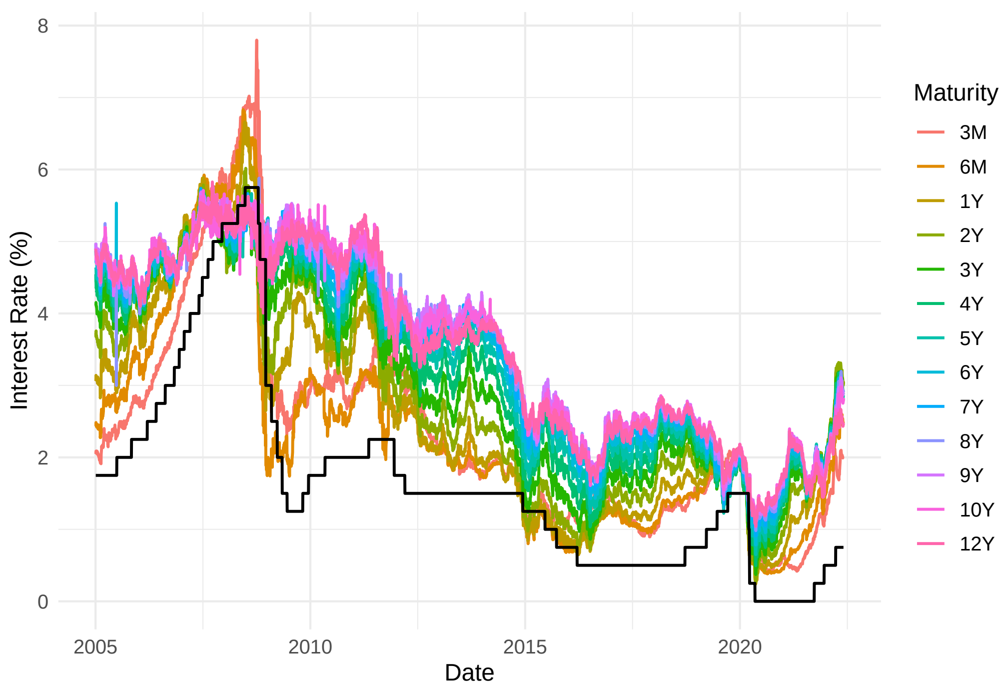

Figure 5.3: Forward rates derived from the NIBOR rates for the period January 2005 to May 2022, with the Norwegian policy rate for the same period displayed by the black line.

The properties of the interest rates in Table 5.2 behave as expected - rates of longer maturities have a higher mean, and a lower standard deviation, than rates of shorter maturities. The mean and standard deviations are almost strictly increasing and decreasing, respectively, for higher maturities, with just a few exceptions.

Maturity	Mean (%)	Sd. (%)
3 months	2.388	1.470
6 months	2.325	1.467
1 year	2.580	1.478
2 years	2.774	1.397
3 years	2.981	1.362
4 years	3.171	1.347
5 years	3.311	1.331
6 years	3.418	1.315
7 years	3.491	1.302
8 years	3.547	1.299
9 years	3.581	1.299
10 years	3.570	1.293
12 years	3.508	1.294

Table 5.2: The means and standard deviations for the Norwegian forward rate series for January 2005-May 2022.

5.3 Portfolio

Risk can be measured directly from the interest rates; however, risk measures of derivatives are often more interesting. Thus, we look at a portfolio of four zero-coupon bonds of different maturities - 3-months, 1-, 5- and 10 years. Bonds are chosen because of their very close relationship to interest rates, which allows us to measure the risk from the interest rate movements with as little noise as possible.

To minimize the risk in such a portfolio, the bonds with longer maturities would be weighted heavier than the shorter maturity ones, because of their less volatile nature. In addition, a dynamic weighting would be preferred. However, we choose to base our modelling on a uniformly weighted portfolio, such that the more volatile nature of the short rates will be better reflected in the measured risk. We are interested in investigating the changes in the bond prices rather than the bond prices themselves, as this will translate to the profit and losses when finding the risk. The logarithmic scales of these changes are used, as often done in quantitative finance.

The bond prices in the portfolio as well as the corresponding forward rates are displayed in Figure 5.4. Here, the left-hand panel shows the prices, while the right-hand panel shows the interest rates. In the former plot, the price of the weighted sum is represented by the black line, located just above the 5-year bond price. We observe that this price reflects properties of all its components, as desired. The separate bond prices behave as expected - bonds of longer maturities entail more risk, and so the price of such bonds will be lower. Also, the prices decrease during more turbulent periods and recessions, when the future expectations are worse. Finally, the different bonds seem to be correlated to a high degree, especially those of similar maturities.

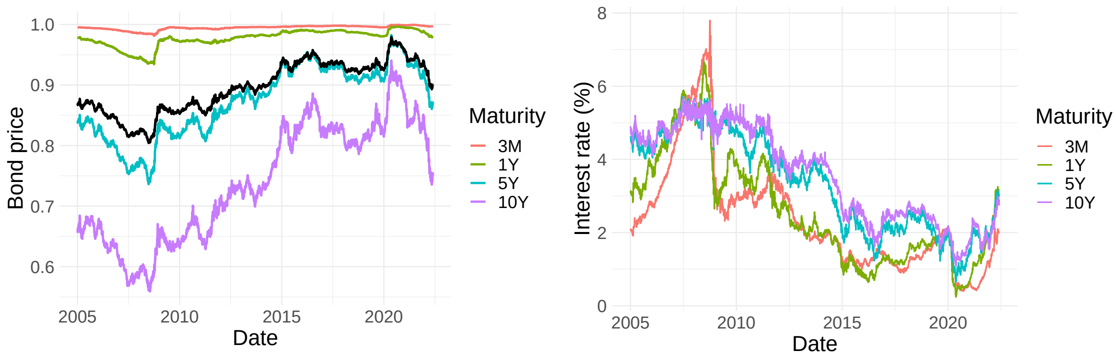

Figure 5.4: The prices of the different bonds in the portfolio and the corresponding forward rates. The black line in the price plot represents the weighted sum of the portfolio.

Furthermore, Figure 5.5 displays the logarithmic daily changes of the portfolio as one unit, and the histogram of these changes. The black vertical line in the histogram shows the mean and is very close to zero. Comparing the daily changes of the portfolio to the swap rates in Figure 5.3, we can confirm that the largest changes of the portfolio do indeed occur when the swap rates are prone to the most drastic changes.

The correlation plot displayed in Figure 5.6 shows the correlations between the daily changes of the bonds in our portfolio. All the correlations are found to be significant at the 0.1% level, i.e., highly significant. The highest correlation is the one of the 5- and 10-year maturity bonds, with a coefficient of 0.912. Furthermore, both the 3-month to 1-year and 1-year to 5-year correlations are larger than 0.5, confirming the apparently high correlation we observed in Figure 5.4.

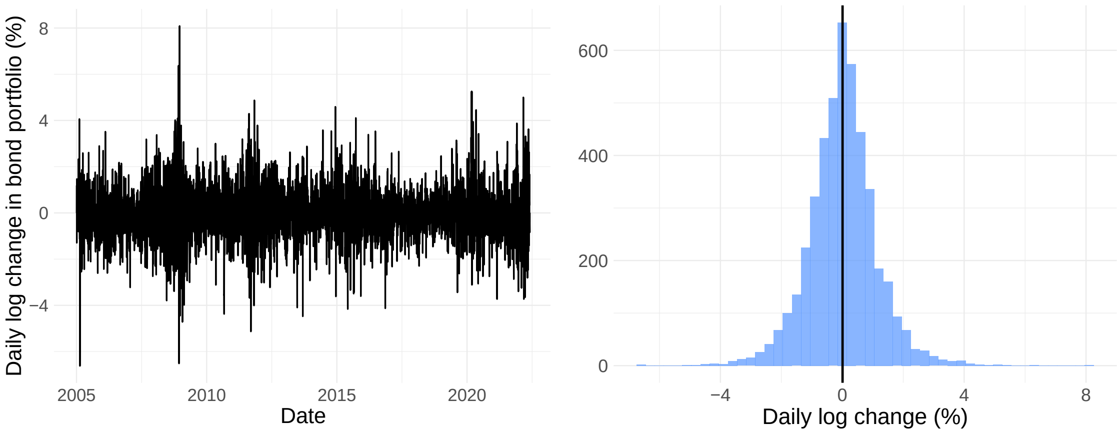

Figure 5.5: The logarithmic daily changes of the portfolio and the corresponding histogram. The black line in the histogram represents the mean of the changes.

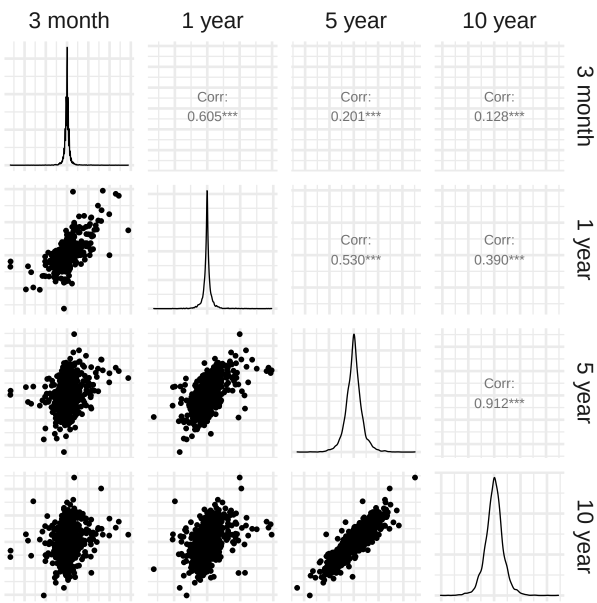

Figure 5.6: Correlation between the daily changes in the portfolio bonds.

5.4 Preliminary PCA Analysis

We apply PCA to the daily changes of the forward rates to analyse the data more thoroughly in a preparatory manner. In particular, we compute the eigenvalues and eigenvectors of the correlation matrix in order to find the volatility factors given by Equation (3.8). These volatility factors are then fitted by polynomials. Applying this process to the entire period would not give very informative results, so instead, we split the period into four different subperiods which represent different characteristics. All periods, except from the last one, are chosen to be approximately two years long to follow the standard convention for financial time series modelling. The dates of the periods are given in Table 5.3.

Period	Dates
Financial crisis	13 June 2008 - 10 May 2010
Normal conditions	1 July 2016 - 1 June 2018
Covid-19 outbreak	26 July 2019 - 5 May 2021
Covid-19 recovery	5 May 2021 - 31 May 2022

Table 5.3: Calibration periods.

Our forward rate data set consists of 13 rates of different maturities, and hence the covariance matrix will be of dimension (13×13) and we will have 13 eigenvalues and eigenvectors. After performing the PCA, we can find the proportion of variance explained by each of the components. The cumulative proportions of variance explained for the first five eigenvalues of the four different subperiods are given in Table 5.4. The faster these proportions converge to 100%, the more correlated is the corresponding data set. Here, we observe that the Covid-19 outbreak and Covid-19 recovery periods have the most prominent correlation structures, while the data indicates weaker correlations in the normal condition and financial crisis periods. We return to Figure 5.3 for a simple explanation; the yield curves are significantly less consistent in the weaker correlation periods than in the outbreak and recovery of the Covid-19 pandemic.

