Erasmus University Rotterdam

Erasmus School of Economics

MSc ECONOMETRICS AND MANAGEMENT SCIENCE:

QUANTITATIVE FINANCE

Market-consistent valuation of inflation-linked liabilities

MSc THESIS

Author

J. Roza

University

Supervisor: Prof. dr. M. van der Wel

Second assesor: dr. M. Grith

Student ID

511222

Supervisors Aegon

F. Roijers

J. Bakker

December 2, 2020

Abstract

This thesis performs a market-consistent valuation of inflation-linked liabilities by combining the Jarrow-Yildrim (JY) model with an existing LIBOR market model (LMM). In the JY model, the nominal interest rates, real interest rates and inflation rate are modelled simultaneously using a HJM framework. In this thesis, however, we replace the nominal short rate projection by a nominal short rate approximation resulting from an LMM. The inflation and real rate component of our model are calibrated to market data of zero-coupon inflation swaps (ZCIIS), year-on-year inflation swaps (YYIIS) and inflation floors. We minimise the sum squared difference of the model- and market prices by adjusting some of our model parameters. The optimised parameters are utilised to generate our final inflation projection. This inflation projection is used to price a dummy portfolio mimicking the inflation-linked liabilities with an embedded floor at 0%, which is commonly found in practice. The model is able to market-consistently price the ZCIIS, YYIIS, and 0% floors and finds that the floor value of the inflation-linked liabilities is 0.3% of the present value without a floor.

Keywords— inflation risk, inflation-linked derivatives, JY model, LMM

Contents

1 Introduction

Insurance companies have both assets and liabilities on their balance sheet and are therefore subject to several financial risks. It is crucial to accurately estimate these risks in order to give a clear insight in their potential impact on the assets and liabilities. The reason for this is twofold. First of all, for the company itself it is important to understand the risks they are facing. By understanding this, not only the impact of a change in one of the financial variables will be known, but also this impact can be used for hedging the risk. Secondly, regulations require insurance companies to be precise about the financial situation. An example is that a requirement of Solvency II, a valuation framework for insurance firms in the EU, is that the valuation needs to be market-consistent. By analysing inflation, this thesis covers one of these financial risks.

Specifically, the purpose of this thesis is to perform a market-consistent valuation of inflation-linked liabilities by extending an interest rate risk model to an inflation model with the focus on estimating the non-linear inflation risk embedded in these liabilities. In previous literature, inflation has had much less attention than for example interest rate risk. The first and main reason is that the exposure to inflation risk is not that large. The second reason is that, historically, inflation has had a low probability of falling below 0%. Recently, however, non-linear inflation risk has become more realistic. Furthermore, exposure to inflation risk can always rise when adapting the portfolio. It thus an interesting topic to investigate. As mentioned above, insurance companies have a lot more exposure to interest rate risk and it is therefore likely that they have some model to capture this risk. By assuming that these companies have an interest in modelling inflation, the need for a compatible inflation model arises. We develop such a model in this thesis. The main contribution of our thesis is thus that it extends an interest rate risk model such that it can also be used to estimate inflation risk, a risk that has become more realistic the last few years and has not received a lot of attention.

Inflation risk is a risk for companies whose portfolio consist of inflation-linked liabilities. Inflation-linked liabilities are liabilities that are indexed with realised inflation (usually an inflation index such as HICPxT or NL CPI) to stabilise the policy holder's purchasing power. Since indices vary over time, the value of the liabilities becomes uncertain. To protect policy holders against the uncertainty, the liabilities are not negatively indexed. This means that when inflation falls below 0%, the liabilities are kept constant rather than being reduced. The policy holder therefore has an embedded guarantee to receive at least 0% inflation; his pay-out can be modelled as the maximum of 0% inflation and the actual inflation. This leads to non-linear inflation risk for the pension insurer.

The presence of the 0% floor affects the value of the liabilities through the guarantee and therefore a stochastic valuation model is required. The guarantee can be split into the intrinsic value and the time value. The intrinsic value refers to an investor's perception of the inherent value of an asset. The time value refers to the portion of an option's premium that is attributable to the amount of time remaining until the expiration of the option. To incorporate this time value, the valuation should consider the possibility that a deflation will occur. It is important to note that therefore the break-even inflation curve cannot be used for the valuation. This curve is always above 0% and thus does not allow for the possibility that inflation will fall below 0% and can therefore not incorporate the effect of this floor. In order to capture this effect, a stochastic valuation is necessary that projects different inflation paths. Each path will have a different evolution of inflation and allows inflation to be negative in some situations.

Several stochastic models are available and to choose our own model, we first review multiple existing inflation models with the main focus on the model by Jarrow and Yildirim (2003). As for interest rate models, the inflation models can essentially be divided into two main groups: market models and short rate models. Short rate models have the purpose of modelling the unobservable instantaneous rate. One of the first short rate models to model inflation is the JY model of Jarrow and Yildirim (2003). In the JY model, the nominal rates, real rates and inflation are modelled simultaneously using a HJM framework. This is exactly the reason we focus on this model when studying short rate models. Since all these three variables are modelled, it is quite intuitive to combine a nominal interest rate risk model with the JY model.

The JY model contains three possible shortcomings or rather simplicities. First of all, the volatility is assumed to be deterministic. Secondly, the model does not allow for jumps in the inflation rate. Thirdly, it is driven by only three factors. Since the introduction of the JY model, more models have been suggested that all try to overcome some of these simplicities. One extension of the JY model is suggested by Hinnerich (2008). This model allows for jumps in the inflation rate. There is strong empirical evidence that interest rates have embedded jumps, mainly caused by information surprises like macroeconomic announcements. Therefore jump-diffusion models should more accurately capture the behaviour. Since this also holds for real interest rates and inflation, it is natural to allow inflation to jump as well. Furthermore, it allows for more than three factors. In other words, the random process describing the instantaneous rates and inflation is driven by a multidimensional Wiener process. Lastly, the model introduces stochastic volatility to capture the smile effect in inflation derivatives.

Market inflation models adopt dynamics for certain observable inflation-related variables, like the CPI or the forward price of real zero-coupon bonds. Well-known market models are those of Belgrade et al. (2004) and Mercurio (2005). The first market model adopts dynamics for the forward CPI. The second one uses lognormal dynamics of the forward CPI together with the assumption that forward rates follow a lognormal LIBOR model to obtain dynamics for the forward price of inflation zero-coupon bonds. Market models have also been adapted in several ways. The market model by Mercurio and Moreni (2006) extends the model of Mercurio (2005) by introducing stochastic volatility instead of deterministic volatility. Later on, a multi-factor SABR model by Mercurio and Moreni (2009) was proposed. This model is based on a multi-factor volatility structure and leads to SABR-like dynamics for forward inflation rates. The advantage of this model is that it is able to jointly price zero-coupon products and year-on-year products.

In our thesis, we assume that a nominal LIBOR market model (LMM) is used by insurance companies to capture interest rate risk and we model the inflation by using the JY model. Short rate models have the drawback that they do not capture the current curve, a feature that market models do contain. On one hand, market models are more advanced. On the other hand, they are usually more complex and require more input parameters. Because of their increased dimensionality, market models may not be best choice in less liquid markets. Or, as van Haastrecht and Pelsser (2011) put it, "due to a lack of calibration instruments in less liquid markets (such as inflation options), hedges and calibrations may become unstable when using market models" (p. 683). This is crucial for our model selection, which is mainly driven by the combination of the given exposure and the liquidity of the specific market. The exposure to interest rate risk is large and hence an advanced model is required. The interest rate market is large and liquid and allows us to make this choice. Inflation risk exposure is small, making it less necessary to use an advanced model. Furthermore, the inflation market is not that liquid and hence market models may possess calibration problems, due to which in this thesis the JY model is studied.

This thesis examines the stochastic valuation by using the JY model, where the nominal evolution is replaced by the LMM. The JY model uses the foreign currency analogy so that nominal dollars correspond to domestic currency, real dollars to foreign currency and inflation index to the spot exchange rate. In that setup, the fluctuations of the real and nominal interest rates and the inflation rate will be correlated. The nominal projection, however, can also be developed by another model and this is exactly what we do. We use the LMM to make a projection of the forward curve for all projection years. The shortest maturity belonging to every forward curve is then extracted to create a nominal short rate evolution which replaces the nominal projection of the initial JY model. In our final model, the real rate and inflation rate will thus depend on the JY model, whereas the nominal interest rate will depend on the LMM.

To be able to market-consistently price the inflation-linked liabilities, we calibrate our model to inflation products. Since we assume that the LMM is calibrated to nominal swap data, the nominal rate projection should be able to market-consistently price these products. The real rate and inflation component are driven by the JY model and do need to be calibrated. Initially, we base the parameters driving these two components on time series of the inflation rate and the real rate. Afterwards, we optimise the values of the volatility parameters by calibrating our model to market data of zero-coupon inflation-indexed swaps (ZCIIS) year-on-year inflation-indexed swaps (YYIIS), and inflation floors. In this calibration process, we minimise the sum squared difference of the model- and market prices by using the Broyden-Flether-Goldfarb-Shanno (BFGS) algorithm.

We use the optimised set of parameters to estimate a dummy portfolio mimicking the inflation-linked liabilities with the embedded floor. After our calibration process, the model should be able to market-consistently price the several inflation products. Our final inflation paths are those used in that valuation. We then adapt the projected indices to incorporate the floor at 0% inflation. Using both the original and adapted index, we index the cash flows belonging to a dummy portfolio that approximates the inflation-linked liabilities. This procedure allows us to finally determine the present value for both the cash flows indexed with the original index and the cash flows indexed with the adapted index. The floor value is the present value without the embedded floor subtracted by the present value incorporating the 0% floor. Lastly, we perform a sensitivity analysis be pricing several other floors. The result of this analysis can eventually be used to hedge the inflation risk.

This thesis finds that, by calibrating to 0% floors only, the model is able to almost perfectly match the market prices of 0 % floors. As an out-of-sample test, we use those parameters to value the ZCIIS, YYIIS and 2% floor. The outcome is that, compared to the market prices, the ZCIIS and YYIIS model prices have a maximum deviation of 1%. The 2% floor prices have a larger deviation from the market prices. To investigate if we can improve this result, we calibrate to ZCIIS, YYIIS and 0 % floor prices simultaneously. It follows that the 0% floor prices remain almost the same, whereas the ZCIIS and YYIIS prices do get closer to the market value, but only by a little. This fit is thus not significantly better than the first fit. As a last fit, we calibrate on the 0% floor and 2% floor. The result is that the 2% floor prices better reflect the prices observed in the market, but that the 0% floor prices deviate away from the market prices compared to the first and second fit. Our model thus struggles to market-consistently price the 0% and 2% floors simultaneously.

Since our focus is on estimating inflation-linked liabilities, with an embedded floor at 0%, we use the parameter set belonging to fit 1 to price the dummy portfolio. The result is that the value of the floor is around 0.3% of the present value without the floor. The sensitivity analysis shows that the valuation is sensitive to a negative change in inflation. More specifically, if the actual inflation turns out to be 1%/2% lower than projected, the floor value is equal to around 1.65% /11% of the present value. The overall conclusion is that the model is able to market-consistently price 0% floors, ZCIIS and YYIIS using a specific set of parameters, but that it does struggle to match the market prices of 0% and 2% floors simultaneously. Using the parameter set of the first 1, the value of the floor in our specific inflation-linked liabilities is 0.3% of the present value without the floor, but the result is quite sensitive to a change in inflation.

The thesis is organised as follows. Section 2 gives an overview of the inflation market and the inflation floor affects our valuation of the liabilities. Section 3 explains the methodology used in this thesis, which covers the chosen model, the calibration procedure and the valuation of the inflation-linked liabilities. The data are discussed in section 4. The obtained results are shown and analysed in section 5. The thesis ends with the conclusion in section 6.

2 Inflation market

This section gives an overview of the inflation market and discusses several inflation products. A crucial part in our thesis is calibrating the model and a key aspect concerning the calibration procedure is the availability of the products and their liquidity. Therefore, in the process of model selection, it is important to have a good oversight of the specific market. The inflation market is a market that enables investors to trade on inflation for several purposes. Nominal instruments embed views on inflation because inflation is incorporated in the interest rate, but inflation-linked instruments allow investors to take inflation specific views. Inflation is used for indexation and is an economic inference, but with the instruments also tradable itself. An inflation instrument is linked to a specific price index. Several price indices exist, such as the harmonised indices of consumer prices (HICP) which is a composite measure of inflation in the Eurozone. Even though the inflation market is a lot smaller in size than the market for nominal instruments, the inflation market has grown in size and liquidity over the recent years.

2.1 Inflation-linked bonds

The simplest inflation product is the inflation-linked bond. In contrast to normal government bonds, where the bond holder loans money to a government in return for an agreed rate of interest, the outstanding principal and thus the face value increases when inflation occurs. The interest paid out by the bonds is also adjusted for inflation. These features make that inflation-linked bonds can diminish the real impact of inflation on the bond holder. This product has several purposes. Governments use it for inflation targeting monetary policy for example, while investors thus use it to protect the real value of their investment.

2.2 Inflation swaps

The most common and most liquid inflation derivative in the market is the zero-coupon inflation indexed swap (ZCIIS). For this product, the inflation buyer pays a predetermined fixed rate and in return receives from the inflation seller inflation-linked payments. In the ZCIIS, there is only one payment time, namely at maturity M. The predetermined fixed rate K is the expected average inflation rate over M years by the seller. Let us assume that the seller and buyer agree upon a particular nominal value N. Using the notation of Mercurio and Moreni (2006), the buyer then pays the seller the following fixed amount at maturity:

$$N[(1+K)^M - 1], \tag{1}$$

In return for the payment above, the buyer receives a floating payment dependent on the inflation rate. Therefore, to determine the payoff of the buyer, we look at the index at maturity I_M and the initial index I_0. So, at maturity, the seller then pays the buyer the following amount:

$$N\Big[\frac{I_M}{I_0} - 1\Big], \tag{2}$$

In an arbitrage-free economy, the two payoffs should be equal to each other. In other words, $(1+K)^M = \frac{I_M}{I_0}$. Hence, intuitively K indeed reflects the average inflation rate over M years the seller expects. As shown, this product allows an investor to secure an inflation-protected return with respect to an inflation index. Furthermore, the ZCIIS is highly flexible with maturities from one to 30 years.

Another product, very much related to the ZCIIS, is the year-on-year inflation indexed swap (YYIIS). This product pays the annual inflation at the end of each year. So, in contrast to the ZCIIS, there are multiple payment dates. The intuition concerning the payoff structure is similar to the ZCIIS, but now there is a payoff each year. Specifically, at each time t, the buyer pays the seller the fixed amount

$$N\psi_t K, \tag{3}$$

where ψ_t is the contract fixed-leg fraction for the interval $[t-1, t]$. At each time t, the seller pays the buyer the following amount:

$$N\psi_t\Big[\frac{I_t}{I_{t-1}} - 1\Big], \tag{4}$$

where I_t is the index in year t.

2.3 Inflation options

Next to swaps, there are options. Financial options can be split into call options and put options. A call option is a financial contract that give the option buyer the right, but not the obligation, to buy a stock, bond, commodity or other asset or instrument at a specified price (strike price) within a specific time period. The buyer would thus earn money when the actual value of the underlying asset is above the strike price. His payoff can be modelled as the maximum of the actual value A minus the strike price κ and 0. Or mathematically, $\max(A - \kappa, 0)$. A put option is a contract giving the purchaser the right, but not the obligation, to sell an underlying security at a predetermined price within a specified time frame. The purchaser would thus earn money when the value of the underlying asset is below the strike price. This payoff can be modelled as the maximum of the strike price minus the actual value and 0. Or in formulae $\max(\kappa - A, 0)$.

When the underlying asset is inflation, we call them inflation options and these consist of inflation-indexed caplets (IIClt) and floorlets (IIFlt). For a caplet, the buyer receives a payment when the inflation rate is above the strike price. For a floorlet, the buyer receives a payment when the inflation rate is below the strike price. Hence, an IIFLt is a put option on the inflation rate, whereas an IICLt is a call option on this same inflation rate. Recall that we want to price inflation-linked liabilities with an embedded floor at 0% inflation. Hence, the inflation floor, and then especially with a strike of 0%, is an interesting product since it is the most similar available product to the one we want to price. Therefore, we focus on this product. Using the payoff structure of the put option described above and that A is the actual inflation rate $\frac{I_t}{I_{t-1}} - 1$ and N is the nominal value, the payoff of the IIFlt at time t is defined as:

$$N\psi_i\Big[\kappa - \Big(\frac{I_t}{I_{t-1}} - 1\Big)\Big]^+, \tag{5}$$

An inflation floor is a stream of floorlets. Therefore, a floor with with maturity M has the payoff of the IIFlt in equation (5) at each time t for $t = 1, 2, ...M$.

3 Methodology

This section explains how the model that was designed for this thesis is used to value the inflation-linked liabilities with the embedded floor. First, we outline the model in section 3.1. Secondly, we explain how our model is calibrated in section 3.2. Finally, section 3.3 shows how our complete model is used to market-consistently price the inflation-linked liabilities.

3.1 The model

Both market and short rate models show (dis)advantages. Contrary to short rate models, market models model the entire curve and explicitly model observable quantities like year-on-year inflation rates and provide more calibration flexibility because of their dimensionality. However, there is also a downside to this particular aspect. The calibration of market models could become unstable in less liquid markets due to a lack of calibration instruments (van Haastrecht and Pelsser (2011)). In other words, a market model is the more advanced model, but may not be best option in a less liquid market due to possible calibration problems. Furthermore, the need for an advanced model is less since the exposure to inflation risk is much smaller than to for example interest rate risk. Taking these aspects into account, we consider a short rate model as the better option to model inflation risk. In particular, because of its analytical tractability and the possibility of tackling the shortcomings by using suggestions from the discussed literature, the JY model is considered as a good basis for modelling the inflation. The model in this thesis indeed uses the JY model as a basis, but does adapt it in some important ways by combining it with a nominal interest rate risk model. The upcoming section 3.1.1 will describe the JY model, section 3.1.2 discusses the particular nominal interest rate risk model used in this thesis and lastly section 3.1.3 explains how these two models are combined to reach our own model.

3.1.1 The JY model

The JY model uses the foreign currency analogy so that nominal dollars correspond to domestic currency, real dollars to foreign currency and inflation index to the spot exchange rate. In that setup, the fluctuations of the real and nominal interest rates and the inflation rate will be correlated. Using the notatation of Jarrow and Yildirim (2003), the specification under the real-world probability space of the nominal and real forward rates, f_n and f_r, respectively and the inflation index I, is as follows:

$$df_n(t,T) = \alpha_n(t,T)dt + \zeta_n(t,T)dW_n^P(t); \tag{6a}$$

$$df_r(t,T) = \alpha_r(t,T)dt + \zeta_r(t,T)dW_r^P(t); \tag{6b}$$

$$dI(t) = I(t)\mu(t)dt + \sigma_I I(t)dW_I^P(t), \tag{6c}$$

where (W_n^P, W_r^P, W_I^P) is a $\mathbb{P}$-Brownian motion with correlations $\rho_{n,r}, \rho_{n,I}$ and $\rho_{r,I}$; α_n, α_r and μ are the drift processes, ζ_n and ζ_r are deterministic functions and σ_I is a positive constant.

The real rate that is modelled is the expected real rate. The actual real rate is only known when the CPI in the corresponding period is known. Therefore, there is no redundancy in modelling the expected real rate together with inflation index and nominal rate. As can be seen, the inflation has its own Brownian motion influencing the dynamics, meaning that Jarrow and Yildirim (2003) believe that the inflation is not just determined by the real and nominal interest rate, but also by other factors that can influence the inflation rate like for example money supply, the exchange rate, and GDP.

These evolutions are abitrage-free and the market is complete if there exists an unique probability measure Q such that $\frac{P_n(t,T)}{B_n(t)}$, $\frac{I(t)P_r(t,T)}{B_n(t)}$ and $\frac{I(t)B_r(t)}{B_n(t)}$ are $\mathbb{Q}$-martingales. According to Jarrow and Yildirim (2003), this is true if and only if the following holds:

$$\alpha_n(t,T) = \zeta_n(t,T)\left(\int_t^T \zeta_n(t,s)ds - \lambda_n(t)\right); \tag{7a}$$

$$\alpha_r(t,T) = \zeta_r(t,T)\left(\int_t^T \zeta_r(t,s)ds - \sigma_I(t)\rho_{rI}\lambda_r(t)\right); \tag{7b}$$

$$\mu_I = r_n(t) - r_r(t) - \sigma_I(t)\lambda_I(t), \tag{7c}$$

where we have that, by Girsanov's theorem, there exists a market price of risk (λ) such that: $\tilde{W}_k(t) = W_k(t) - \int_0^t \lambda_k(s)ds$ for $k \in \{n, r, I\}$ are $\mathbb{Q}$-Brownian motions.

The first restriction is the arbitrage-free restriction as in the original nominal interest HJM model. The second one is the analogous abitrage-free forward rate drift restriction for real forward rates. The third restriction is the Fisher equation relating the expected nominal rate to the expected real rate and the expected inflation rate. Using these restrictions and the fact that $d\tilde{W}_n(t) = dW_n(t) - \lambda_n(t)$, we write the evolutions under the $\mathbb{Q}$-measure as:

$$df_n(t,T) = \zeta_n(t,T)\int_t^T \zeta_n(t,s)ds + \zeta_n(t,T)d\tilde{W}_n(t); \tag{8a}$$

$$df_r(t,T) = \zeta_r(t,T)[\int_t^T \zeta_r(t,s)ds - \rho_{r,I}\sigma_I(t)]dt + \zeta_r(t,T)d\tilde{W}_r(t); \tag{8b}$$

$$dI(t) = I(t)[r_n(t) - r_r(t)]dt + \sigma_I I(t)d\tilde{W}_I(t). \tag{8c}$$

In his paper, Mercurio (2005) shows that the above specification can also be written in the equivalent short rate formulation:

$$dn_{JY}(t) = [\theta_n(t) - a_n n(t)]dt + \sigma_n d\tilde{W}_n(t); \tag{9a}$$

$$dr_{JY}(t) = [\theta_r(t) - \rho_{r,I}\sigma_I\sigma_r - a_r r(t)]dt + \sigma_r d\tilde{W}_r(t); \tag{9b}$$

$$dI_{JY}(t) = I(t)[n_{JY}(t) - r_{JY}(t)]dt + \sigma_I I(t)d\tilde{W}_I(t), \tag{9c}$$

where $(\tilde{W}_n, \tilde{W}_r, \tilde{W}_I)$ are Q-Brownian motions with correlations $\rho_{n,r}$, $\rho_{n,I}$ and $\rho_{r,I}$; $\theta_x(t) = \frac{\partial f_x(0,t)}{\partial T} + a_x f_x(0,t) + \frac{\sigma^2 x}{2a_x}(1 - e^{-2a_x t})$, for $x \in \{n, r\}$. The "JY" underscript is to clarify that these are

JY projections.

We briefly mentioned the model parameters before, but we believe that a more elaborate explanation is necessary to intuitively understand the model. A first important note is that all the expressions consist of a drift (the term before dt) and a diffusion term (the term before $d\tilde{W}$). The drift is the value the model moves around. The diffusion term determines by how much the model moves around this drift. The parameters a_r and a_n are called the mean reversion speed parameters. These show how quickly the values of the variable of interest return back to the initial model curve. Therefore, the parameters a_r and a_n affect the behaviour of the real short rate and nominal rate, respectively. As Moysiadis et al. (2019) states it: "A small value would produce more trending simulation paths, while a larger value can result in steady evolution of interest rate" (p.10). θ_r and θ_n are the functions to exactly fit the term structures of the real and nominal short rate. Hence, these parameters, together with the mean reversion speed parameters, make sure that we get our initial model curve back of the variable of interest. So indeed, these two parameters make up the so-called drift for the interest rates. Do note that there is a small correction term of $\sigma_r \sigma_i \rho_{r,i}$ for the real interest rate. For the inflation rate, the drift is defined by $n - r$.

The fact that the model moves around the drift and is not exactly equal to this value is because of the σ_y for $y \in \{n, r, i\}$ parameters in our model. These parameters are the diffusion terms and reflect the volatility in parameter y. So, for example σ_r is the volatility in the real interest rate. In equation (9b), we can see that it is multiplied by the Brownian motion and it thus influences how much the real short rate changes per one step forward in the simulation. The last set of parameters are the correlations. Since these parameters determine how the Brownian motions affect each other, they have an influence on the diffusion term.

3.1.2 Nominal LMM with displaced diffusion

Many insurance companies have a lot more exposure to interest rate risk compared to inflation risk. It is therefore safe to assume that they have some model to capture this interest rate risk. By assuming that these companies have an interest in modelling inflation as well, the need for a compatible inflation model arises. Therefore, to make it as practical as possible, this thesis extends an interest rate risk model to a model that is also able to capture inflation risk. Do note that the focus in this thesis is on developing the inflation model and not on the interest rate risk model. Hence, some exact choices within that model are beyond the scope of this thesis and will not be motivated or explained in further detail. Rather, we just assume that the nominal interest rate risk model used in this thesis is used by the company.

The nominal model used in our thesis is a LMM with displaced diffusion (LMM-DD). The LMM directly models the entire forward rate curve. This makes it intuitively easy to understand. The downside of the LMM is that it cannot produce negative rates, which is inconsistent with the market nowadays. This can, however, be solved by adding a displacement. Under the LMM-DD, the forward rates are driven by the following stochastic differential equation:

$$dF_i(t) = \tilde{\sigma}_i(t)(F_i(t) + \alpha_{DD})dW_i^{Q_{i+1}}(t) \quad \text{for } \ i \in \{1, 2, ..., N\}, \tag{10}$$

where $F_i(t)$ is the forward rate for period (t_i, t_{i+1}) observed at time t; α_{DD} is the displacement parameter; $\tilde{\sigma}_i(t)$ is the volatility of forward rate in LMM-DD model; Q_{i+1} is the measure using $P(t, t_i + 1)$ as numeraire; $W_i^{Q_{i+1}}$ is the standard Wiener process under measure Q_{i+1} and $\rho_{i,j}$ is the linear correlation between increments of the i-th and j-th forward rates.