Eigenvalue	Financial crisis	Normal conditions	Covid outbreak	Covid recovery
λ_1	24.3	51.6	73.4	66.8
λ_2	44.3	63.8	86.3	75.6
λ_3	57.2	71.9	90.2	82.7
λ_4	68.3	78.8	93.4	88.0
λ_5	76.9	84.9	96.1	91.8

Table 5.4: Cumulative proportion of variance explained for the first five eigenvalues of the four subperiods, given in percent.

The discrete fitted volatilities are displayed in Figures 5.7 and 5.8. Just like the yield curves we presented in Chapter 5, different shapes of the volatilities have different implications; however, the interpretation here varies slightly from the yield curve interpretation. The first volatility factors are all fitted by straight, constant lines, corresponding to a parallel shift movement being the most important characteristic of the model. That is, from one time to another, the interest rates of different maturities will change by the same factor.

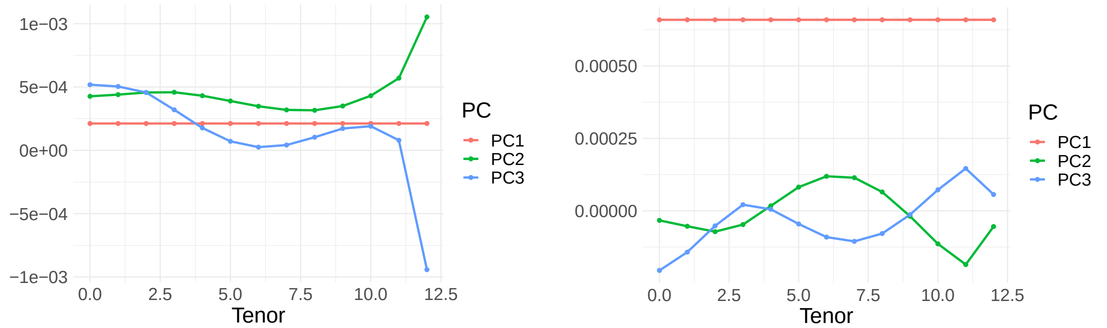

(a) 13 June 2008 to 10 May 2010 (b) 1 July 2016 to 1 June 2018

Figure 5.7: The discrete fitted volatilities of two calibration periods: the financial crisis and normal market conditions.

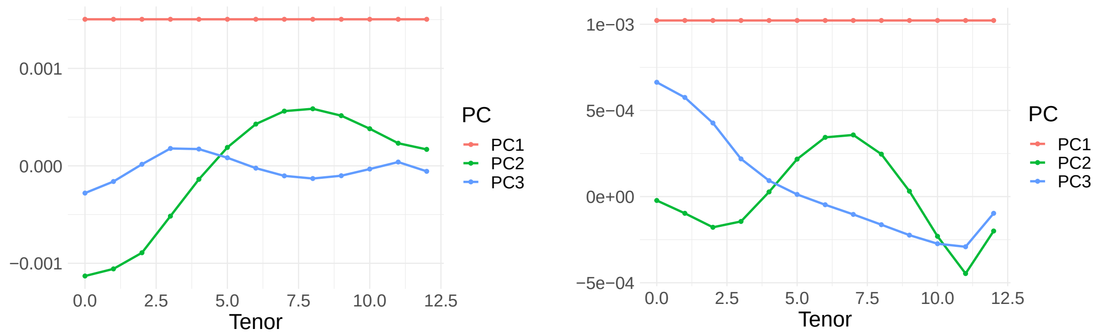

(a) 26 July 2019 to 5 May 2021 (b) 5 May 2021 to 31 May 2022

Figure 5.8: The discrete fitted volatilities of two calibration periods: the Covid-19 outbreak and the past year of an increased policy rate.

The second volatility components of Figure 5.7a and Figure 5.8a are approximately two different linear movements, with one change of direction over the tenors. This corresponds to a twist in the yield curve, a phenomenon where the short and long rates change their relative position. We can recognise these movements with ease in the forward rate curve displayed in Figure 5.3; the twist is particularly visible during the financial crisis, when the 3 month rate is almost 8% in the fall of 2008. The second volatility components of Figure 5.7b and Figure 5.8b are not as easy to interpret, and likewise, they do not indicate as much about the yield curve.

Looking at the forward rate plot in Figure 5.3 once again, we can see why; the calibration periods simply do not contain any yield curve inversions in the two right-hand panel calibration periods.

The third volatility components are related to the curvature of the yield curve. In the financial crisis and the Covid-19 recovery periods we see a declining trend from short to long maturities, while the two other calibration periods have less consistent trends along the tenors [61].

Chapter 6

Implementation

The aim of this thesis is to investigate the effect of changes of the Norwegian policy rate on interest rate risk and we thus want a model from which we can easily measure the risk. The model should be sensitive to the changes in the financial time series we are modelling, while not being too complex in terms of implementation.

The chosen model is a stochastic volatility model with an autoregressive model of order 1 (AR(1)) as the latent model, which can easily be implemented in INLA. The basic version of this model was presented by Taylor in 1982 [62], and is given by

$$r_t = \exp\left\{\frac{\eta_t}{2}\right\}\epsilon_t \tag{6.1}$$

$$\eta_t = \mu + \phi(\eta_{t-1} - \mu) + \sigma h_t \tag{6.2}$$

where r_t is the return on day t, η_t is the logarithmic variance on day t and $\epsilon_t, h_t \sim \mathcal{N}(0,1)$. Furthermore, η_t is an AR(1) process with three parameters; mean μ, persistence $\phi \in (0,1)$ and Gaussian noise σ [62].

In our case, we rewrite Equations (6.1) and (6.2) as the following for simplicity:

$$\begin{aligned} r_t|\eta_t &\sim \mathcal{N}(0, \exp\{\eta_t\}) \\ \eta_t &= \mu + f_t \\ f_t|f_{t-1}, \ldots, f_1, \lambda_f, \phi &\sim \mathcal{N}(\phi f_{t-1}, \lambda_f^{-1}) \end{aligned}$$

where $\lambda_f = 1/\sigma^2$ is the inverse volatility, i.e., the precision, of the AR(1) process given in Equation (6.2). When the common mean μ is assigned a vague Gaussian prior, our latent field is given by

$$\boldsymbol{x} = (\eta_1, \ldots, \eta_n, \mu) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{Q}^{-1}(\boldsymbol{\theta}_1)),$$

where the precision matrix $\boldsymbol{Q}(\boldsymbol{\theta})$ is a sparse matrix satisfying the Markov property, such that $\boldsymbol{x}$ is a GMRF. Note that μ is fixed to zero by default in INLA; so

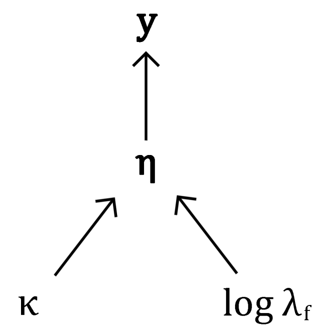

Figure 6.1: A simple figure of the model structure.

is it in our model. When we write the model on this form, $\boldsymbol{\theta}$ consists of two hyperparameters:

$$\boldsymbol{\theta} = (\log \lambda_f, \phi).$$

The structure of this model is shown in Figure 6.1. Here, the persistence parameter ϕ is transformed by

$$\kappa = \text{logit}\left(\frac{\phi + 1}{2}\right).$$

Using this parametrisation, the prior for ϕ will be roughly uniform in $(0, 1)$, which gives a stationary time series. Furthermore, it ensures that the parameters are well-defined on $\mathbb{R}$, to avoid restricting the area where the optimisation is performed. We want both hyperparameters to have vague priors, and assign

$$\log \lambda_f \sim \text{LogGamma}(1, 0.0005)$$
$$\kappa \sim \mathcal{N}(0, (0.0001)^{-1})$$

To test the effect of the policy rate, we add a categorical variable to the latent process. We denote this variable by z_t, and define it in the following way:

$$z_t = \begin{cases} 0 & ; \quad \Delta_{PR,t} = 0 \\ 1 & ; \quad \Delta_{PR,t} = \pm 0.25 \\ 2 & ; \quad |\Delta_{PR,t}| > 0.25 \end{cases}$$

where $\Delta_{PR,t}$ is the change of the policy rate from day $t-1$ to day t in percentage points. The latent process is then modelled by

$$\eta_{t,i} = \mu_{i,z} + f_t, \qquad i = 0, 1, 2.$$

The coding is implemented so that the levels of z are compared to a reference group where the z_t variable is omitted. Hence, we obtain one coefficient of μ for each level of z_t, i.e., $\mu_{0,z}$ is the coefficient if z_t is fixed to be zero, and so on.

Chapter 7

Results and Discussion

We start by finding the 97.5% and 99% VaR and ES using the historical method, measuring the risk based on the 100 last days. These measures are displayed in Figure 7.1 and Figure 7.2 for the 97.5% and 99% risk measures, respectively. Because we are using the historical method, these plots are quite uneven, as the risk can take the same value multiple days in a row with such small quantiles. However, we get an indication of which properties we should expect to see in our modelled risk later. The plots behave like expected; the ES takes higher values than the VaR in both figures, and the 99% risk measures are at a higher level than the 97.5% measures. The most distinct differences occur in periods we know are more volatile. Also, the ES stays at a higher level for a longer time than the VaR. However, the historical method does not picture the risks during the Covid-19 outbreak and recovery as well as we would expect. Because of the used method, the risk observations are lagged compared to the interest rate series. That is, if a multi-day event increasing the risk happens at time t, we will not observe the increased risk clearly before some weeks have passed.

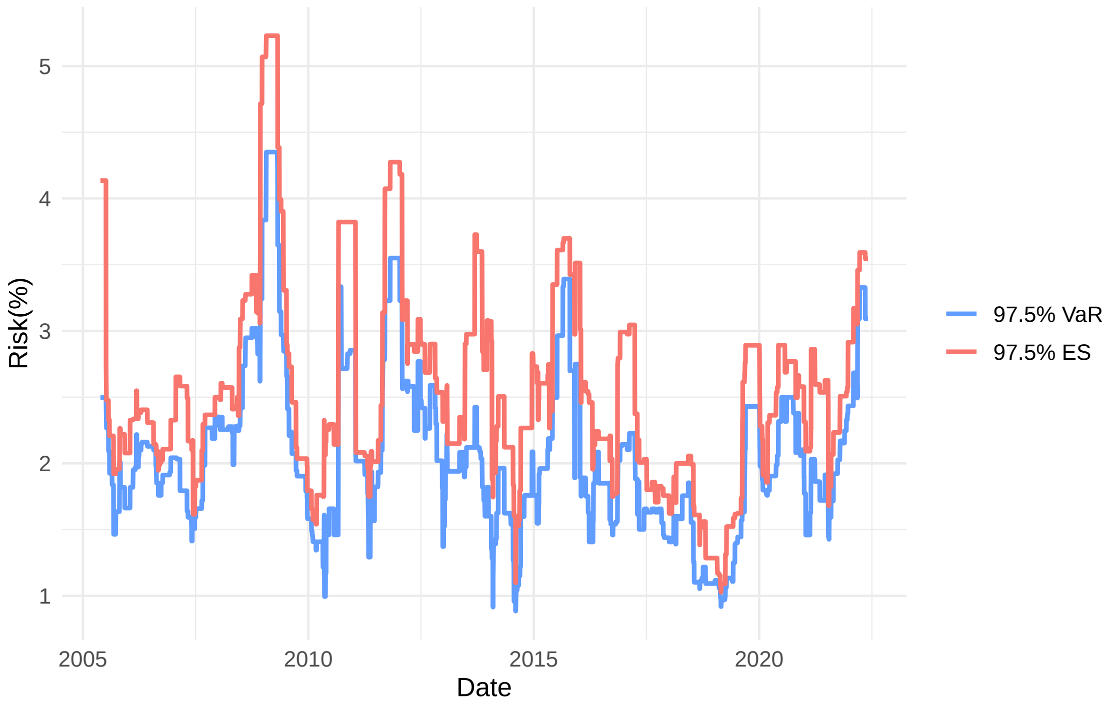

Figure 7.1: The 97.5% VaR and ES of the portfolio measured by the historical method.