The problem in this situation, however, is that the number of drift terms is different for different forward rates. For the process to have a steady number of drift terms, the numeraire should move along with time. This is obtained by changing the numeraire to the spot measure. The process of $F_i(t)$ under the spot measure is defined as:

$$dF_i(t) = \tilde{\sigma}_i(t)(F_i(t) + \alpha_{DD})\left(\sum_{j=m(t)}^{i} \frac{\rho_{j,i}\tilde{\sigma}_i(t)(F_i(t) + \alpha_{DD})}{1 + F_j(t)} + dW_i^{Q_{i+1}}(t)\right). \tag{11}$$

In the LMM, a principal component analysis (PCA) is performed and the eigenvectors corresponding to the three largest eigenvalues are extracted. These eigenvectors explain most of the movements in the entire curve. In the end, the LMM is thus captured by three principal components (PCs), where each PC is driven by a Brownian motion.

The interest rate risk model used in this thesis is characterised by a good balance between complexity and forecasting performance and it will be satisfactory for medium-sized insurance companies. There do exist several other possibilities for a nominal interest rate risk model. As an example, one could use nominal short rate models. An important drawback of these models, however, is that they do not model the evolution of the whole curve. To model long term liabilities with a large exposure, a more advanced model is needed. Another possibility is to extend the LMM by adding stochastic volatility to achieve an improved fit to the market. We believe that this improved fit does not outweigh the added complexity and calibration time. Another alternative model is a SABR model, which is a stochastic volatility model. It attempts to capture the volatility smile in the derivatives market. The limitation of this model is that it is a single forward model. For a more complete overview and a more detailed description of the models, we refer to other articles/books, i.e. Brigo and Mercurio (2006).

3.1.3 Final Model

The final model is the JY model where the nominal evolution is replaced by the LMM. Furthermore, the real rate will be adjusted. The reason for this is that the projection of the real rate in the JY model is very much related to the nominal rate projection. If we replace the nominal projection by the projected based on the LMM, we need to adjust the real rate to preserve the relation the real and nominal rate of the initial JY model.

First, we replace the nominal JY projection by the nominal LMM projection. From the nominal forward curve projections, constructed using equation (11), the nominal short rate evolution that we define as $n_{LMM}(t)$ is retrieved. For every projection year, we construct a forward curve and use the forward rate belonging to the lowest maturity as a proxy for the instantaneous forward rate. From theory we know that this is equal to the short rate. Going through this procedure for all projection years yields the evolution of the short rate. This exact procedure will be clarified with an example in section 5.2.1.

Second, the real rate is adjusted in order to preserve a correct relation between the nominal and real interest rate. The LMM possesses log-normal dynamics, whereas the JY model is characterised by normal dynamics. If we left the real rate unchanged, the real rate would thus be normally distributed and the nominal rate log-normally. The inflation rate, depending on both these rates, would then be driven by a process containing some extreme numbers. Therefore, we decide to use a straightforward solution. Namely, we adapt the real rate in the following way. We use the nominal and real rate projection of the JY model from equations (9a) and (9b) and take the difference between these two projections: $n_{JY} - r_{JY}$. Hence, we make the assumption that, after the LMM projection replaces the nominal JY projection, the difference between the real rate and nominal rate stays the same as in the JY model. In this way, we reach a new real rate projection which we define as r_{LMM}. In formula: $r_{LMM} = n_{LMM} - (n_{JY} - r_{JY})$. Or we can rewrite it as $r_{LMM} = r_{JY} + n_{LMM} - n_{JY}$. Hence, our final model can be defined as:

$$dn_{JY}(t) = [\theta_n(t) - a_n n(t)]dt + \sigma_n d\tilde{W}_n(t); \tag{12a}$$

$$dF_i(t) = \tilde{\sigma}_i(t)(F_i(t) + \alpha_{DD})\left(\sum_{j=m(t)}^{i} \frac{\rho_{j,i}\tilde{\sigma}_i(t)(F_i(t) + \alpha_{DD})}{1 + F_j(t)} + dW_i^{Q_{i+1}}(t)\right); \tag{12b}$$

$$dr_{LMM}(t) = [\theta_r(t) - \rho_{r,I}\sigma_I\sigma_r - a_r r(t)]dt + \sigma_r d\tilde{W}_r(t) + dn_{LMM} - dn_{JY}; \tag{12c}$$

$$dI_{LMM}(t) = I(t)[n_{LMM}(t) - r_{LMM}(t)]dt + \sigma_I I(t)d\tilde{W}_I(t). \tag{12d}$$

As stated in section 3.1.2, the parameters of the LMM are not explained in more detail. The model parameters relating to the JY model are all explained at the end of section 3.1.1. An important aspect that changes, however, is the correlation structure. If we would have used the JY model as it was, there would have been three correlation variables, namely $\rho_{n,I}, \rho_{n,r}$ and $\rho_{r,I}$. In our situation we replace the nominal JY part by the LMM. Note that we do use the nominal JY part to adapt our new real rate. Therefore, that Brownian motion still influences our model. In other words, we have three extra Brownian motions belonging to the three principal components because of this replacement. Hence, our correlation matrix will be a 6 by 6 matrix with the variables $n, I, r, pc1, pc2$ and $pc3$. The exact correlation matrix is shown in section 5.1.

3.2 Model calibration

This section explains how the model is calibrated such that it is able to market-consistently price inflation products. A general observation for all variables is that we make a distinction between the short term and the long-term. Market products accurately capture, for example, the behaviour of the volatility, but only have a maturity up to 30 years. Pricing inflation-linked liabilities like pensions require a forecast of more than 30 years. This means that for the years higher than 30, we need some other way of estimating the parameters. Possibilities include historical data of the variable of interest or the targets of the European Central Bank (ECB). One can assume that, in the long-term, the variable of interest converges to these particular values and that these values should reflect some realistic number. Because of this, we decide to use the long-term mean as our initial value for all parameters, after which we adjust some parameter values by calibrating to market prices of inflation products. Section 3.2.1 describes how we initially determined the parameter values exactly. Section 3.2.2 explains how we adjust some of these parameters to reach an optimised model.

3.2.1 Initial parameters

Recall from section 3.1.1 that σ_y for $y \in \{n, r, i\}$ reflects how much variable y changes per one step in the model. A year as the step size in the model would thus reflect the yearly change in the particular variable y. In that situation, this parameter can be estimated by using a time series of variable y after taking the first differences on a yearly basis. Of those observations, we just calculate the volatility to obtain σ_y. Hence, we obtain σ_r, σ_n and σ_i by calculating the volatility of the first differences on a yearly basis of the real short rate, nominal short rate and inflation rate, respectively.

The mean reversion speed parameters a_n and a_r can be estimated in several ways. Possible estimation methods of this parameter are mentioned in the paper of Moysiadis et al. (2019). We opt for an regression procedure. To explain this procedure, we focus on a_n. Recall that we have $d_n(t) = (\theta_n(t) - a_n n(t))dt + \sigma_n dW_n(t)$. The solution of this model is generally known to be: $n(t) = e^{-a_n t}n(0) + \frac{\theta_n}{a_n}(1 - e^{-a_n t}) + \sigma_n e^{-a_n t}\int_0^t e^{a_n u}dW_n(u)$. This is the solution for $r_t|F_0$. We can, however, also take $r_t|F_{t-1} = r_t|r_{t-1}$ and write it as a regression. This leads to the following expression: $r(t) = \mu + \psi n(t-1) - \varepsilon(t)$, where $\mu = \frac{\theta_n}{a_n}(1 - e^{-a_n})$, $\phi = e^{-a_n}$ and $\varepsilon \sim N(0, \frac{\sigma_n^2}{2a_n}(1 - e^{-2a_n}t))$. From this regression formula, it follows that $a_n = -log(\phi)$. Hence, we use the linear regression coefficient to determine the mean reversion speed. For a_r, the only aspect that differs from a_n is that the solution of the Hull-White model is defined as $d_r(t) = (\theta_r(t) - \sigma_i\sigma_r\rho_{r,i} - a_r r(t))dt + \sigma_r dW_r(t)$. Note, however, that this only changes the value of μ in the linear regression formula and that we can thus still use the regression coefficient in the same way as we did for a_n.

Concerning θ_x for $x \in \{n, r\}$, we have that, as stated in section 3.1.1, $\theta_x(t) = \frac{\partial f_x(0,t)}{\partial T} + a_x f_x(0,t) + \frac{\sigma^2 x}{2a_x}(1 - e^{-2a_x t})$,, where $f_x(0,t)$ is the forward rate of variable x observed at time 0 for instantaneous borrowing at maturity t. This means that θ_x is a parameter that does not need to be estimated on its own. Rather, we plug in the estimated parameters a_r and σ_r and use the forward curve of variable x to determine $f_x(0,t)$ for t = 1,2,..,T.

Regarding the correlations in our final model, we notice that the three Brownian motions of the JY model are multiplied by their respective volatilities and are therefore closely linked to these variables; see formulas 9a until 9c. Therefore, we also use the first differences of the data used for the volatility parameters. So, for example, to calculate $\rho_{i,r}$, we calculate the correlation between the first differences of the real short rate and the first differences of the inflation rate. For the correlations relating to the LMM Brownian motions, we use the historical values of the principal components. So, for example, $\rho_{i,pc1}$ is the correlation between the first differences of the inflation rate and the historical values of the first principal component. By construction, the three principal components have a correlation of zero. Based on this analysis, we derive a particular correlation matrix which is shown in table 3 in section 5.1.

3.2.2 Optimising the parameters

As mentioned, we calibrate our model on market prices of several inflation products. More specifically, we use the ZCIIS, YYIIS, 0% floor and 2% floor. To that end, we first need to obtain model prices so that we can compare these to the market prices. There are two possibilities to do so. Either we use pricing expressions or we use a simulated-based method. The pricing expressions for the JY model are given in the paper of Mercurio and Moreni (2006). However, since we replace the nominal short rate by an LMM, these pricing expression are not longer valid. Although a lot of effort and time was spent, it follows that it is hard to make the specific adaptations to these pricing expressions to derive them ourselves for our own model. Therefore, we choose to use a simulation-based method to obtain model prices.

The simulation-based method uses several inflation paths that are in the end used to determine the model price. First of all, we plug in all initial parameter values in equations (12a) until (12d) and construct 1000 different inflation paths. Using the payoff structures of all four inflation products explained in section 2, we determine the payoff at each year in each path. One should see this as a multi-dimensional binomial tree. We then discount all these payoffs back in order to obtain the price of the particular product we want to price. The model price of the product is then the average price of all 1000 paths.

After determining our model price, we compare this to the market price and use some optimise target and optimise algorithm to reach this target. For the optimise function we opt for the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm. The BFGS method belongs to quasi-Newton methods, a class of hill-climbing optimization techniques that seek a stationary point of a function. In addition, we implement boundaries to the parameters we want to optimise. We will specify these later on. As an optimise target, we use that we want to minimise the sum squared difference of the model- and market prices. We use three different sets of our inflation products. Specifically, we take the sum squared differences between the model and market prices for the 0% floor only, then of the ZCIIS, YYIIS and the 0% floor simultaneously and lastly of the 0% floor and 2% floor together. In all three situations, our goal is thus to minimise this difference by adjusting some of our model parameters.

Recall that the LMM is already calibrated to nominal market data. Furthermore, as explained, the LMM is not the focus of this thesis and we therefore take these parameters as given. Therefore, the only parameters of interest are the real interest rate volatility, the inflation volatility and the mean reversion speed of the real interest rate. We choose to only optimise our model by adapting the two volatility parameters. The other parameter are kept the same as their initial value. For these volatility parameters we make an important assumption. Namely, the first 30 years are captured by six different values. In other words, for the first 5 years we assume we have some volatility, then for the next 5 years a particular volatility and so on until the years 26 until 30. Since we have both the inflation volatility and the real interest rate volatility, we thus have 12 model parameters that we want to optimise with the boundary that all need to be positive. The results are shown in section 5.3.

3.3 Valuation of the inflation-linked liabilities

In this section we discuss step by step how the inflation-linked liabilities of the insurance company are priced. In section 3.2.2 we explained how we calibrate our model to three different sets of products, namely firstly the 0% floor only, secondly the ZCIIS, YYIIS and 0% floor simultaneously and thirdly on the 0% floor and 2% floor together. In this section, we focus on the calibration on the 0% floor only. The final parameter set belonging to that particular minimisation problem is used to construct our final set of inflation paths. To price our inflation-linked liabilities, we need to perform several calculations. First, we need to adapt these projected indices to incorporate the floor. This is explained in section 3.3.1. The remaining steps of the valuation are discussed in section 3.3.2.