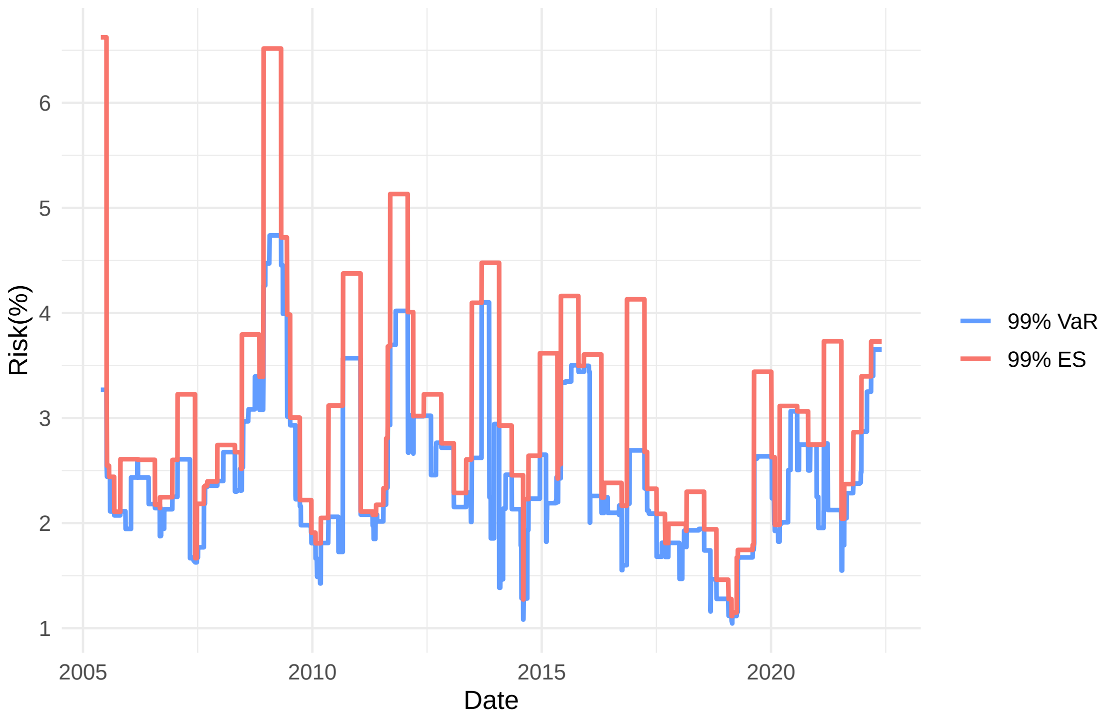

Figure 7.2: The 99% VaR and ES of the portfolio measured by the historical method.

7.1 INLA

We fit the models given in Chapter 6 to the entire data set, from January 2005 to May 2022. The hyperparameters of the two models, without and with the policy rate factor variable, are given in Table 7.1. For both models, the persistence is large and close to 1, with low standard deviations. A high persistence indicates a strong dependence between the daily observations. Thus, shocks in such time series will significantly affect the future predictions. In addition, the precision parameters are similar to each other, both in terms of the mean and the standard deviation.

	ϕ		$\log \lambda_f$	
	Mean	Sd.	Mean	Sd.
Without factor variable	0.999	0.008	2.178	0.306
With factor variable	0.969	0.007	2.184	0.310

Table 7.1: The means and standard deviations for the hyperparameters of the two models.

For the model with the factor variable, each level of μ has its own distribution. The means and standard deviations for these levels are given in Table 7.2. We note how the $z = 0$ group has a significantly lower standard deviation than the two other groups. This is as expected, simply because of the occurrence of each group; there are relatively few occurrences of $z = 1$, and even fewer of $z = 2$.

	Mean	Sd.
$z = 0$	-0.024	0.087
$z = 1$	0.111	0.275
$z = 2$	0.732	0.521

Table 7.2: The means and standard deviations for the different policy rate change groups.

The posterior mean of η for the simple model is displayed in Figure 7.3, with the corresponding 2.5% and 97.5% quantiles. Similarly, Figure 7.4 displays the posterior mean of η for the three factor levels. Here, the quantiles are omitted for better visibility. Plots with quantiles for each factor level are given in Figure A.1 in the appendix. In summary, the posterior mean and quantiles of the $z = 0$ group are similar to the results from the simple model, but with slightly larger quantile intervals. For the $z = 1$ and $z = 2$ groups, the uncertainty caused by the few occurrences is reflected in the much larger quantile intervals. In addition, this model, and in particular the $z = 2$ group, seems to underestimate the ups and downs we expect to see in the posterior mean of η.

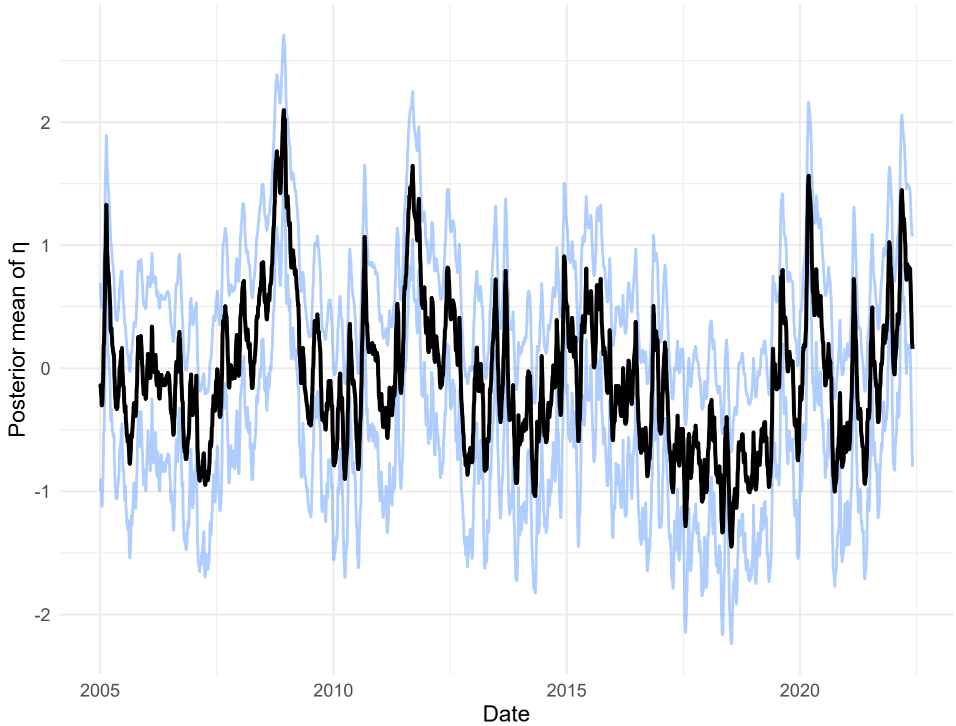

Figure 7.3: The posterior mean of η for the simple model. The blue lines represent the 2.5% and 97.5% quantiles.

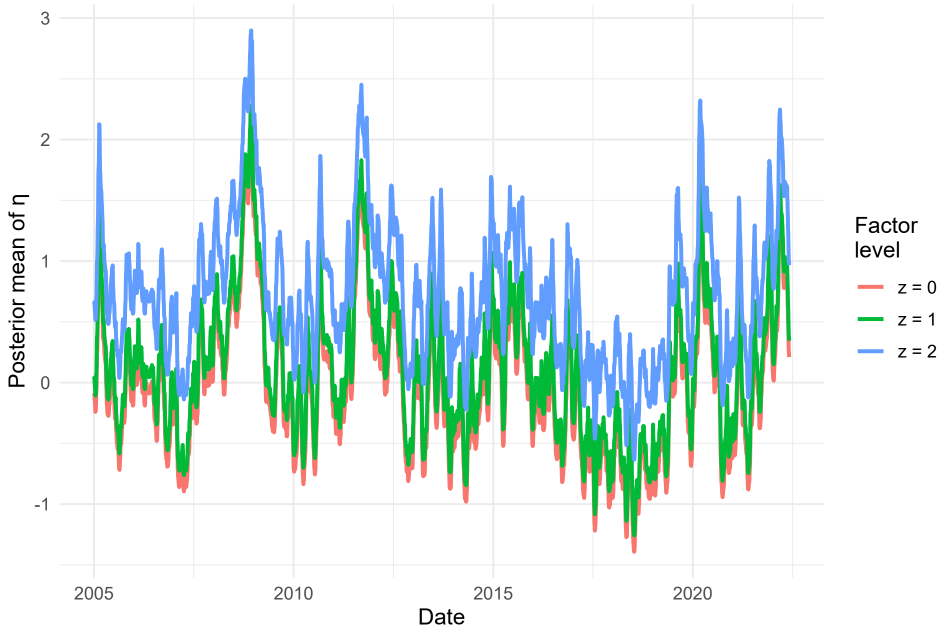

Figure 7.4: The posterior mean of η for the factor variable model. Each line represents a different level of the factor, and the quantiles are omitted from the figure for a neater presentation.

Risk can be measured by a parametric or a non-parametric method, as discussed in Chapter 4. When modelling in INLA, we obtain a distribution for posterior marginal of η for each of the data points. From this, we can find both the VaR and ES. In our case, η is an AR(1) process, which follows a Gaussian distribution. To confirm that this is indeed the case so that we can use the parametric method for measuring the risks, we plot the posterior marginal of η for three different dates in our model data set - the first, one in the middle, and the last. These posterior marginals are displayed in Figure 7.5.

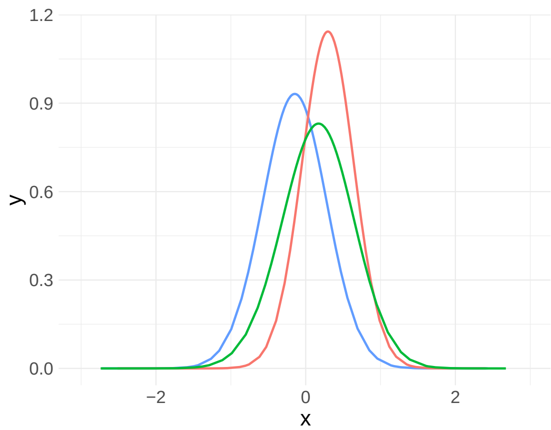

Figure 7.5: The posterior marginal of η for three arbitrarily chosen dates from the data set.