3.3.1 Incorporating the floor in the inflation-linked liabilities

Since inflation is used for indexation, the process of adjusting a particular value based on the changes in the inflation rate, it is incorporated in the liabilities of many insurance companies. Inflation, however, varies over time and can even fall below 0%. To protect policy holders against this, the liabilities are not negatively indexed. Hence, the policy holder has the guarantee of the maximum of the actual inflation rate i and 0%. Although the 0% floor has the focus in this thesis, other floor values can be considered. To make it more general, we state that the guarantee of the policy holder is the maximum of the actual inflation rate and the floor value (f): max(i,f). We recognise this as the payoff of a floor with a strike price at f% as defined in section 2.3 and also price it like this.

When using a floor, there are several possibilities concerning the behaviour of the outgoing cash flows. We first explain the general construction considered in this thesis, after which we will explain this in more detail by looking at the 0% floor. In general, when considering a particular floor value, the policy holder has the guarantee of receiving at least that floor value as indexation. Hence, at every year t, the index incorporating the floor at year t is the maximum of the original index (IO) at year t and the index incorporating the floor (IF) at year $t-1$ multiplied by 1 plus the the floor value (f) in percentages. Mathematically, this can be denoted as $IF_t = \max(IO_t, IF_{t-1} * (1 + f/100))$, where $IO_0 = IF_0 = 100$.

To give the intuition behind this construction, we consider the 0% floor as an example. In case of deflation, the liabilities will not be decreased. A subsequent indexation will be given when the future inflation, if any, has exceeded the deflation. Hence, we keep the index constant if a deflation occurs and we only increase the index after the index later in time has surpassed the index value belonging to the year before the deflation occurred. In order to illustrate this construction, a graphical representation of one possible inflation path is shown. In figure 1a, we focus on the 0% floor. In this figure, the initial index is shown as the black line, whereas the red line refers to the adapted index incorporating the 0% floor as described above. A good example can be found in the year 32 where inflation is below 0%. Instead of decreasing the index, we keep the index constant and we increase the index in year 34 since it has surpassed the index value of year 31. Figure 1b shows the indices incorporating the floors at -2%, -1%, 0%, 1% and 2%.

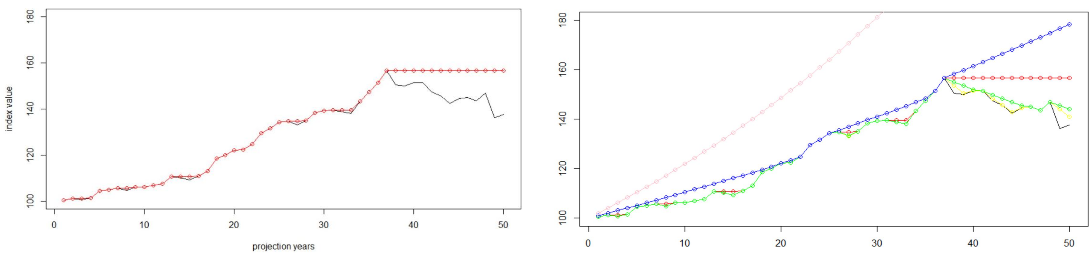

(a) Index behaviour 0% floor

(b) Index behaviour multiple floors

Figure 1: Index behaviour incorporating embedded floor

Note: This graph shows the effect on the index behaviour of incorporating a particular floor. To obtain the index incorporating the floor at a specific year, we take the maximum of the initial index at that specific year and the index incorporating the floor at the previous year multiplied by 1 plus the floor value in percentages. To clearly illustrate the effect of a 0% floor, the focus of our thesis, figure 1a only shows the effect of that floor. Figure 1b shows four different floor values to give a more complete understanding. The black line is our initial index. The yellow line, green line, red line, blue line, and pink line reflect how the index behaves incorporating a floor at -2%, -1%, 0%, 1% and 2%, respectively.

3.3.2 Pricing the floor embedded in the inflation-linked liabilities

After adapting our final paths of the inflation index in order to incorporate the floor, we are able to price the floor embedded in these liabilities. First, we index the cash flows using both the original index IO and the adapted index IF for all 1000 paths. We refer to the cash flows indexed by the original indexed as $CFIO$ and to the cash flows indexed by the adapted indexed as $CFIF$. Second, we discount both sets of cash-flows for all paths by using the nominal rate projection and take the sum of all discounted cash-flows to obtain the present value (PV) per path. Finally, we calculate the final present value by taking the average of all paths. The floor value (FV) is the result of the present value of cash flows indexed with the adapted index subtracted by the present value of the cash flows indexed with the original index: $FV = PV_{CFIF} - PV_{CFIO}$.

4 Data

This section provides an overview of all the data used in this thesis, in which we require data for calibrating the model and for pricing the inflation-linked liabilities. For the latter, we only need cash flows. For the calibration of the model, we need several data sets. For the initial estimation of the volatility parameters and the mean reversion speed parameters, we need a time series of the particular variable. For θ_n and θ_r we then already know σ_n and a_n and σ_r and a_r and only need to specify the forward rates for which we need the forward curves. Lastly, for the correlations we then only require the historical data of the principal components as extra data. For optimising the data we only need market prices. We refer to section 3.2 where we explained all of this in detail. A clear overview of all the data necessary to calibrate the model is given in table 1 below.

parameter	determining initial value	optimising value
σ_n	time series nominal interest rate	none
σ_i	time series inflation rate	market prices
σ_r	time series real interest rate	market prices
a_n	time series nominal interest rate	none
a_r	time series real interest rate	none
θ_n	nominal forward curve	none
θ_r	real forward curve	none
correlations	historical data principal components	none

Table 1: Data table

Note: This table shows all the data necessary to calibrate the model.

The upcoming section elaborate on the data, where section 4.1 first states some model assumption relating to the data used in this thesis. Section 4.2 discusses the inflation data. The data needed for the nominal and real interest rate is reviewed in section 4.3. Section 4.4 analyses the data used to construct the nominal and real forward curve. The data of the principal components is considered in section 4.5. The market prices are stated in section 4.6. Finally, the cash flows of the inflation-linked liabilities are provided in section 4.7.

4.1 Data assumptions

Before going into the specific data used in this thesis, we need to specify three model assumptions relating to the data, namely the start date of our simulation, the projection horizon and the time steps of the model. In section 3, we kept it general so that one can carry out the same research with any data set. In this thesis, we use 31/01/2020 as our start date. The motivation for this is twofold. Normally, one would use the end of a quarter, but due to the corona virus we find those possible start dates unrealistic. Next to this, the end of the year is also a date that does reflect the most realistic scenario. As a projection horizon, we decide to use 50 years. In the model, we want to forecast a sufficient number of years to price long term liabilities. Furthermore, we want our model to make the distinction between the short term (less than 30 years) and the long term (more than 30 years). If the model can forecast 50 years, it is also able to forecast for example 80 years. Last, we use time steps of a year. Or, in analytical form, $\Delta t = 1$. One of the motivations for this choice is that seasonality is something that is very much present in inflation. When taking time steps of a month, this would have a significant effect. For example, the price of inflation products would be affected by the seasonality and this is not realistic. When taking time steps of a year, this seasonality effect is no longer an issue. We will also illustrate this in section 4. Another motivation is that the LMM uses years, making it easier to combine the two models.

4.2 Inflation rate

The inflation data is obtained from the official website of the European Central Bank (ECB), where we use the overall HICP index for the euro area. This is an inflation index (I) in which a weighted average is taken among the most common 19 European countries. The data set consists of monthly observations of the inflation index from 01/31/1996 to 05/31/2020 with 2015 as a base year (index 03/31/2015 is equal to 100). So, in total, the data set consists of 293 observations. In order to retrieve the inflation rate in percentages at time t relative to t-1 ($i(t, t-1)$) we use the following formula:

$$i(t, t-1) = \frac{I(t) - I(t-1)}{I(t-1)} * 100 \tag{13}$$

In our model, we opt for $\Delta t = 1$ and hence use yearly steps. Therefore, we take the yearly inflation rate. Because of this operation, we lose the first 12 observations. After we have calculated the yearly inflation rates, we take the yearly differences and lose 12 more observations. Hence, in the end, we are left with 269 observations of the yearly differences of the realised inflation rate. To illustrate the motivation to use $\Delta_t = 1$, the monthly inflation is shown in figure 2a and the yearly inflation is shown in figure 2b. In these two figures it can be seen that the monthly inflation is a lot less stable than the yearly inflation rate. The first differences of the yearly inflation rate are depicted in figure 2c.

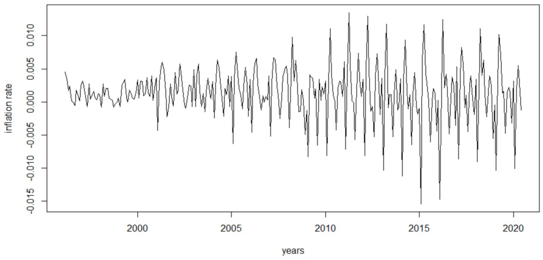

(a) Monthly inflation rate

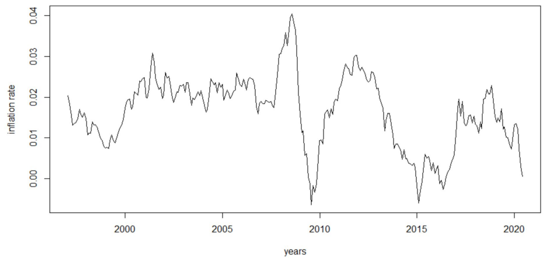

(b) Yearly inflation rate

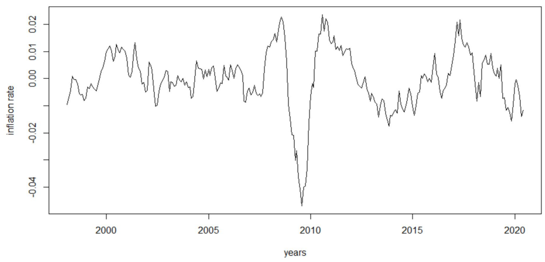

(c) Yearly differences of the yearly inflation rate

Figure 2: Inflation data

Note: This figure shows the inflation data we considered in this thesis. Figure 1a shows the monthly inflation rate, figure 1b the yearly inflation rate and 1c the yearly differences of the yearly inflation rate. These figures are all based on the data found at the ECB website.

4.3 Nominal and real interest rate

We retrieved the data of the nominal and real interest rate from the website of the Organisation for Economic Co-operation and Development (OECD). The nominal interest rate we take is the short term interest rate. Similar to the previous data set, we use the observations from 01/31/1996 to 05/31/2020. To obtain the real interest rate, we use the Fisher equation stating that $(1+n) = (1+i) * (1+r)$. If we expand this, we have $1+n = 1+i+r+ir$. Since the term ir is generally very small (percentage of a percentage), we can state the approximation $n = i+r$, or $r = n - i$. Hence, we subtract the realised inflation from the nominal interest rate to obtain the real interest rate. Subsequently, we take the yearly differences. The nominal interest rate is shown in figure 3a and the yearly differences of the nominal interest rate in figure 3b. The real interest rate is depicted in figure 4a, whereas the yearly differences of the real interest rate in figure 4b.

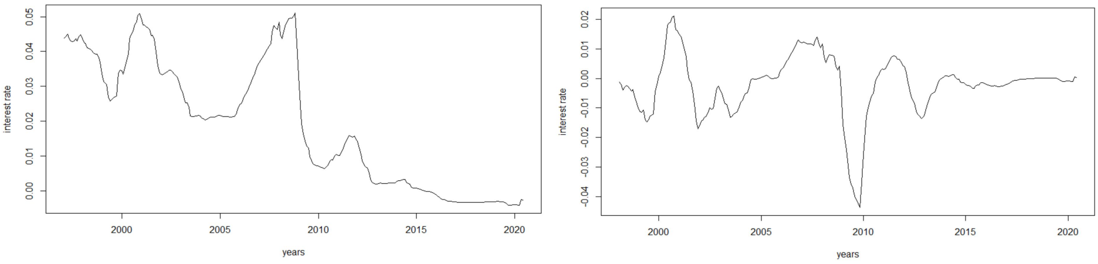

(a) Yearly nominal interest rate

(b) Yearly differences of yearly nominal interest rate

Figure 3: Nominal interest rate data

Note: This figure shows the nominal interest rate data we considered in this thesis. Figure 1a shows the yearly nominal interest rate, whereas figure 1b shows the yearly differences of that data. These figures are all based on the data found at the OECD website.

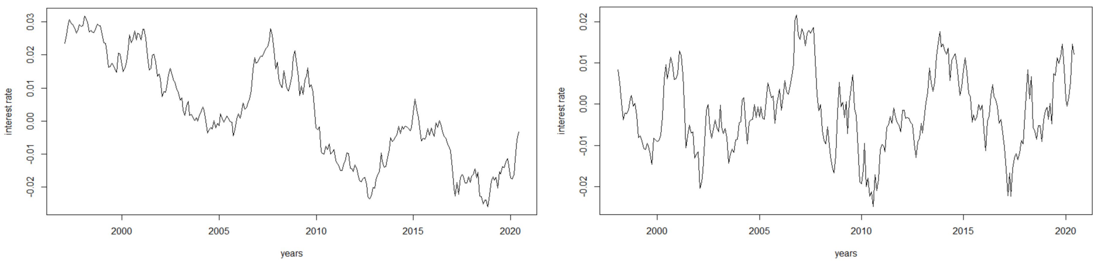

(a) Yearly real interest rate

(b) Yearly differences of yearly real interest rate

Figure 4: Real interest rate data

Note: This figure shows the real interest rate data we considered in this thesis. Figure 1a shows the yearly real interest rate, whereas figure 1b shows the yearly differences of that data. These figures are constructed using the Fisher equation and the nominal interest rate data and the inflation date.

4.4 Nominal and real forward curve

To construct the nominal forward curve and real forward curve, we use the zero rates of the nominal interest rate and the inflation rate. We obtained that data from Bloomberg. Using the term structures of the nominal interest rate and the inflation rate, we construct the term structure of the real zero interest rate. We then use that as a basis to retrieve the term structure of the real forward rate by making use of the following relation:

$$f_r(t) = \frac{(1 + r(t))^t}{(1 + r(t-1))^{t-1}} - 1 \tag{14}$$

Note that the real forward rate for year 1 is equal to the real zero rate for year 1. The constructed nominal forward curve and the constructed real forward curve are shown in figures 5a and 5b depicted below:

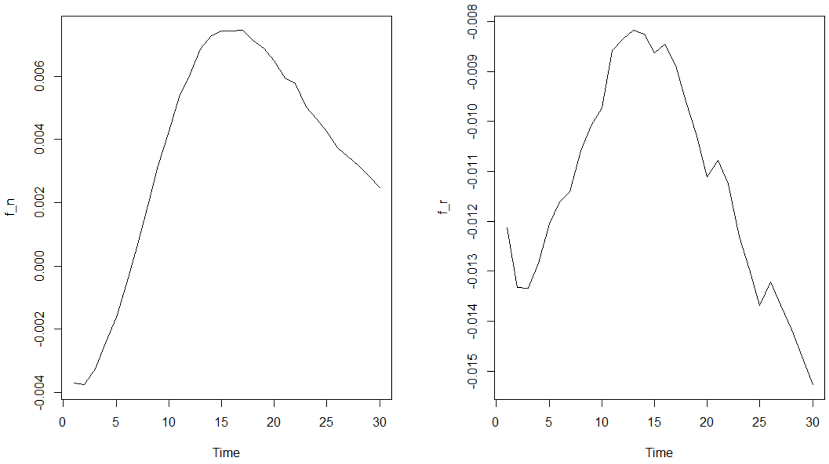

(a) Nominal forward curve (b) Real forward curve

Figure 5: Forward curves

Note: This figure shows the forward curves used in this thesis to estimate the parameters θ_n and θ_r. The nominal forward curves are constructed using equation (14), where the nominal zero rates are obtained from Bloomberg. For the real forward curve we use the same approach, but the zero rates follow from the nominal zero rate and inflation zero rate. The inflation zero rates are also retrieved from Bloomberg.

4.5 The principal components

To analyse the correlation, as stated at the end of section 3.1.3, we need historical information of the three principal components resulting from the LMM-DD. These are monthly observations and are from 31/01/2000 until 31/05/2020. Recall that we do not go into detail on how these are obtained since it is not the focus of our thesis. The data are shown in figure 6 below.

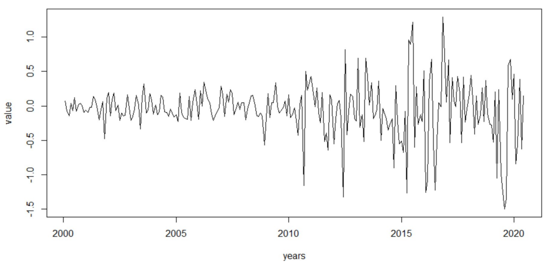

(a) Historical data PC1

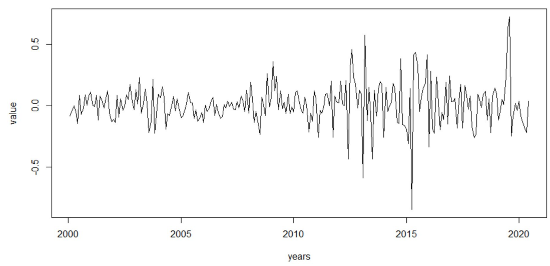

(b) Historical data PC2

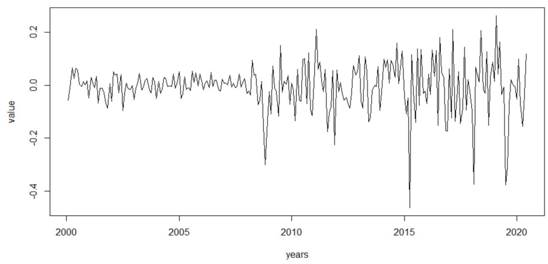

(c) Historical data PC3

Figure 6: Data principal components

Note: This figure shows the data used of the principal components (PC) to obtain the correlation matrix in our model. Figure 6a, 6b and 6c show the historical values of the first PC, second PC and third PC, respectively. For all these sub-figures, the data is from 31/01/2000 until 31/05/2020 and is obtained from the principal component analysis performed on the LMM. Recall that we do not go into detail on how these are obtained since it is not the focus of our thesis.

4.6 Market prices

The inflation products we use in our thesis are ZCIIS, YYIIS, 0% floors and 2% floors. The payoffs of these products are discussed in section 2. To obtain the market prices of ZCIIS, we use equation (2) and the inflation term structure. For the prices of the YYIIS, we use equation (4). The ratio of the inflation indices can be derived from the ZCIIS. Hence, to derive market prices for ZCIIS and YYIIS, we only need the data consisting of the inflation term structure. The market prices of the 0% and 2% floors have been extracted from Bloomberg. All market prices are shown in table 7 below, where the black line shows the ZCIIS prices, the red line the YYIIS prices, the blue line the 0% floor prices and the green line the 2% prices. The exact prices can be found in table 8 in Appendix A.

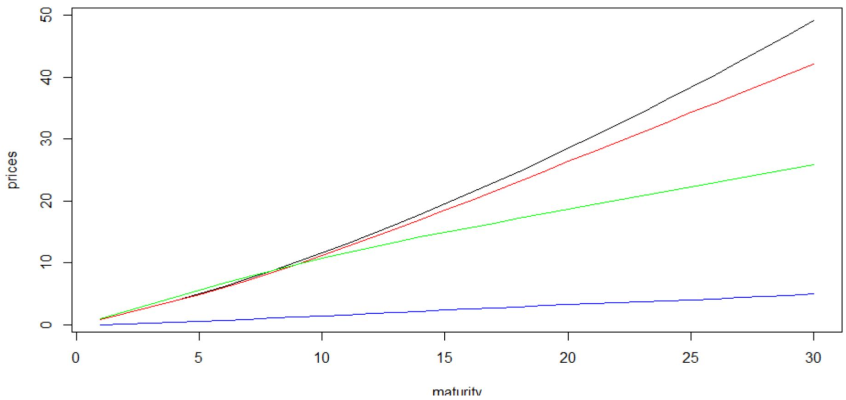

Figure 7: Market prices

Note: This graph shows the market prices of the ZCIIS (black), YYIIS (red), 0% floor (blue) and 2% floor (green). The market prices of the ZCIIS and YYIIS are constructed using the term structure of the inflation rate and equation (2) and (4), respectively. The market prices of the 0% and 2% floors have been extracted from Bloomberg.

4.7 Cash flows liabilities

To valuate the embedded floor in the inflation-linked liabilities, we need the cashflows of the inflation-linked liabilities. As mentioned, we use a dummy portfolio approximating the inflation-linked liabilities. This dummy portfolio consists of six different contracts, where each contract has cash flows for 50 years. These cash flows are shown in figure 8 below, where the exact numbers for the first 30 years are provided in table 12 in Appendix A. Note that the behaviour of the cash flows differs among the several contracts. This is done on purpose to investigate whether or not this affects the floor value and if so by how much.

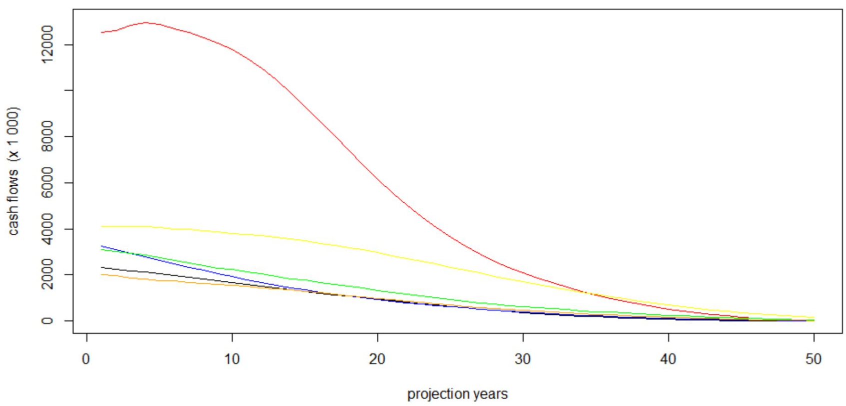

Figure 8: The cash flows of the inflation-linked liabilities

Note: This graph shows the cash flows of the dummy portfolio approximating the inflation-linked liabilities, where we use six different contracts. Contract 1 to 6 are shown as the black line, red line, blue line, green line, yellow and orange line, respectively.

5 Results

This section states the results obtained in this thesis. Section 5.1 shows the results of the underlying parameters. The projections of the nominal and real interest rate and the consequential inflation rate using these parameter values are considered in section 5.2. Section 5.3 analyses our calibration results. Finally, section 5.4 discusses the results of the valuation of the inflation-linked liabilities.

5.1 Results underlying parameters

This section states the parameter values used for the projections of the nominal interest rate, the real interest rate and the inflation rate. In our model, we have three volatility parameters and two mean-reversion parameters we need to estimate. Furthermore, we have the correlation matrix that shows the correlation between the six Brownian motions. In the table below, we first show the values of the volatility and mean-reversion parameters. The volatility parameters are based on time series data and the mean reversion speed parameters follow from the regression coefficient, where the regression is also based on the time series data. We refer to section 3.2.1 for the detailed explanation on how these parameters are estimated.

parameter	value
σ_n	0.0103
σ_r	0.0093
σ_i	0.0108
a_n	0.82
a_r	0.88

Table 2: Initial values volatility and mean reversion speed parameters

Note: This table shows the initial values for the volatility and mean reversion speed parameters before optimising the model. The volatility parameters are based on time series data and the mean reversion speed parameters follow from a regression coefficient, where the regression also uses the time series data. We refer to section 3.2.1 for the detailed explanation on how these parameters are estimated.

It thus follows that all three volatility parameters are equal to approximately 1%. This means that the diffusion around the drift in our model is not that high. The values of the mean reversion speed parameters for the nominal interest rate and real interest rate are equal to 0.82 and 0.88, respectively. Since a mean reversion speed parameter is always between 0 and 1, these values are quite high. This means that our interest rates quickly return back to the initial model curve. Given the values of the volatility parameters and mean reversion speed parameters, we can thus conclude that we have a rather steady evolution of the interest rates. For the inflation rate the same holds since it is driven by these two interest rates. Note that this statement refers to the situation before we replace the nominal rate projection by the LMM approximation. After this replacement, the projection will be based on the LMM which possesses lognormal dynamics, which automatically implies a less steady evolution.

The correlation matrix is constructed using the time series data and the historical values of the principal component analysis. Again, we refer to section 3.2.1 for the detailed explanation. Using that approach, the correlation matrix is analysed to be as follows:

	n	i	r	PC1	PC2	PC3
n	1	0.2409	0.3069	0.0683	-0.1875	0.1291
i	0.2409	1	-0.8498	0.1786	-0.0669	0.1171
r	0.3069	-0.8498	1	-0.1379	-0.0362	-0.0447
PC1	0.0683	0.1786	-0.1379	1	0	0
PC2	-0.1875	-0.0669	-0.0362	0	1	0
PC3	0.1291	0.1171	-0.0447	0	0	1

Table 3: Correlation matrix

Note: This table shows the six by six correlation matrix for the Brownian motions in our model belonging to the nominal rate (n), the inflation rate(i), the real rate(r), the first principal component (PC1), second principal component (PC2) and third principal component (PC3), respectively. It is constructed using the time series data and the historical values of the principal component analysis. We refer to section 3.2.1 for the detailed explanation. Using that approach, the correlation matrix is analysed to be as follows:

From the obtained correlation matrix, we notice that the first differences of the yearly real interest rate is positively correlated with the first differences of the nominal interest rate and negatively correlated with the first differences of the inflation rate. This makes intuitive sense since, in theory, the real rate is equal to the nominal rate minus the inflation rate. Furthermore, it follows that the nominal interest rate is positively correlated with inflation. This also makes intuitive sense since that if interest rates are lowered, people tend to spend more money instead of saving which results in an increase in the prices. One would expect to the see the same relation between the principal components capturing the nominal interest rate and the real interest rate and inflation rate. This is, however, not true. It could be the case that the relation is disturbed because of the many operations applied when performing a principal component analysis. The principal components have a correlation of zero by construction. Lastly, the variables of course have a correlation of 1 with themselves. We conclude that the Brownian motions of the principal components have a small effect on the other Brownian motions. The opposite is true for the Brownian motion of the real interest rate and the inflation rate since their correlation is close to -1.