The Gaussian distribution of the posterior marginals of η are confirmed by Figure 7.5, and thus we can use the parametric method for finding the risk measures. The 97.5% and 99% VaR and ES of the simple model are given in Figure 7.6 and Figure 7.7, respectively. The same plots for each level of the factor variable model can be found in the appendix, in Figure A.2. Comparing to the historically measured risk in Figure 7.1 and Figure 7.2, we observe two main differences. Firstly, the INLA models seems to be better at identifying risk in low-rate environments, for instance as seen in 2020 and 2022. Furthermore, because risk is measured for just one date at the time in INLA, we avoid the irregularities and lagging of the historical method.

A summary of the subfigures of Figure A.2 are given in Figure 7.8, which displays the 97.5% ES for the $z = 1$ and $z = 2$ levels. This corresponds to the estimated ES if the policy rate change was kept constant at $\Delta_{PR,t} = \pm 0.25$ or $|\Delta_{PR,t}| > 0.25$, respectively. As expected, these risks are higher than the ones observed for the simple model in Figure 7.6. We observe that the ES for the different levels behave in a similar manner as the simple model, in terms of upward movements around recessions and turbulent periods. However, the factor variable model seems to contain slightly more noise. This is an inherent consequence of the uncertainty caused by the few observations of policy rate changes.

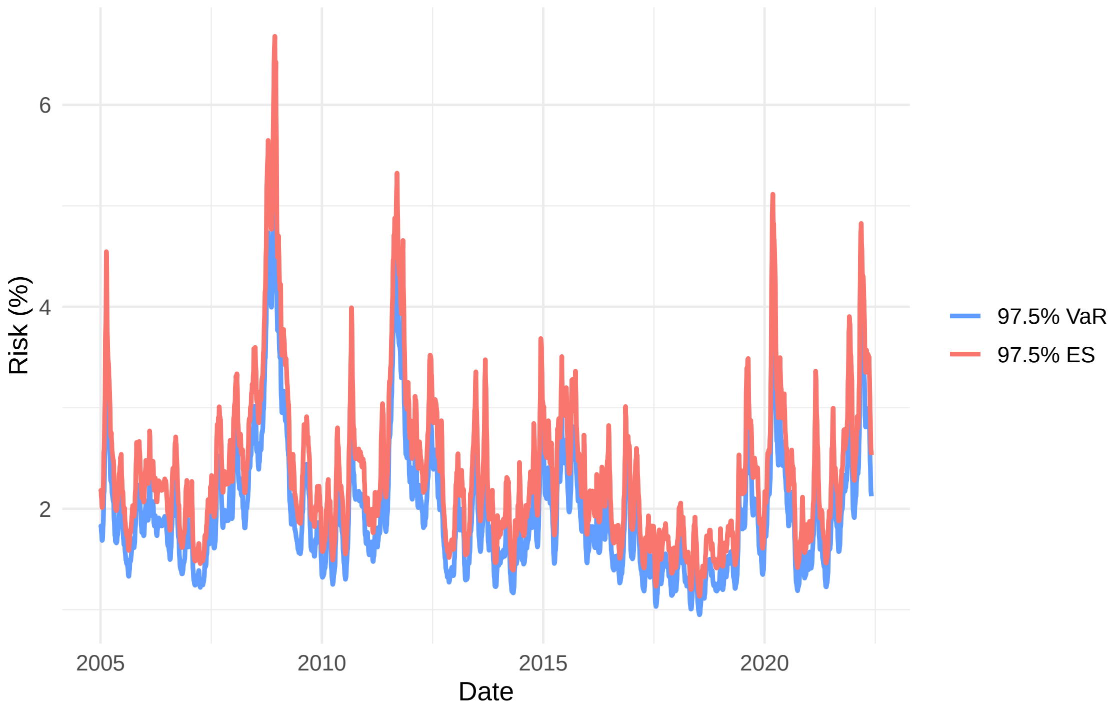

Figure 7.6: The 97.5% VaR and ES for the simple model.

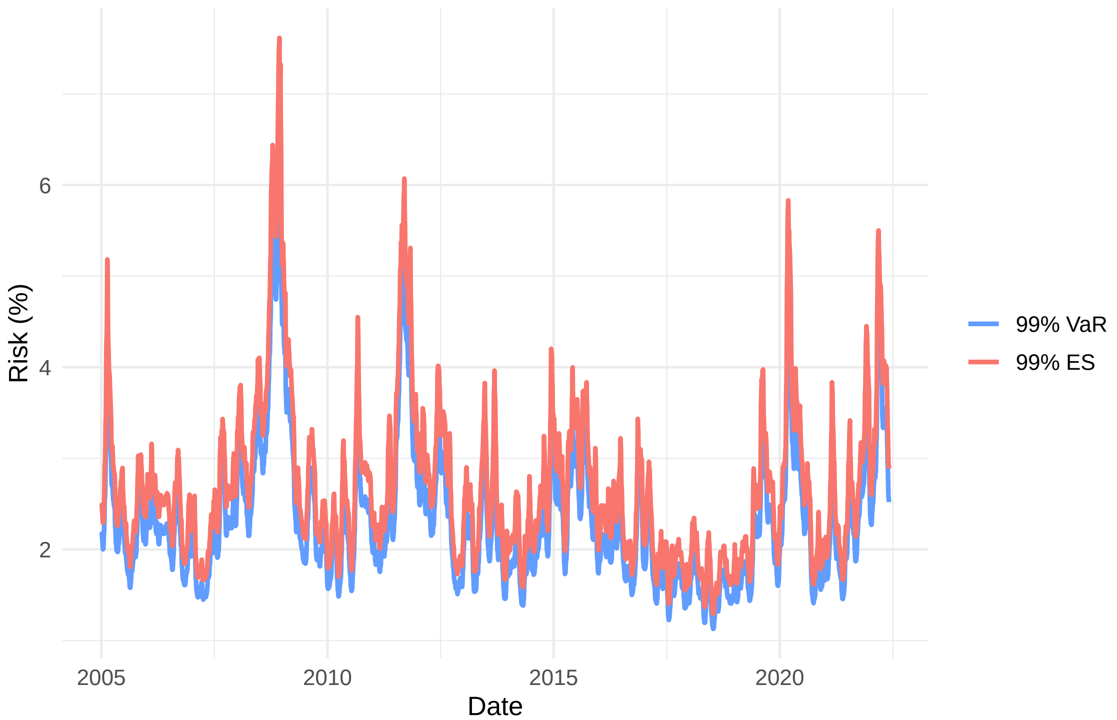

Figure 7.7: The 99% VaR and ES for the simple model.

For both risk levels, the VaR and ES follow each other closely. There are two main reasons for this; the similarity of the formulas used to calculate the risk measures, and the distributions from which we measure the standard deviations. These distributions are Gaussian and have no heavy tails. Hence, the ES and VaR will not be very different.

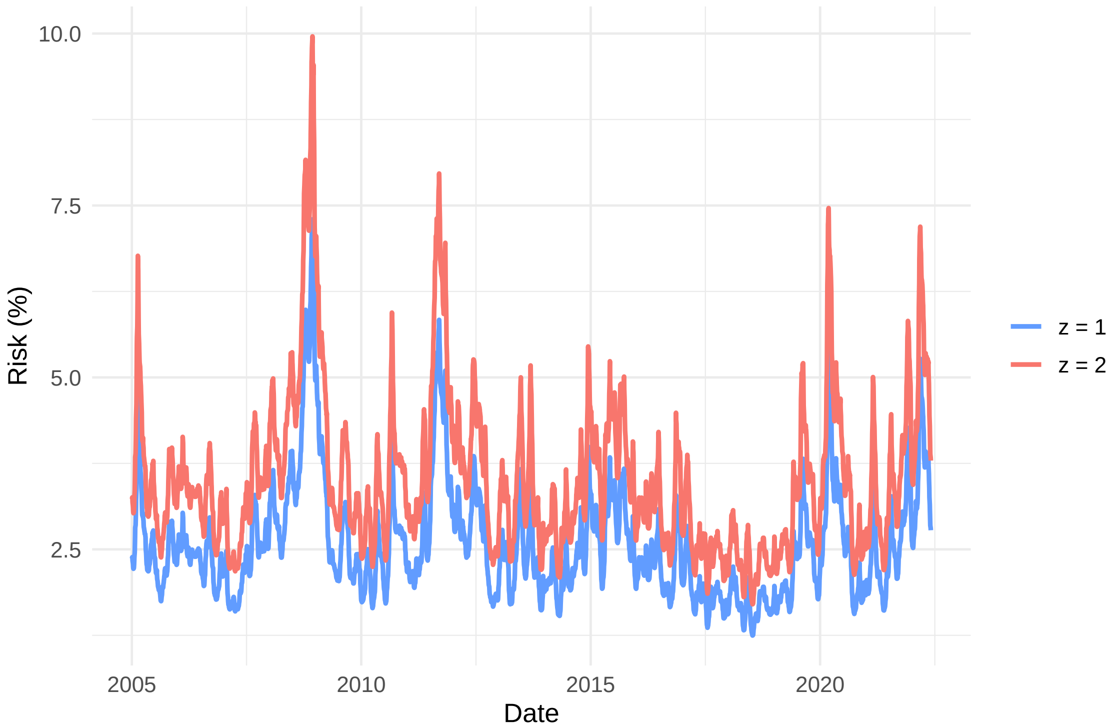

Figure 7.8: The 97.5% ES for the $z = 1$ and $z = 2$ levels of the factor variable model.

To compare the differences of the Basel II and Basel III framework standards, we plot the difference between the 99% VaR and the 97.5% ES of the simple model in Figure 7.9. Here, a negative value corresponds to a larger 97.5% ES than 99% VaR. We observe that the ES is in fact larger than the VaR for all the dates, and that the difference is more prominent during recessions and other turbulent economic periods. This confirms why ES is better at identifying tail risk, and thus why it is - and should be - preferred over VaR.

Moreover, the risk measure plots and the difference displayed in Figure 7.9 show another important property of the ES. While the VaR can recover relatively fast after a shock, the tail contains more risk over a longer period of time. Year 2008-2009 in Figure 7.9 displays this property particularly well; if the ES had recovered as fast as the VaR, the spike would have been narrower than the spike observed here.

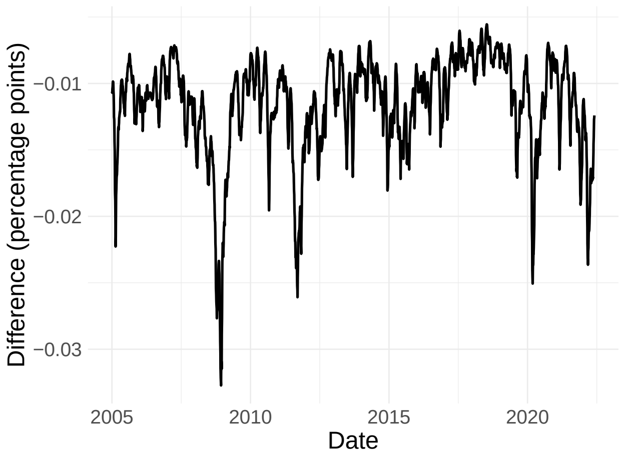

Figure 7.9: The 99% VaR minus the 97.5% ES for the simple model.

We then compare the models with and without the policy rate variable to investigate the effect of policy rate changes on interest rate risk. A summary of the computed values of the model selection criteria from Section 3.8 is given in Table 7.3. Furthermore, a lagged factor variable model is included in this table. This lagged model is identical to the factor variable model, but the observations of the policy rate are pushed ten trading days back in time. That is,

$$z_{\text{new}} = \left(z_{11,\text{orig}}, z_{12,\text{orig}}, \dots, z_{n-10,\text{orig}}, 0, \dots, 0\right)$$

where z_{new} denotes the lagged variable, $z_{i,\text{orig}}$ is the i'th observation of the original policy rate change variable and there are 10 zeros at the end of z_{new}. Hence, we allow for some absorption time of the policy rate changes into the market prices.

	DIC	WAIC	CPO	Mlik
Without factor variable	13052.17	13126.61	1.470	-6614.61
With factor variable	13054.69	13131.76	1.473	-6628.36
With lagged factor variable	13054.23	13129.99	1.471	-6627.77

Table 7.3: Model selection criteria values for the two models, as well as for a lagged factor variable model.

From Table 7.3, all criteria but the marginal likelihood favour the simple model without the policy rate variable. The differences between the criteria of the three models are very small, implying that adding the variable does not improve the model, either if it is lagged or not. We know that the policy rate influences the interest rates in real life. The reason why our model does not indicate the same can be that the information following policy rate changes is already incorporated by the market. The lagged model seeks to test this more thoroughly, but it seems the interest rate series already contain this information even when the variable is lagged. Furthermore, the criteria slightly favour that if the factor variable should be included in the model, it should be lagged. We choose to use the non-lagged variable in our analysis because of the very small and insignificant difference.

7.2 Comparison of INLA and HJM

We proceed with a comparison of INLA and HJM, where the latter is more commonly used in the financial industry. The two methods are built upon quite different strategies, which is important to keep in mind when comparing the results presented in this section.

The evolution of the yield curve is modelled as described in Section 3.7 for a period of one year. For one such simulation, we find the bond prices and the daily return in the portfolio from day 251 to 252. The last two dates of the simulation are chosen to include as much variability as possible into our risk measures. The HJM simulation is run $n = 100$ times, and the process is repeated $m = 1000$ times to find the return distribution from which we can measure the risk.

The simple model of the previous section is chosen as the INLA model for the subperiods. There is a possibility that the more complex model would have been preferred over the simple model for the subperiods, in terms of the model selection criteria. However, we have seen that these models give very similar results, and the essence of the INLA method will certainly be reflected in the results also for the simple model.

The following subsections present results from our four different subperiods. For the INLA models, the posterior means of η and the 97.5% VaR and ES are presented. The 2.5% and 97.5% quantiles are included in the posterior mean plots to display the corresponding uncertainty. For the HJM simulations, the distributions of the returns on the last days for each simulation period are presented by histograms. These results are summarised and discussed in Section 7.2.5. Here, the risk measures for the HJM simulations are also presented. Finally, the inherent differences of the INLA and HJM methods are discussed based on our results.

When reading these results and the following discussion, it is important to remember the differences of the two methods. The INLA results come from a regression-style model, which does not make any predictions. Contrarily, the HJM simulations are predictions, estimating the price after simulating one year forward in time. Since these predictions incorporate the trend of the corresponding calibration period, they can be thought of as the risk of the trend.

The mechanisms behind the INLA and HJM methods used in this thesis are sketched in Figure 7.10 and Figure 7.11, respectively. The INLA model is very local; each value of η is constructed by given neighbourhood observations because of the sparsity of the Gaussian Markov random field. The construction of the autoregressive time series ensures that present shocks are propagated into future observations of the process. On the other hand, the HJM model gives one single global risk measure from the return distribution. In addition to the computational efficiency, this local versus global level of measurement is the most prominent difference of the two methods.

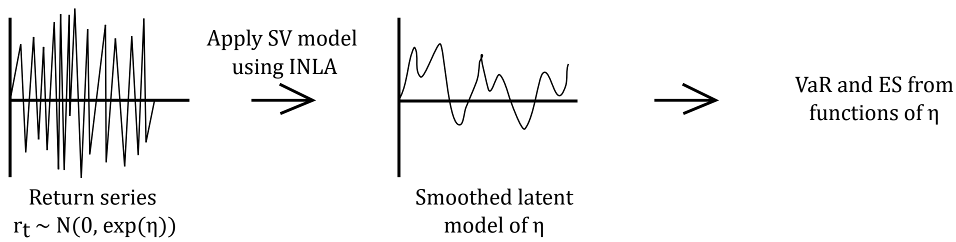

Figure 7.10: Sketch of the INLA method. The INLA framework models the latent process η of a time series, here the returns of the bond portfolio. Since the exponential of this latent process is the variance of the return series, and since we obtain a Gaussian distribution of η for each time step, we can find the VaR and ES using the parametric method.

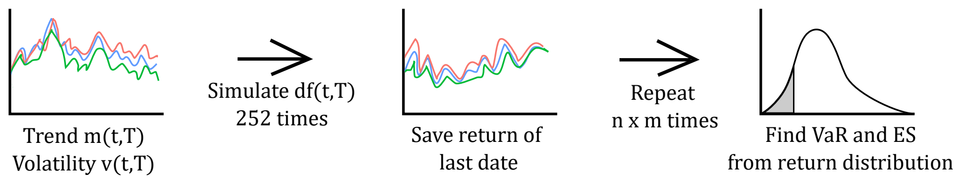

Figure 7.11: Sketch of the HJM method. The trend and volatility of some interest rate series are found, and then used to simulate the increments $df(t,T)$. This yields a new, simulated interest rate series. The relevant bond returns are found for the last date of this simulation, and the process is repeated n times to price the bond portfolio. Then, this is repeated m times to obtain a return distribution from which we can find the VaR and ES using the non-parametric method.

Comparing these two methods despite their differences is interesting because they take different approaches to reach the same goal. HJM is an established method for finding the risk of financial time series, modelling the term structure and the market expectations of the future interest rate. If INLA gives similar results as the HJM framework in terms of the order of magnitude of the risks, it indicates that INLA is an appropriate method within quantitative finance - although it does not preserve the term structure. The appropriateness of INLA is what we investigate by the following sections.

7.2.1 Financial crisis period

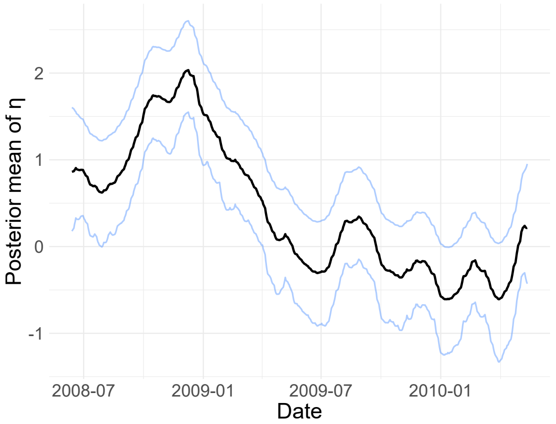

(a) The posterior mean of η, where the blue lines represent the 2.5% and 97.5% quantiles.

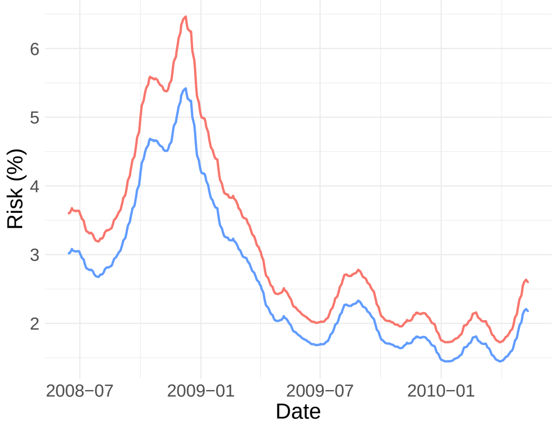

(b) The 97.5% VaR and ES. The red line represents the ES, and the blue line represents the VaR.

Figure 7.12: INLA results for the model calibrated on the financial crisis period.

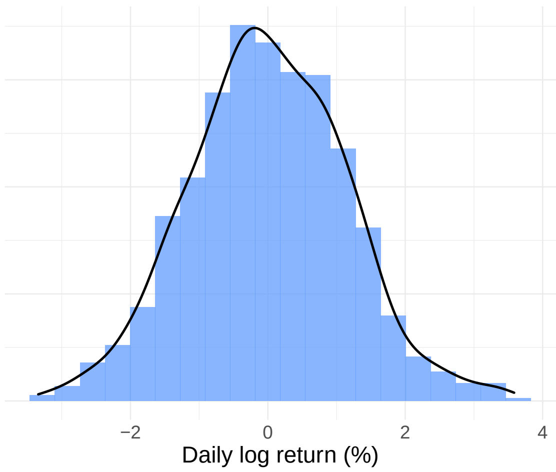

Figure 7.13: The distribution of the returns from the HJM model calibrated on the financial crisis period. The black line represents the smoothed distribution of the blue histogram.

7.2.2 Normal conditions period

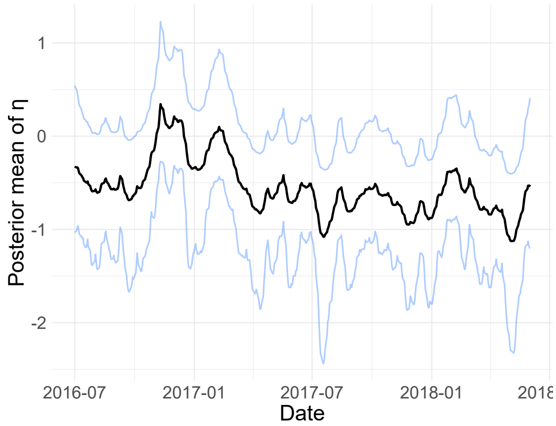

(a) The posterior mean of η, where the blue lines represent the 2.5% and 97.5% quantiles.

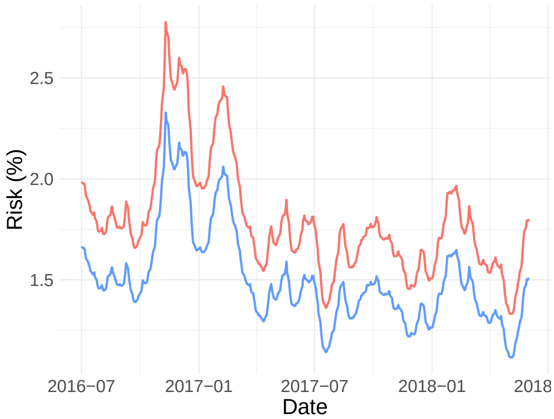

(b) The 97.5% VaR and ES. The red line represents the ES, and the blue line represents the VaR.

Figure 7.14: INLA results for the model calibrated on the normal conditions period.