5.2 Results interest rate and inflation rate

This section analyses the interest rate projections from the initial JY model, the replaced nominal rate projection from the LMM and the adapted real interest rate following from this replacement.

5.2.1 Nominal and real interest rate

In section 3.1.3, we explained how our nominal short rate evolution is retrieved from the nominal forward curve constructed by the LMM. The LMM projects a forward curve for all 50 projections years and for each curve we take the shortest maturity. To clarify, we use an example in which we analyse the realised evolution of the first path, which is shown in figure 9d. To roughly show that it is constructed in the way we described, we take the forward curves of projection year 1, year 25 and year 50 for the first path. These are shown in figures 9a to 9c. As one can see, the short rate evolution begins with the value of around -0.005, has a value at year 25 of -0.028 and ends with a value of approximately -0.01. These correspond to the zero rate of the shortest maturity of the zero curves of projection years 1, 25 and 50, respectively.

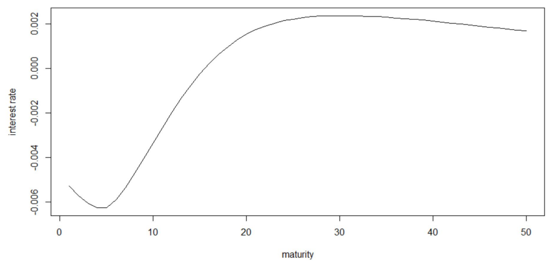

(a) Forward curve for projection year 1

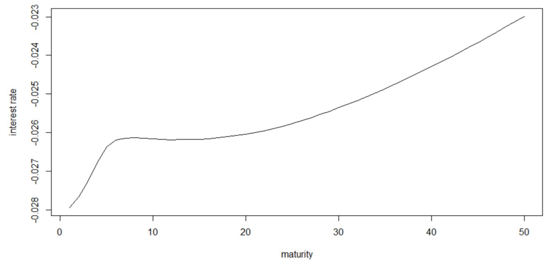

(b) Forward curve for projection year 25

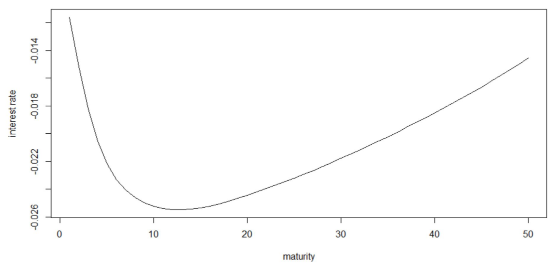

(c) Forward curve for projection year 50

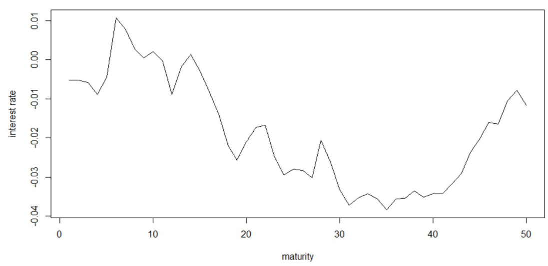

(d) Nominal short rate evolution

Figure 9: Approximation of nominal short rate evolution

Note: These graphs show how we use the LMM to approximate the nominal short rate evolution. Figures 9a, 9b and 9c show the forward curves for the projection years 1,25 and 50, respectively. The resulting nominal short rate approximation from all forward curves is shown in figure 9d. It is constructed by extracting the shortest maturity belonging to the forward curve for each projection year. This is explained in more detail in section 3.1.3

As mentioned before in section 3.1, the nominal rate and real rate are first projected using the JY model. The nominal rate, however, is replaced by the LMM and we therefore adapt the projected real rate in order to maintain a solid relationship between the real and nominal interest rate. To show how this changes our model, the distribution of the nominal and real rate for both the initial JY model and our own model can be found in figures 10 and 11.

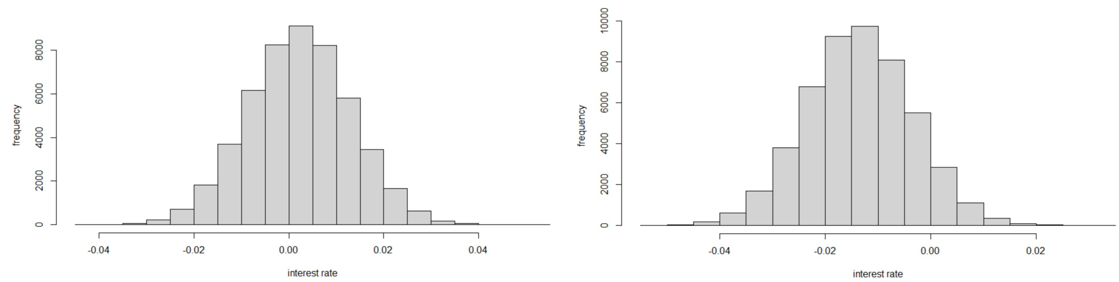

(a) Distribution of the nominal rate projected by the JY model

(b) Distribution of the real rate projected by the JY model

Figure 10: Distribution interest rates JY model

Note: This figure shows the distribution of the interest rates projected by the JY model. More specifically, figure 10a displays the nominal rate distribution and figure 10b presents the real rate distribution.

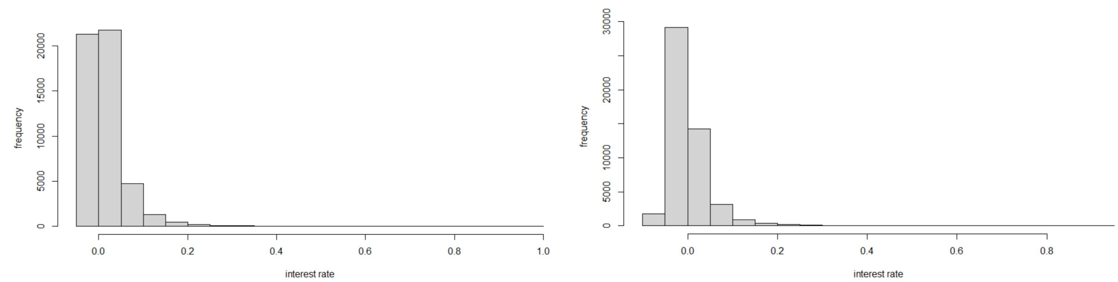

(a) Distribution of the nominal rate projected by the LMM model

(b) Distribution of the real rate projected by the JY model

Figure 11: Distribution interest rates after LMM replacement

Note: This figure shows the distribution of the interest rates after the LMM replacement. More specifically, figure 11a displays the distribution of the projected nominal rate by the LMM and figure 10b presents the distribution of the adapted real rate.

Recall that we stated that the projection of the interest rates would be a rather steady evolution because of our volatility and mean reversion speed parameter values. This statement is supported by the distributions shown in figure 10. After the LMM replacement, the interest rate projection is indeed more fluctuating as mentioned before. This was exactly the reason why we adapted the real rate. The relation between the nominal rate projected by the LMM and the projected real rate of the JY model would not be the correct one and would lead to an inflation projection consisting of extreme numbers.

5.2.2 Inflation rate

Our final nominal and real rate projections, meaning the projected nominal rate of the LMM and the adapted real rate, are used to make projections of the inflation rate. We refer to formula 12d for the exact specification. To give an idea of our forecast inflation rate, the resulting distribution can be found in figure 12.

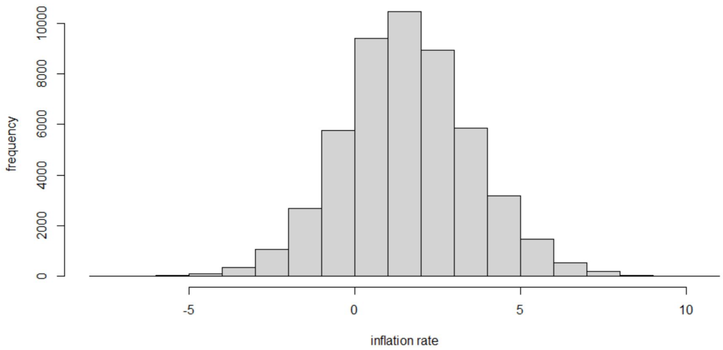

Figure 12: Distribution inflation rate

Note: This graph shows the distribution of the inflation rate using the nominal rate projection of the LMM and the adapted real rate projection.

The result we want to highlight is that the resulting inflation rate projection is a steady evolution without many extreme values. This is achieved because we adjusting our real rate projection after the nominal rate projection replacement. This is exactly what we wanted to achieve since we want our model to resemble the JY model.

5.3 Calibration results

This section analyses the performance of our model calibration. As mentioned in section 3.2, we try to minimise the squared sum difference between model and market prices for our four inflation products by adapting 12 volatility parameters. Six parameters for the real rate and six parameters for the inflation rate, where each value reflects a time-span of 5 years. Since our goal is to price inflation-linked liabilities with a floor at 0%, we first try to minimise the difference between these model prices and market prices. We call this fit 1. The results are shown in figure 13 below. The exact numbers can be found in table 10 in Appendix A.

For the interested reader, the results obtained when using our initial parameters are shown in table 9 in Appendix A. This serves to demonstrate the performance of our model when using only the time series estimation discussed in section 3.2.1. Further, by comparing our optimised results to these initial results, one can see how important it is to base the short term volatility on market prices rather than on historical data.

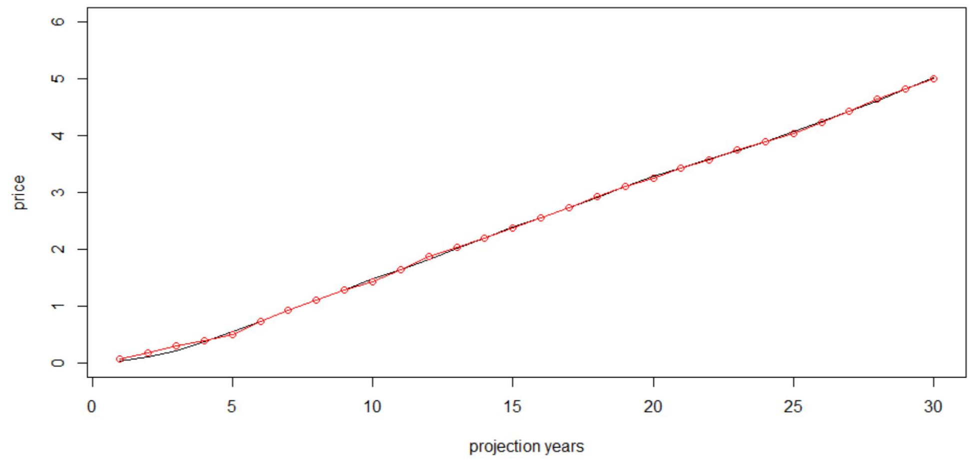

Figure 13: Fit 1 results 0% floor

Note: This graph shows the model and market prices of the 0% floor.

From this figure, it is clear that our model is able to approximate the market prices very closely by adapting the volatility parameters. Do note that this is mainly the case from maturity 6 onwards. Apparently, the model does struggle to match the prices in the very short term. An explanation could be that those prices do not contribute a lot to the sum squared differences since prices are low and that therefore the algorithm mainly focuses on the higher maturities. One could investigate whether the fit would improve if we split the volatility belonging to the first 5 years into 5 separate values. Overall, the conclusion is that our model is almost perfectly able to market-consistently price the 0% floor. The parameters that belong to fit 1 can be found in table 5.

To investigate whether the optimised parameters obtained by only fitting to 0% floors are able to market-consistently price other inflation products, we perform an out-of-sample test. Note that is a different out-of-sample test then is usually utilised. Normally, one would use the first so many maturities to calibrate the model and then test whether these obtained parameters are able to match the market prices for later maturities. In this thesis, we calibrate our model on a specific inflation product and use the obtained parameters to price another inflation product. More specifically, as an out-of-sample test in this thesis, we also calculate the ZCIIS, YYIIS prices and 2% floor prices using the parameter values obtained after calibrating to 0% floor prices. The result of this test can be found in figure 14. The exact prices are shown in table 11 in Appendix A.