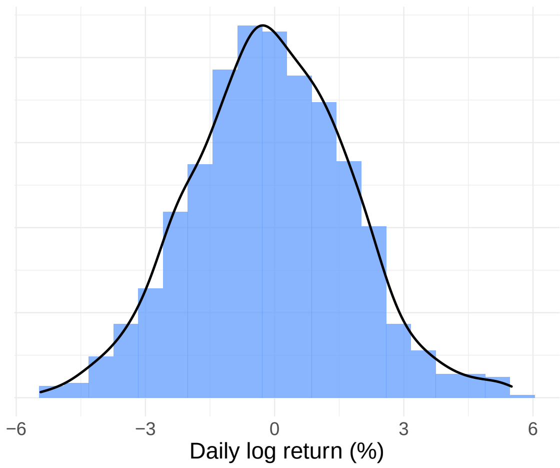

Figure 7.15: The distribution of the returns from the HJM model calibrated on the normal conditions period. The black line represents the smoothed distribution of the blue histogram.

7.2.3 Covid-19 outbreak period

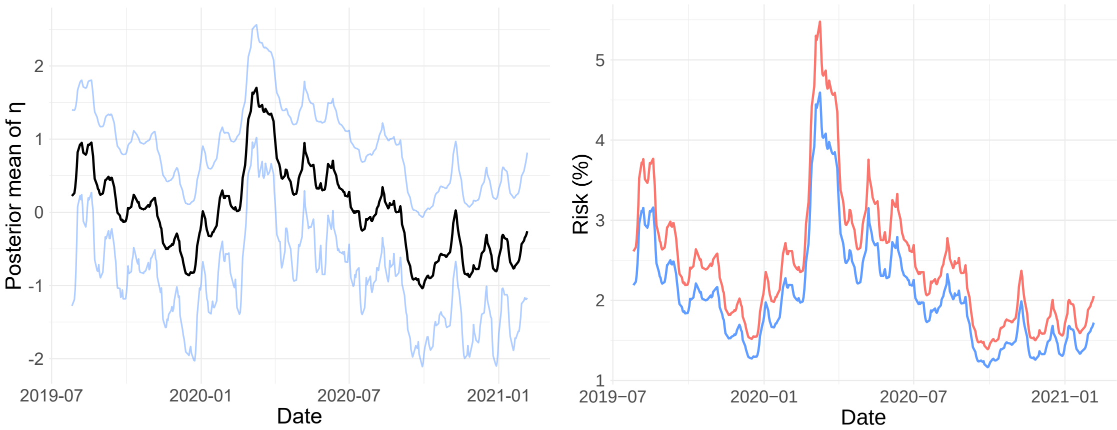

(a) The posterior mean of η, where the blue lines represent the 2.5% and 97.5% quantiles.

(b) The 97.5% VaR and ES. The red line represents the ES, and the blue line represents the VaR.

Figure 7.16: INLA results for the model calibrated on the Covid-19 outbreak period.

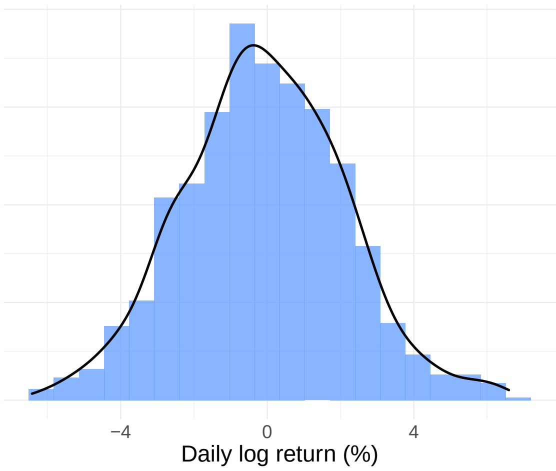

Figure 7.17: The distribution of the returns from the HJM model calibrated on the Covid-19 outbreak period. The black line represents the smoothed distribution of the blue histogram.

7.2.4 Covid-19 recovery period

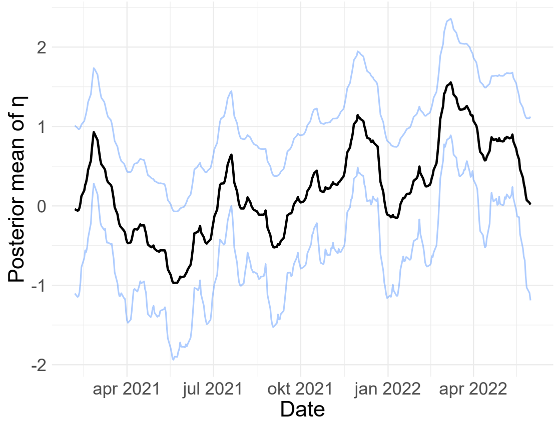

(a) The posterior mean of η, where the blue lines represent the 2.5% and 97.5% quantiles.

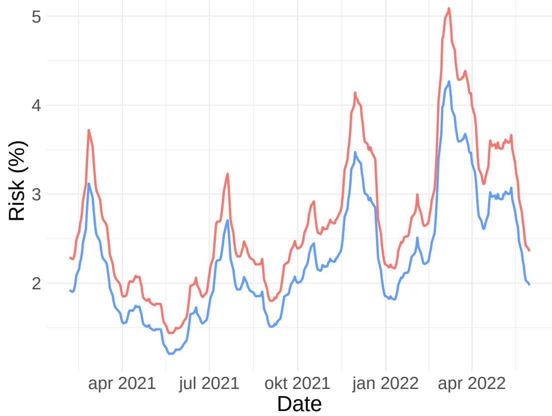

(b) The 97.5% VaR and ES. The red line represents the ES, and the blue line represents the VaR.

Figure 7.18: INLA results for the model calibrated on the Covid-19 recovery period.

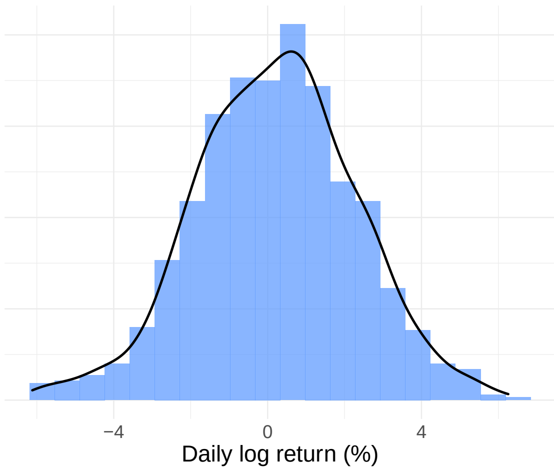

Figure 7.19: The distribution of the returns from the HJM model calibrated on the Covid-19 recovery period. The black line represents the smoothed distribution of the blue histogram.

7.2.5 Discussion

In the previous sections, INLA and HJM results for each subperiod are displayed. For INLA, the posterior means with the corresponding 2.5% and 97.5% quantiles, as well as the 97.5% VaR and ES, are presented. These quantities are observed every day of the subperiods. For the HJM simulations, histograms of the daily logarithmic returns are presented. As previously mentioned, these histograms represent the daily change from the second last to the last day of the simulation period, to include more variability. With the chosen method, we only obtain one measure of VaR and one of ES for the entire period for each of the HJM simulations. These risk measures are given in Table 7.4.

		VaR (%)	ES (%)
Financial crisis	97.5%	2.172	2.650
	99%	2.635	2.951
Normal conditions	97.5%	3.634	4.282
	99%	4.284	4.739
Covid-19 outbreak	97.5%	4.395	5.149
	99%	5.232	5.706
Covid-19 recovery	97.5%	4.147	4.991
	99%	5.074	5.644

Table 7.4: The 97.5% and 99% VaR and ES from the HJM simulations for the four subperiods.

Furthermore, the mean risk measures of the INLA subperiod models are given in Table 7.5.

		VaR (%)	ES (%)
Financial crisis	97.5%	2.598	3.098
	99%	3.083	3.532
Normal conditions	97.5%	1.506	1.796
	99%	1.787	2.047
Covid-19 outbreak	97.5%	2.069	2.468
	99%	2.456	2.813
Covid-19 recovery	97.5%	2.278	2.717
	99%	2.703	3.097

Table 7.5: The mean 97.5% and 99% VaR and ES from the INLA models of the four subperiods.

A HJM simulation gives an estimate for the forward rates, and hence also an estimate for the portfolio returns in our case. These estimates rely heavily on the trend of the corresponding calibration period, which explains the somewhat surprising HJM results of the financial crisis period. The calibration period includes the crash of the financial crisis and ends in May 2010. In the latter months of this period, the shock had dampened, and the yield curve had decreased to a low level. From Figure 7.13 and the risk measures in Table 7.4, we observe relatively low risks. This can be a consequence of the falling trend which has been picked up by the simulations, combined with low rates at the starting date of the simulation.

Comparing the financial crisis HJM results to the normal conditions results, the effect discussed in the previous paragraph is made even clearer. It seems reasonable that the financial crisis model would have the highest risks of these two periods, yet this is not the case. From the INLA results of the normal conditions period, displayed in Figure 7.14, we observe that there is a much weaker trend in this subperiod. This can be the explanation of why the daily returns of the normal conditions period take on a larger range of values, giving higher risks as well.

Because of the different properties of the two strategies used to measure the risk of the two methods, Table 7.4 and Table 7.5 should not be directly compared, as previously mentioned.

To compare the risk measurements of the two different strategies, we should instead rely on a combination of the daily risks and the mean of these for the INLA model. For the financial crisis period, the HJM results are fairly similar to the last year of the corresponding INLA model. Moreover, the maximal 97.5% ES from the INLA models of the two Covid-19 periods are similar to the 97.5% ES of the HJM models.

The normal conditions period is where the two strategies differ the most. Here, the maximal 97.5% ES of the INLA model is approximately 2.75%, while the corresponding measure of the HJM model is 4.28%. As previously discussed, this subperiod has the least trend of all the periods, explaining why the HJM simulations take on a much broader interval for the daily returns.

The INLA framework is very flexible and can easily fit complex models with many variables. The flexibility allows for a local level operation of the model which still takes the previous history into account, yielding high-quality results for a range of different models. Even though the framework is very flexible, the models are also easy to interpret. More complex models will certainly be somewhat harder to interpret, but the overall structure of an INLA model makes interpretation easier than similar methods. For instance, the interpretation of the stochastic volatility model of this thesis, with or without the factor variable, is straightforward.

Even though we have not done so here, INLA can be used for prediction. When simulating forward in time using INLA, we still obtain local, daily risk measures. However, these measures will be much less precise than the measures from a regression. This is because the distributions of the parameters will have significantly larger variances, introducing more uncertainty to the daily predictions.

Furthermore, INLA is based on the observations of the given data set and does not have any interest rate properties embedded like the HJM framework. That is, if we base an INLA prediction on a set of observations with no clear trend, we have no guarantee that the predictions will be reasonable in terms of interest rate properties. However, if there are clear trends in the observations, e.g., a periodicity, the INLA predictions can be as accurate as the HJM predictions. In the end, no model can precisely predict the future. Hence, if the desired properties are preserved, there are no reasons why a HJM prediction should be qualitatively better than an INLA-based prediction.

Chapter 8

Conclusion

This thesis aimed to investigate the effect of policy rate changes on interest rate risk. Implementing a stochastic volatility model under the INLA framework, we modelled the volatility of the daily returns of a bond portfolio based on NIBOR rates. Two risk measures related to the size of the losses, VaR (value at risk) and ES (expected shortfall), were calculated by a parametric method. Comparing models with and without a policy rate change variable, the selection criteria indicated that including this variable did not improve the ability to model the risk.

INLA was chosen and investigated because of its computational efficiency, with a much faster runtime than MCMC algorithms. The framework is not established within quantitative finance as of today. Thus, this thesis studied how INLA can be used in a quantitative setting by applying the framework to typical problems within the subject. We found that INLA provided a flexible and user-friendly framework, simplifying our modelling and analysis because of its high interpretability. The results from the stochastic volatility models implemented in INLA were found to be sensible and well-performing, according to our expectations, the historical risk and the economic declines of the data set. Both stochastic volatility models were found to have a high persistence, implying that previous shocks were incorporated in the daily volatilities found by INLA.

The ES was found to be higher than the VaR for all models and dates, as expected. The most distinct differences occurred during and right after recessions. We found that the VaR recovered faster than the ES after a shock, because the tail contains some risk that only the latter is able to reveal.

To compare the standards of the previous and prevailing framework standards of Basel II and Basel III, we compared the 99% VaR to the 97.5% ES for the simple model. For all dates, the 97.5% ES was larger than the 99% VaR. The biggest differences were observed in times where the overall risk levels were high, e.g., during the financial crisis and the Covid-19 outbreak. We thus concluded that the ES is indeed a better risk indicator than the VaR, as it contains some risk that the VaR ignores.

We found that adding a policy rate variable did not improve the model. However, the criteria values for the two models were very close, indicating that adding the factor variable did not make the model significantly worse either. For instance, the DIC of the simple model was 13052.17, while the DIC of the more complex model was 13054.69. We know that the policy rate does have an impact on the interest rates in reality. A possible explanation as of why the model did not indicate this could be because the policy rate information was already absorbed and incorporated by the interest rate series. Then, adding the same information through the policy rate variable could not improve the model.

INLA was compared to the more established HJM by splitting the data set into four different subperiods. For three of the four subperiods, the two methods gave similar results. The fourth subperiod gave more divergent results, but still of the same order of magnitude. These results indicated that INLA could be an alternative method to measure the risk of financial time series.

Future work on the same topic could compare INLA to other methods, as there exist MCMC algorithms which can be used in the same manner as INLA. The MCMC model can then be on the same hierarchical form as our INLA model and can thus include the factor variable as we did in this thesis. Martino, Aas, Lindqvist, Neef and Rue [4] found that INLA and the MCMC program OpenBUGS gave very similar results for financial time series, with the MCMC error as the only quantitative difference. However, the runtimes of the two methods are not comparable. For computational efficiency, INLA will be better choice than MCMC.

In future similar studies, it would be interesting to test the predictive abilities of INLA for the risk measures. This would give a closer comparison to the HJM framework, yielding a more detailed answer on how suitable INLA is for quantitative finance. For a prediction model, another underlying process might be more appropriate than the AR(1) process suggested here. To approach the problem in a more similar way as the HJM and MCMC methods, a random walk model would probably be appropriate.

Introducing other variables to the model could provide insight to mechanisms behind interest rate risk. Such variables could be the policy rates of the FED and the European Central Bank, or the US 10 Year Treasury Rate, as we know that these rates are important in the global economy. Furthermore, investigating how market expectations influence the risk could also give valuable and interesting information. As far as we know, no such data exists for the Norwegian market today; this research would thus require another study collecting the needed data.

Finally, analysing the lagging between the policy rate changes and the following changes in the interest rate risk could be researched more thoroughly in the future. If a connection is found, we could improve the prediction of interest rate risk, which would be valuable for all investors and risk managers. Optimising the portfolio could also improve the predictive powers of the model. The portfolio used in this thesis was static and uniformly weighted in order to simplify the complexity of the model. This is sufficient when we wish to study overall connections, but a more optimal portfolio could help the model to find more complex correlations.

In conclusion, INLA seems to be a good alternative to established methods. The numerical errors are small and comparable to MCMC errors. Moreover, the INLA algorithm is exceptionally efficient; the runtime is almost negligible compared to HJM. The flexibility of INLA allows for prediction as well as regression, providing a complete framework for quantitative finance. Furthermore, the method gives daily measurements, precisely picturing the stochastic nature of the financial risk.

Bibliography

[1] Bachelier L. Théorie de la spéculation. *Annales scientifiques de l'École Normale Supérieure*, 3(17):21–86, 1900.

[2] Black F and Scholes M. The Pricing of Options and Ccorporate Liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

[3] Chopin N, Martino S, and Rue H. Approximate Bayesian Inference for Latent Gaussian Models Using Integrated Nested Laplace Approximations. *Journal of the Royal Statistical Society*, 71(Part 2):319–392, 2009.

[4] Martino S, Aas K, Lindqvist O, Neef LR, and Rue H. Estimating stochastic volatility models using integrated nested laplace approximations. *The European Journal of Finance*, 17(7):487–503, 2011.

[5] Fisher I. *The Theory of Interest: As Determined by Impatience to Spend Income and Opportunity to Invest It.* MacMillan, 1930.

[6] Fama EF. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2):383–417, 1970.

[7] van der Wijst N. *Finance: A Quantitative Introduction.* Cambridge University Press, 2013.

[8] Wilmott P, Howison S, and Dewynne J. *The Mathematics of Financial Derivatives: A Student Introduction.* Cambridge University Press, 1995.

[9] Lov om Norges Bank og pengevesenet mv. (sentralbankloven)[Central Bank Act]. LOV-2019-06-21-31. URL `https://lovdata.no/dokument/NL/lov/2019-0621-31`. Last accessed 2 February 2022. In Norwegian.

[10] Bestemmelse om pengepolitikken [Regulation on Monetary Policy]. FOR-2019-12-13-1775. URL `https://lovdata.no/dokument/INS/forskrift/2019-12-13-1775`. Last accessed 2 February 2022. In Norwegian.

[11] Central Bank News. Inflation Targets, 2022. URL `http://www.centralbanknews.info/p/inflation-targets.html`. Last accessed 2 February 2022.

[12] Norges Bank. Norges banks pengepolitiske strategi, 2021. URL https://www.norges-bank.no/tema/pengepolitikk/pengepolitisk-strategi/. Last accessed 2 February 2022.

[13] Corporate Finance Institute. Liquidity Trap, 2019. URL https://corporatefinanceinstitute.com/resources/knowledge/economics/liquidity-trap/. Last accessed 24 February 2022.

[14] Norges Bank. Renter, 2021. URL https://app.norges-bank.no/query/#/no/interest. Last accessed 2 February 2022.

[15] Norges Bank. Styringsrenten, 2021. URL https://www.norges-bank.no/tema/pengepolitikk/Styringsrenten/. Last accessed 2 February 2022.

[16] Norges Bank. How does the policy rate influence other interest rates?, 2019. URL https://www.norges-bank.no/en/knowledge-bank/the-policy-rate/how-does-the-policy-rate-influence-other-interest-rates/. Last accessed 16 May 2022.

[17] Norges Bank. How does the policy rate influence the economy and inflation?, 2019. URL https://www.norges-bank.no/en/knowledge-bank/the-policy-rate/how-does-the-policy-rate-influence-the-economy-and-inflation/. Last accessed 16 May 2022.

[18] National Bureau of Economic Research. Business Cycle Dating Committee Announcement January 7, 2008, 2008. URL https://www.nber.org/news/business-cycle-dating-committee-announcement-january-7-2008. Last accessed 30 March 2022.

[19] McBride J. The U.S. Financial Crisis, 2018. URL https://www.cfr.org/timeline/us-financial-crisis. Last accessed 30 March 2022.

[20] Johansson P and Meldrum A. Predicting Recession Probabilities Using the Slope of the Yield Curve, FEDS Notes, 2018. URL https://doi.org/10.17016/2380-7172.2146. Last accessed 30 March 2022.

[21] Miller DS. There is No Single Best Predictor of Recessions, 2019. URL https://doi.org/10.17016/2380-7172.2367. Last accessed 30 March 2022.

[22] International Monetary Fund. Policy Responses to Covid-19, 2021. URL https://www.imf.org/en/Topics/imf-and-covid19/Policy-Responses-to-COVID-19. Last accessed 30 March.

[23] Bougroug A, Kjos ØK, and Sletten P. Økonomisk utvikling gjennom Covid-19. En sammenligning av utvikling i Norge, Sverige og Danmark. Technical report, Statistisk sentralbyrå.

[24] Statistisk sentralbyrå. KPI opp 5.3 prosent siste tolv måneder, 2022. URL `https://www.ssb.no/priser-og-prisindekser/konsumpriser/statistikk/konsumprisindeksen/artikler/kpi-opp-5-3-prosent-siste-tolv-maneder`. Last accessed 16 May 2022.

[25] Bache IW, Almås I, Børsum Ø, and Fjære-Lindkjenn J. Pengepolitisk rapport med vurdering av finansiell stabilitet. Technical report, Norges Bank, March 2022.

[26] Statistisk sentralbyrå. Price index for existing dwellings, 2022. URL `https://www.ssb.no/en/priser-og-prisindekser/boligpriser-og-boligprisindekser/statistikk/prisindeks-for-brukte-boliger`. Last accessed 16 May 2022.

[27] Finanstilsynet. Boliglånsundersøkelsen 2021, 2021. URL `https://www.finanstilsynet.no/nyhetsarkiv/pressemeldinger/2021/boliglansundersokelsen-2021/`. Last accessed 16 May 2022.

[28] Trading Economics. United States Inflation Rate, 2022. URL `https://tradingeconomics.com/united-states/inflation-cpi`. Last accessed 29 May 2022.

[29] Cox J. Fed raises rates by half a percentage point - the biggest hike in two decades - to fight inflation, 2022. URL `https://www.cnbc.com/2022/05/04/fed-raises-rates-by-half-a-percentage-point-the-biggest-hike-in-two-decades-to-fight-inflation.html`. Last accessed 9 May 2022.

[30] International Monetary Fund. World Economic Outlook. Technical report, April 2022.

[31] Investopedia. Financial Instrument. URL `https://www.investopedia.com/terms/f/financialinstrument.asp`. Last accessed 21 March 2022.

[32] Durrett R. Brownian Motion. In *Probability: Theory and Examples*. Cambridge University Press, 5th edition, 2019.

[33] Wilmott P. *Paul Wilmott On Quantitative Finance*, volume 1. John Wiley & Sons, Incorporated, 2nd edition, 2006.

[34] Itô K. Stochastic Integral. *Proc. Imp. Acad. Tokyo*, 20(8):519–524, 1944.

[35] Allenby GM, McCulloch RE, and Rossi PE. Hierarchical Bayes Models: A Practitioners Guide, 2005. URL `http://faculty.washington.edu/bajari/iosp07/rossi1.pdf`. Last accessed 9 March 2022.