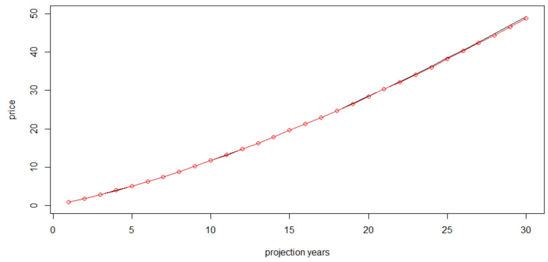

(a) ZCIIS prices

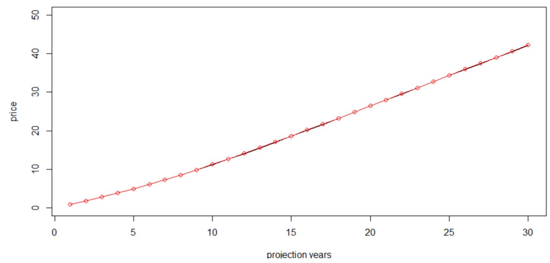

(b) YYIIS prices

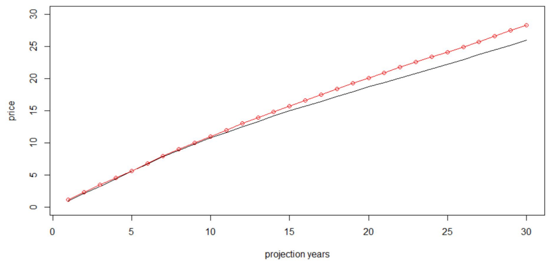

(c) 2% floor prices

Figure 14: Out-of-sample test results fit 1

Note: This figure shows the results of the out-of-sample test of fit 1. More specifically, it tells us how well we are able to market-consistently price other inflation products using the parameter results of fit 1. Figure 14a presents the ZCIIS prices, figure 14b displays the YYIIS prices and figure 14c shows the 2% floor prices.

From this figure it is clear that the ZCIIS and YYIIS prices are close to the market prices. To be precise, from table 11, it follows that the prices of the ZCIIS and the YYIIS have a deviation of maximum 1%. The 2% floor prices, however, deviate more from the market prices, especially for the higher maturities. This means that our first fit shows promising results concerning the ZCIIS and YYIIS. Unfortunately, the 2% floor prices do not show perfect results.

To investigate whether we can get any closer for the ZCIIS and YYIIS, we calibrate on the ZCIIS, YYIIS and 0% floor prices together. The results of this calibration is presented in table 4 below. We choose to show this in a table rather than a figure since the numbers are very close to each other, which means that a figure would not be able to clearly illustrate the difference. Recall that fit 1 is on the 0% floor only and fit 2 is on the three products together.

maturity	ZCIIS fit 1	YYIIS fit 1	0% floor fit 1	ZCIIS fit 2	YYIIS fit 2	0% floor fit 2
1	0.84	0.84	0.07	0.84	0.84	0.11
2	1.82	1.81	0.18	1.82	1.81	0.21
3	2.87	2.84	0.30	2.87	2.84	0.31
4	3.97	3.89	0.39	3.97	3.89	0.38
5	5.07	4.95	0.50	5.07	4.95	0.48
6	6.24	6.07	0.74	6.25	6.07	0.71
7	7.52	7.29	0.93	7.53	7.29	0.91
8	8.85	8.56	1.11	8.85	8.56	1.10
9	10.26	9.89	1.29	10.26	9.90	1.28
10	11.74	11.30	1.44	11.75	11.30	1.43
11	13.21	12.68	1.65	13.22	12.68	1.66
12	14.72	14.11	1.87	14.73	14.11	1.90
13	16.28	15.57	2.04	16.30	15.57	2.08
14	17.91	17.09	2.20	17.93	17.09	2.26
15	19.62	18.65	2.37	19.64	18.65	2.44
16	21.30	20.19	2.54	21.32	20.19	2.59
17	22.96	21.69	2.73	22.99	21.69	2.74
18	24.67	23.21	2.93	24.71	23.22	2.89
19	26.51	24.81	3.11	26.55	24.82	3.03
20	28.44	26.47	3.25	28.49	26.48	3.14
21	30.29	28.03	3.42	30.34	28.05	3.33
22	32.14	29.58	3.58	32.20	29.59	3.52
23	34.03	31.14	3.75	34.09	31.15	3.73
24	36.02	32.74	3.90	36.08	32.76	3.90
25	38.18	34.41	4.03	38.25	34.44	4.06
26	40.24	35.99	4.23	40.33	36.02	4.26
27	42.31	37.53	4.42	42.40	37.56	4.44
28	44.34	39.02	4.63	44.43	39.06	4.64
29	46.47	40.56	4.82	46.58	40.59	4.82
30	48.80	42.18	5.00	48.91	42.22	4.99

Table 4: Model prices based on two different fits

Note: This table shows the model prices of the 0% floor, ZCIIS and YYIIS for two different fits. Fit 1 refers to the calibration on 0% floors only, whereas fit 2 is the calibration on the 0% floors, ZCIIS and YYIIS simultaneously.

From the table, we can conclude that the prices of the ZCIIS and YYIIS come closer to the market prices, but that the 0% floor prices move a bit away from the market prices, especially for the first maturities. Again, the reason could be that, since the algorithm tries to minimise the sum squared difference, it automatically puts more weight on the products/maturities containing the higher prices. In this situation, the algorithm would focus on the ZCIIS and YYIIS prices.

Note that we did not include the 2% floor prices, but we can assume that these also do not change a lot. The parameters that belong to fit 2 are also displayed in table 5.

Finally, to find out if we can improve the 2% floor results, we calibrate on 0% floors and 2% floors together. The results of this calibration is shown in figure 15 below, where the parameter values belonging to this fit are also shown in table 5. Comparing these results to fit 1, or more specifically to figures 13 and 14c, we conclude that the 2% floor prices come closer to the market prices, but that the 0% floor prices deviate away from the market prices. More specifically, notice that the 2% floor prices were initially too high compared to the market prices. We are now able to overcome this issue, but the 0% floor prices are now too low compared to the market prices. The influence of the volatility on the model prices is apparently not exactly the same for these two products.

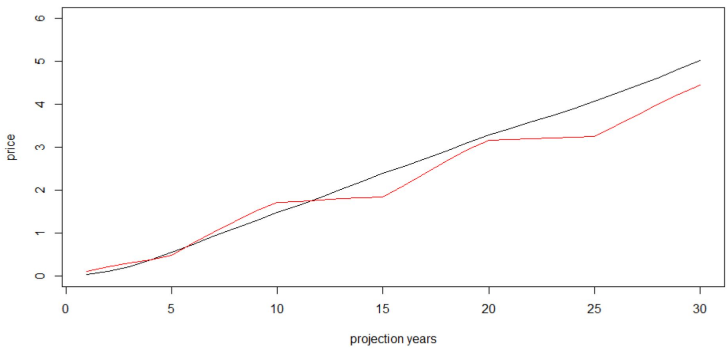

(a) 0% floor prices

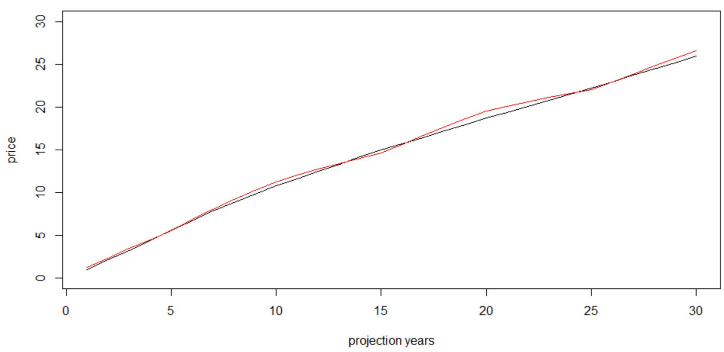

(b) 2% floor prices

Figure 15: Results of fit 3

Note: This figure shows the results of fit 3, where fit 3 refers to calibration on 0% floors and 2% floors simultaneously. Figure 15a present the results of 0% floor, whereas figure 15b the results of the 2% floor.

Our model thus struggles to market-consistently price the 0% floor and 2% floor at the same. Initially, we calibrated on 0% floors only and used the obtained parameter set to price the 2% floors as well. The result was that the model prices of the 0% floor prices very closely approximated the market prices, but that the model prices of the 2% floor prices deviated away from the market prices for higher maturities. The same result was obtained when calibrating on ZCIIS, YYIIS and 0% floors together. In these situations, however, we did not include the 2% floor prices in the calibration procedure, but only as an out-of-sample test. When calibrating on 0% floors and 2% floors, we find that the model prices of the 2% floor prices come closer to the market prices, but that the 0% floors are now further away from the market prices. The model is thus not able to find a volatility structure that works for both these products. This could be the case since the 2% floors are less liquid than 0% floors or that the volatility structure in our model is not advanced enough to capture the behaviour of both these products at the same time.

years	initial σ_r	initial σ_i	σ_r fit 1	σ_i fit 1	σ_r fit 2	σ_i fit 2	σ_r fit 3	σ_i fit 3
1-5	0.0093	0.0108	0.0000	0.0020	0.0000	0.0000	0.0000	0.0000
6-10	0.0093	0.0108	0.0000	0.0099	0.0000	0.0101	0.0000	0.0124
11-15	0.0093	0.0108	0.0000	0.0122	0.0000	0.0130	0.0032	0.0000
16-20	0.0093	0.0108	0.0007	0.0123	0.0000	0.0108	0.0000	0.0165
21-25	0.0093	0.0108	0.0019	0.0123	0.0000	0.0148	0.0037	0.0000
26-30	0.0093	0.0108	0.0034	0.0124	0.0000	0.0141	0.0037	0.0143

Table 5: Parameter values for the three different fits

Note: This table shows the volatility of the real rate and inflation rate in the initial situation and for our two different fits. For fit 1, these parameter values are obtained by using the BFGS algorithm and minimising the sum-squared difference of the model- and market prices of the 0% floor. For fit 2, we minimise the sum-squared difference of the model- and market prices of the 0% floor, ZCIIS and YYIIS together.

From table 5 we notice that the algorithm mainly adjusts the parameter value of the real rate volatility and the inflation volatility for the first five years. Especially the adjustment to the real rate volatility is striking. For the first fit, the volatility of the real rate is set to 0 for the first 15 years and for fit 2 it is even set to 0 for all 30 years. We can conclude two different things. Either our model is just able to price inflation products by only considering a volatility in the inflation rate or this was just the easiest way to reach an optimum. For both fits, the volatility structure of the inflation rate does feel intuitive in the sense that it mainly increases over time and does not make any weird jumps. The opposite can be stated of the volatility structure of fit 3. This supports the statement of before that the model struggles to match the 0% floor prices and 2% floor prices together.

5.4 The inflation-linked liabilities

This section reviews the valuation results of the inflation-linked liabilities. For this valuation, we choose to use the parameter values of fit 1, i.e. the set of parameters when minimising the sum squared differences between the model prices and market prices of 0% floors only. We do so since in that situation we focused on the 0% floor only and we have this exact same product embedded in the liabilities. Using the parameters belonging to fit 1, we can calculate the final projection of our inflation index for all paths. This index is adapted in the way discussed in section 3.3.1 to incorporate the 0% floor. These two indices are then used to price the floor value as explained in section 3.3 for all six contracts. The results are shown in table 6 below:

contract	floor value	average PV	floor value in % PV
1	123481.62	43897995.92	0.28%
2	864475.59	302008243.31	0.29%
3	150384.33	54280336.27	0.28%
4	201317.05	68179312.61	0.30%
5	225061.05	71234515.25	0.32%
6	140457.00	47279989.99	0.30%

Table 6: Floor value for all 9 contracts

Note: This table shows the floor value, the present value of the liabilities and the floor value in percentages of the present value and the 6 different contracts.

Recall from figure 8 that the behaviour of cash flows differs among the contracts. Even though this is the case, the floor value for all contracts is around 0.3% of the present value without the embedded floor.

In theory, pricing a floor 1 or 2 percent higher/lower is the same as pricing a 0% floor when the inflation rate would be 1 or 2 percent lower/higher. Therefore, to show the sensitivity of our analysis, we also show the results of the liabilities when incorporating a floor at -2%,-1%,1% and 2% In table 7, the floor value as a percentage of the present value is shown for these five floors.

contract	-2% floor	-1% floor	0% floor	1% floor	2% floor
1	0.01%	0.06%	0.28%	1.63%	10.07%
2	0.01%	0.06%	0.29%	1.67%	10.22%
3	0.01%	0.06%	0.28%	1.60%	9.69%
4	0.02%	0.07%	0.30%	1.66%	10.88%
5	0.02%	0.08%	0.32%	1.72%	12.29%
6	0.02%	0.07%	0.30%	1.67%	11.04%

Table 7: floor value in percentage of present value

Note: This table shows the floor value in percentage of the present value for the floors at -2%, -1%, 0%, 1% and 2%.

From the table, we see that the difference between the -1% and 0% floor is a lot smaller than the difference between the 0% floor and the 1% floor. This is even more so for the difference between the -2% floor and the 0% floor and the difference between the 0% floor and the 2% floor. This makes intuitive sense since floor only has a value when falling below 0%. Therefore, the result is less sensitive to an underestimation of the inflation rate than an overestimation of inflation rate.

6 Conclusion

This thesis has performed a market-consistent valuation of inflation-linked liabilities with an embedded floor at 0% inflation. Because of this 0% floor, a stochastic valuation was necessary. To perform this stochastic valuation, a new inflation model was created in this thesis. We used the JY model of Jarrow and Yildirim (2003), where we replaced the nominal part by a LIBOR market model with displaced diffusion (LMM-DD). Furthermore, since the LMM possesses log-normal dynamics, whereas the JY model is characterised by normal dynamics, we adapted the initial projected real rate of the JY model. To be more precise, we assumed that the difference between the nominal and real interest rate projection of the JY model stayed the same. Therefore, the new real rate was the projected nominal rate of the LMM minus this difference.