[36] Wilmott P. *Frequently Asked Questions in Quantitative Finance*. John Wiley & Sons, Incorporated, 2nd edition, 2010.

[37] Givens GH and Hoeting JA. *Computational Statistics*. John Wiley & Sons, Incorporated, 2nd edition, 2012.

[38] Martino S. Lecture 11: Integrated Nested Laplace Approximation (INLA), 2021. URL `https://www.math.ntnu.no/emner/TMA4300/2021v/Slides/Lecture11.pdf`. Last accessed 3 May 2022.

[39] Martino S and Riebler A. Integrated Nested Laplace Approximations (INLA), 2019. URL `https://www.math.ntnu.no/emner/TMA4300/2020v/LectureNotes/inla_paper.pdf`. Last accessed 3 May 2022.

[40] Jolliffe IT. *Principal Component Analysis*. Springer, 2nd edition, 2002.

[41] European Central Bank. Key ECB interest rates, 2019. URL `https://www.ecb.europa.eu/stats/policy_and_exchange_rates/key_ecb_interest_rates/html/index.en.html`. Last accessed 24 November 2021.

[42] Heath DD, Jarrow R, and Morton A. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica*, 60(1):77–105, 1992.

[43] Glasserman P. *Monte Carlo Methods in Financial Engineering*. Springer, 2003.

[44] Akaike H. Information theory and an extension of the maximum likelihood principle. *Second International Symposium on Information Theory*, pages 267–281, 1973.

[45] Spiegelhalter DJ, Best NG, Carlin BP, and Van Der Linde A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society*, 64 (4):583–639, 2002.

[46] Gelman A, Hwang J, and Vehtari A. Understanding predictive information criteria for bayesian models. *Statistics and computing*, 24(6):997–1016, 2014.

[47] Pettit LI. The Conditional Predictive Ordinal for the Normal Distribution. *Journal of the Royal Statistical Society*, 52(1):175–184, 1990.

[48] How to use CPO?, 2011. URL `https://groups.google.com/g/r-inla-discussion-group/c/cWwXBO11jWM`. Last accessed 20 May 2022.

[49] Verma E. Financial Risk and Its Types, 2022. URL `https://www.simplilearn.com/financial-risk-and-types-rar131-article`. Last accessed 25 May 2022.

[50] Artzner P, Delbaen F, Jean-Marc E, and Heath DD. Coherent Measures of Risk. *Mathematical Finance*, 9(3):203–228, 1999.

[51] Martin Haugh. IEOR E4602: Quantitative Risk Management. Basic Concepts and Techniques of Risk Management, 2019. URL `http://www.columbia.edu/~mh2078/QRM/BasicConceptsMasterSlides.pdf`. Last accessed 19 May 2022.

[52] Embrechts P and Wang R. Seven Proofs for the Subadditivity of Expected Shortfall, 2015. URL `https://people.math.ethz.ch/~embrecht/ftp/Seven_Proofs.pdf`. Last accessed 16 December 2021.

[53] Bank for International Settlements. History - overview, 2018. URL `https://www.bis.org/about/history.htm?m=1004`. Last accessed 25 May 2022.

[54] Bank for International Settlements. BIS member central banks, 2014. URL `https://www.bis.org/about/member_cb.htm?m=1%7C2%7C601`. Last accessed 25 May 2022.

[55] Europex. Capital Requirements Directive (CRD) IV and Capital Requirements Regulation (CRR), 2022. URL `https://www.europex.org/eulegislation/crd-iv-and-crr/`. Last accessed 25 May 2022.

[56] Bank for International Settlements. History of the Basel Committee, 2021. URL `https://www.bis.org/bcbs/history.htm`. Last accessed 1 February 2022.

[57] Basel Committee on Banking Supervision. The market risk framework - In Brief. Technical report, Bank for International Settlements, January 2019.

[58] Basel Comittee on Banking Supervision. Revisions to the Basel II market risk framework. Technical report, Bank for International Settlements, January 2009.

[59] Basel Comittee on Banking Supervision. Minimum capital requirements for market risk. Technical report, Bank for International Settlements, January 2009.

[60] Investopedia. Yield curve, 2022. URL `https://www.investopedia.com/terms/y/yieldcurve.asp`. Last accessed 25 May 2022.

[61] Gilles C. Volatility and the Treasury yield curve, 2017. URL `https://www.bis.org/publ/confp01m.pdf`. Last accessed 30 November 2021.

[62] Taylor SJ. Financial Returns Modelled by the Product of Two Stochastic Processes - A Study of Daily Sugar Prices 1961-79. In Anderson OD, editor, *Time series analysis: Theory and practice*, volume 1, pages 203–226. Amsterdam, 1982.

Appendix: INLA Results

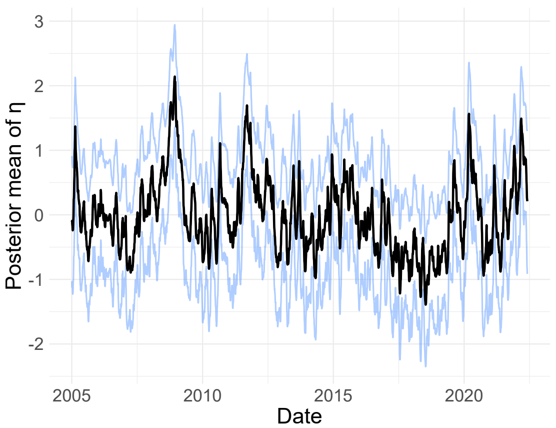

(a) Factor group $z = 0$

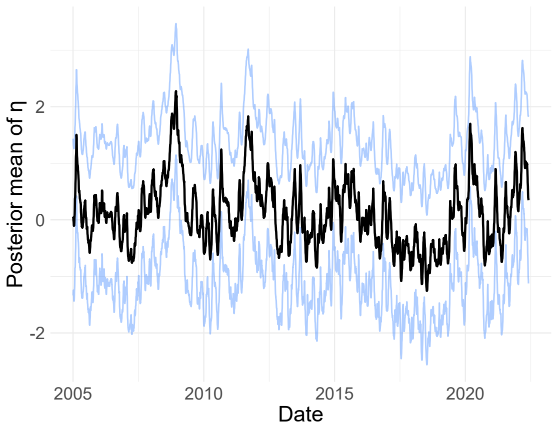

(b) Factor group $z = 1$

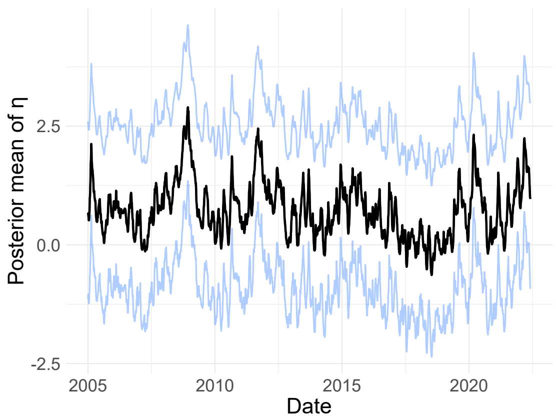

(c) Factor group $z = 2$

Figure A.1: The posterior mean of η for each of the three factor levels with their respective 2.5% and 97.5% quantiles.

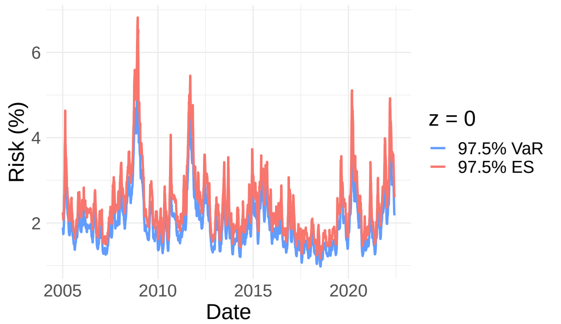

(a) Factor group $z = 0$

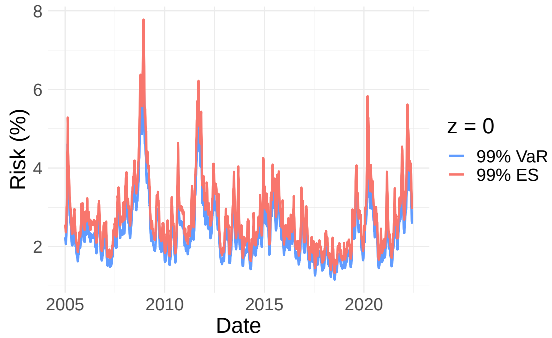

(b) Factor group $z = 0$

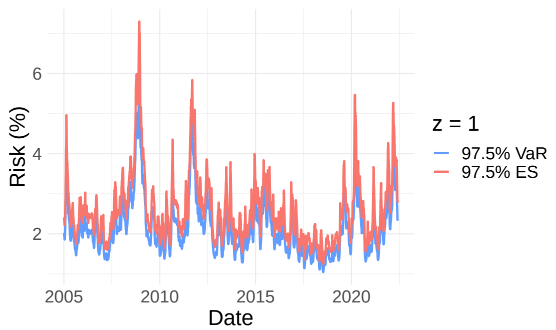

(c) Factor group $z = 1$

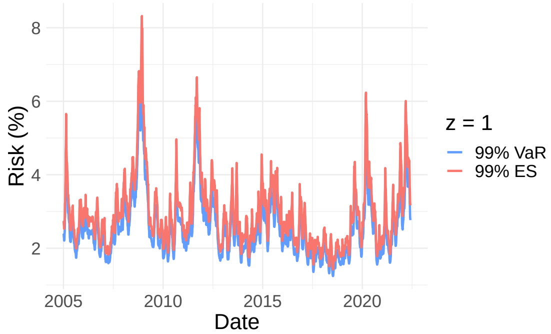

(d) Factor group $z = 1$

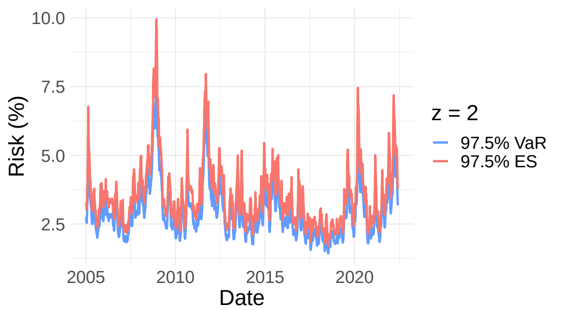

(e) Factor group $z = 2$

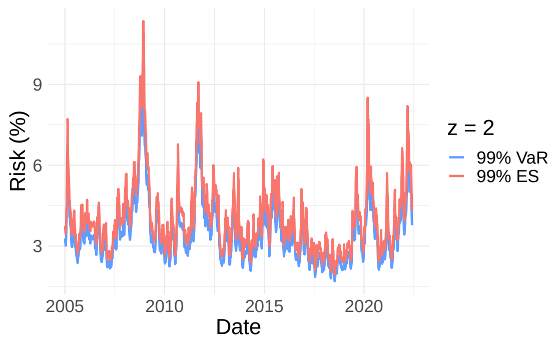

(f) Factor group $z = 2$

Figure A.2: The 97.5% and 99% VaR and ES for the three different factor levels of the factor variable model.

NTNU

Norwegian University of
Science and Technology