In order to market-consistently price the inflation-linked liabilities, we calibrated our model to market data. The particular LMM used in our thesis is already calibrated to nominal swap data. So, only the real rate and inflation component of our model needed to be calibrated. We decided to calibrate on several inflation products, namely ZCIIS, YYIIS, 0% inflation floors and 2% inflation floors. Since the JY model is adapted, the existing pricing expressions in previous literature were no longer valid and were intuitively hard to derive after this adaptation. Therefore, the calibration used a simulation-based approach. The optimise algorithm used in the calibration procedure is the BFGS algorithm with the target to minimise the sum squared differences of the model- and market prices. The parameters that could be changed by finding an optimum were the volatility of the real interest rate and the inflation rate, where we assumed that the volatility was split into blocks of 5 years.

This thesis finds that, by calibrating to 0% floors only, the model is able to almost perfectly match the market prices of 0 % floors. As an out-of-sample test, we use those parameters to value the ZCIIS, YYIIS and 2% floor. The outcome is that, compared to the market prices, the ZCIIS and YYIIS model prices have a maximum deviation of 1%. The 2% floor prices have a larger deviation from the market prices. To investigate if we could improve this result, we calibrated to ZCIIS, YYIIS and 0 % floor prices simultaneously. It followed that the 0% floor prices remain almost the same, whereas the ZCIIS and YYIIS prices did get closer to the market value, but only by a little. This fit was thus not significantly better than the first fit. As a last fit, we calibrated on the 0% floor and 2% floor. The result is that the 2% floor prices better reflect the prices observed in the market, but that the 0% floor prices deviated away from the market prices compared to the first and second fit. Our model thus struggled to market-consistently price the 0% and 2% floors simultaneously.

Since our focus was on estimating inflation-linked liabilities, with an embedded floor at 0%, we used the parameter set belonging to fit 1 to price the dummy portfolio. The result is that the value of the floor is around 0.3% of the present value without the floor. The sensitivity analysis showed that the valuation is sensitive to a negative change in inflation. More specifically, if the actual inflation turns out to be 1%/2% lower than projected, the floor value is equal to around 1.65% /11% of the present value. The overall conclusion is that the model is able to market-consistently price 0% floors, ZCIIS and YYIIS using a specific set of parameters, but that it does struggle to match the market prices of 0% and 2% floors simultaneously. Using the parameter set of the first 1, the value of the floor in our specific inflation-linked liabilities is 0.3% of the present value without the floor, but the result is quite sensitive to a change in inflation.

Although our model has shown very promising results, there are two main suggestions for further research. A limitation of the current model, for instance, is that it is built upon the assumption that the difference of the nominal and real interest rate projected by the JY model is kept constant when replacing the JY projection by the LMM projection. We did this to preserve a correct relation between the nominal and real interest rate, but it would be interesting to examine a more sophisticated solution when making this adaptation. Alternatively, one could investigate the possibility of combining the LMM with another inflation model rather than the JY model.

References

Baxter, M. and Rennie, A. (1996). *Financial Calculus: An introduction to derivative pricing.*

Belgrade, N., Benhamou, E., and Koehler, E. (2004). A market model for inflation. *SSRN Electronic Journal.*

Brigo, D. and Mercurio, F. (2006). *Interest Rate Models — Theory and Practice: With Smile, Inflation and Credit.*

Dam, H., Macrina, A., Skovmand, D., and Sloth, D. (2018). Rational models for inflation-linked derivatives. *SSRN Electronic Journal.*

Hinnerich, M. (2008). Inflation-indexed swaps and swaptions. *Journal of Banking Finance*, 32:2293–2306.

Jarrow, R. and Yildirim, Y. (2003). Pricing treasury inflation protected securities and related derivatives using an hjm model. *Journal of Financial and Quantitative Analysis*, 38(2):337–358.

Mercurio, F. (2005). Pricing inflation-indexed derivatives. *Quantitative Finance*, 5(3):289–302.

Mercurio, F. and Moreni, N. (2006). Pricing inflation-indexed options with stochastic volatility. *SSRN Electronic Journal.*

Mercurio, F. and Moreni, N. (2009). A multi-factor sabr model for forward inflation rates. *SSRN Electronic Journal*, pages 1–19.

Moysiadis, G., Anagnostou, I., and Kandhai, D. (2019). Calibrating the mean-reversion parameter in the hull-white model using neural networks: Methods and protocols. pages 23–36.

van Haastrecht, A. and Pelsser, A. (2011). Generic pricing of fx, inflation and stock options under stochastic interest rates and stochastic volatility. *Quantitative Finance*, 11(4):665–691.

Appendices

A Tables

Maturity	ZCIIS	YYIIS	0% floor	2% floor
1	0.84	0.84	0.03	0.99
2	1.82	1.81	0.11	2.12
3	2.86	2.82	0.22	3.24
4	3.94	3.88	0.37	4.40
5	5.04	4.93	0.56	5.59
6	6.21	6.05	0.74	6.72
7	7.48	7.26	0.93	7.83
8	8.81	8.51	1.10	8.82
9	10.21	9.84	1.29	9.80
10	11.69	11.23	1.49	10.76
11	13.15	12.61	1.65	11.64
12	14.66	14.03	1.83	12.52
13	16.24	15.49	2.01	13.34
14	17.87	16.99	2.20	14.17
15	19.56	18.54	2.40	15.01
16	21.23	20.05	2.56	15.73
17	22.97	21.59	2.73	16.46
18	24.76	23.17	2.91	17.20
19	26.61	24.77	3.10	17.95
20	28.54	26.40	3.29	18.71
21	30.39	27.94	3.43	19.40
22	32.30	29.50	3.58	20.10
23	34.27	31.08	3.73	20.80
24	36.31	32.68	3.89	21.52
25	38.43	34.30	4.06	22.24
26	40.46	35.83	4.24	22.96
27	42.55	37.37	4.42	23.69
28	44.71	38.92	4.61	24.43
29	46.92	40.49	4.81	25.17
30	49.23	42.07	5.02	25.93

Table 8: all market prices

Note: This table shows market prices of the ZCIIS, YYIIS, 0% floor and 2% floor. The market prices of the ZCIIS and YYIIS are constructed using the term structure of the inflation rate and equation 2 and 4, respectively. The market prices of the 0% and 2% floors have been extracted from Bloomberg.

maturity	ZCIIS	YYIIS	0% floor	2% floor	maturity	ZCIIS	YYIIS	0% floor	2% floor
1	0.84	0.84	0.10	1.22	16.00	20.83	19.98	5.70	20.20
2	1.80	1.79	0.54	2.69	17.00	22.48	21.47	6.01	21.27
3	2.83	2.80	1.02	4.17	18.00	24.14	22.96	6.31	22.30
4	3.90	3.83	1.50	5.66	19.00	25.94	24.55	6.60	23.29
5	4.99	4.88	1.94	7.10	20.00	27.83	26.20	6.84	24.22
6	6.13	5.99	2.38	8.52	21.00	29.65	27.75	7.09	25.19
7	7.38	7.20	2.77	9.84	22.00	31.48	29.29	7.33	26.14
8	8.69	8.45	3.17	11.17	23.00	33.34	30.84	7.57	27.09
9	10.06	9.77	3.56	12.44	24.00	35.30	32.43	7.79	27.98
10	11.51	11.16	3.88	13.62	25.00	37.45	34.11	7.99	28.80
11	12.95	12.54	4.25	14.84	26.00	39.47	35.66	8.25	29.74
12	14.42	13.95	4.62	16.05	27.00	41.53	37.21	8.49	30.65
13	15.93	15.40	4.89	17.12	28.00	43.50	38.68	8.77	31.61
14	17.52	16.90	5.17	18.16	29.00	45.60	40.22	9.02	32.52
15	19.17	18.45	5.44	19.18	30.00	47.90	41.83	9.25	33.40

Table 9: Initial model prices

Note: This table shows the obtained model prices when using the initial parameters.

maturity	optimised	market	maturity	optimised	market
1	0.07	0.03	16	2.54	2.56
2	0.18	0.11	17	2.73	2.73
3	0.30	0.22	18	2.93	2.91
4	0.39	0.37	19	3.11	3.10
5	0.50	0.56	20	3.25	3.29
6	0.74	0.74	21	3.42	3.43
7	0.93	0.93	22	3.58	3.58
8	1.11	1.10	23	3.75	3.73
9	1.29	1.29	24	3.90	3.89
10	1.44	1.49	25	4.03	4.06
11	1.65	1.65	26	4.23	4.24
12	1.87	1.83	27	4.42	4.42
13	2.04	2.01	28	4.63	4.61
14	2.20	2.20	29	4.82	4.81
15	2.37	2.40	30	5.00	5.02

Table 10: 0% floor prices fit 1

Note: This table shows the 0% floor prices of model fit 1.

maturity	optimized ZCIIS	market ZCIIS	optimised YYIIS	market YYIIS	optimized 2% floor	market 2% floor
1	0.84	0.84	0.84	0.84	1.19	0.99
2	1.82	1.82	1.81	1.81	2.34	2.12
3	2.87	2.86	2.84	2.82	3.46	3.24
4	3.97	3.94	3.89	3.88	4.53	4.40
5	5.07	5.04	4.95	4.93	5.62	5.59
6	6.24	6.21	6.07	6.05	6.81	6.72
7	7.52	7.48	7.29	7.26	7.91	7.83
8	8.85	8.81	8.56	8.51	8.99	8.82
9	10.26	10.21	9.89	9.84	10.01	9.80
10	11.74	11.69	11.30	11.23	10.97	10.76
11	13.21	13.15	12.68	12.61	12.00	11.64
12	14.72	14.66	14.11	14.03	13.02	12.52
13	16.28	16.24	15.57	15.49	13.93	13.34
14	17.91	17.87	17.09	16.99	14.82	14.17
15	19.62	19.56	18.65	18.54	15.70	15.01
16	21.30	21.23	20.19	20.05	16.60	15.73
17	22.96	22.97	21.69	21.59	17.52	16.46
18	24.67	24.76	23.21	23.17	18.42	17.20
19	26.51	26.61	24.81	24.77	19.26	17.95
20	28.44	28.54	26.47	26.40	20.04	18.71
21	30.29	30.39	28.03	27.94	20.89	19.40
22	32.14	32.30	29.58	29.50	21.73	20.10
23	34.03	34.27	31.14	31.08	22.56	20.80
24	36.02	36.31	32.74	32.68	23.34	21.52
25	38.18	38.43	34.41	34.30	24.05	22.24
26	40.24	40.46	35.99	35.83	24.91	22.96
27	42.31	42.55	37.53	37.37	25.73	23.69
28	44.34	44.71	39.02	38.92	26.62	24.43
29	46.47	46.92	40.56	40.49	27.45	25.17
30	48.80	49.23	42.18	42.07	28.23	25.93

Table 11: ZCIIS, YYIIS and 2% floor prices

Note: This table shows the out-of-sample test of fit 1.

maturity	contract 1	contract 2	contract 3	contract 4	contract 5	contract 6
1	2307763	12540714	3242109	3108796	4091563	2021697
2	2237276	12604543	3094330	3030934	4119349	1961985
3	2160021	12843602	2938833	2950270	4087378	1878038
4	2112988	12961137	2780614	2846864	4101860	1808571
5	2062612	12855328	2626487	2751797	4056684	1759252
6	1976282	12681374	2471634	2631686	3989097	1728606
7	1895655	12506535	2334319	2529999	3968687	1684263
8	1823987	12274112	2198037	2415613	3916228	1631373
9	1745828	12075118	2059774	2288914	3857198	1597581
10	1666088	11788756	1923127	2228579	3805274	1553554
11	1584482	11418705	1793690	2133286	3742216	1502959
12	1507744	10984195	1672324	2041916	3706821	1452984
13	1437460	10468200	1559233	1935541	3654181	1402841
14	1364844	9896511	1452495	1836545	3572256	1351376
15	1284646	9289743	1349700	1772242	3478346	1286732
16	1211456	8664473	1253585	1672531	3364604	1228003
17	1143184	8034039	1161945	1597638	3278407	1166229
18	1075386	7402221	1076092	1506320	3182890	1105191
19	1007985	6771908	995303	1419871	3091658	1040996
20	941271	6161733	920452	1330911	2972301	977748
21	875411	5580953	851068	1243241	2844163	915819
22	810779	5038821	786561	1155823	2718790	855770
23	747595	4531244	726385	1079295	2592999	797955
24	686334	4065403	669930	999652	2463986	741939
25	627223	3635449	616808	929231	2337070	688192
26	570521	3248767	566620	862075	2208251	636748
27	516515	2907542	519087	796506	2076665	587832
28	465114	2602935	474072	736834	1948048	541460
29	416491	2329749	431447	680106	1824446	497839
30	370491	2083861	391143	627291	1703025	456943

Table 12: cashflows for first 30 years and first 6 contracts