Erasmus University Rotterdam

MSc Econometrics and Management Science

Erasmus School of Economics

Pricing Inflation-Linked Derivatives by a Two-Process Hull-White Model: Comparing Different Calibration Methods

Author:

Mark Molenaar

Student Number:

576764

Supervisors:

dr. Hanno Reuvers - ESE

drs. Pieter Bouwknegt - NN

drs. Thijs Feenstra - NN

Second Assessor:

prof.dr. Michel van der Wel - ESE

August 29, 2022

Abstract

This paper investigates the risk-neutral pricing of inflation-linked caps and floors by a correlated two-process Hull-White model. Proper calibration is of crucial importance, but often overlooked in the literature. Therefore we lay special emphasis on different calibration methods, in which we investigate the use of regularization and weight functions to more accurately price a specified range of maturities. We develop a new method for capturing the volatility smile, by using a strike-dependent volatility parameter in the calibration step. On average, this correlated two-process HW model performs 0.92% better over a weighted range of 108 different products, compared to the uncorrelated two-process HW model as currently in use by NN. Moreover, by using a strike-dependent volatility, we are able to retain a full analytical solution, circumventing the need for numerical simulation. The single drawback of our method, is that the no-arbitrage conditions are violated when pricing exotic derivatives that depend on more than one strike price, in which case we have to resort to the original model.

List of Abbreviations

ATM	At the money
ATSM	Affine term structure model
CIR	Cox–Ingersoll–Ross
CPI	Consumer price index
Cpl	Caplet
EURIBOR	Euro interbank offered rate
Fll	Floorlet
HICPxT	Harmonized consumer price index excluding tobacco
HJM	Heath-Jarrow-Morton
HW	Hull-White
ILB	Inflation-linked bond
IR	Interest rate
ITM	In the money
JY	Jarrow-Yildirim
LIBOR	London interbank offered rate
OTC	Over the counter
OTM	Out of the money
RMSE	Root mean square error
TIPS	Treasury inflation-protected security
VAR	Vector autoregressive
ZBC	Zero-bond call
ZBP	Zero-bond put
ZCB	Zero-coupon bond

Contents

1 Introduction

For insurance companies, risk-neutral valuation of insurance liabilities is one of the corner stones for effective balance-sheet management. This includes hedging risks, and estimating future cash flows. It also plays an essential role in financial reporting to regulating agencies. For example in Solvency II, where inflation risk can have a significant impact on the solvency ratio when not properly hedged (David, 2008). Moreover, these new reporting standards require market-consistent reporting of the liabilities, underlining the need for risk-neutral valuation models.

In the past years, inflation levels have risen dramatically, increasing the exposure of insurance companies to inflation risk. In order to stay fully hedged, insurance companies need to acquire a larger amount of inflation-linked products, which has caused an increased interest for the valuation of these products. Some of the most used products for inflation hedging are inflation-linked swaps, caps and floors. These products provide a direct inflation exposure, whereas inflation-linked bonds also expose the buyer to the bond market, bringing unwanted risks that require extra capital to be hedged.

A large source of inflation risk for many insurance companies is the inflation indexation offered in the pension contracts. In the case of the Dutch insurer nationale-nederlanden (NN), most of their customers are guaranteed an indexation with an upper limit of either 3%, 4% or 5%. In the case of deflation, the insurer is often not allowed to cut the pension payments. Thus from the viewpoint of the pension insurer, this is equivalent to giving out a contract with an inflation floor of 0%. In the same way, we can view the upper indexation limit as an inflation cap with a strike price of either 3%, 4% or 5%. Thus proper valuation of inflation caps and floors is also essential for the valuation of the insurance companies own contracts.

Risk-neutral valuation of inflation caps and floors has two main difficulties. First, the products need to be discounted. This means that next to the inflation rate, we also include a dependence on the interest rate. These two underlying factors are often correlated, which we wish to account for. Second, we often observe an implied volatility smile that can cause significant mispricing for the far OTM and ITM products when not accounted for.

Currently the valuation at NN is performed using a scenario generation based on two independently modelled Hull-White processes for the interest rate and inflation. After independent calibration for both processes, the correlations are calculated based on the past data. The main problem with this approach is that in recent years the correlation between inflation rates and interest rates has changed dramatically as inflation levels peaked, while the interest rates remained close to zero as of 1/6/2022. This poses a real problem as it makes the valuation model less accurate in a time where it is the most needed.

To tackle this problem, we examine a modernized approach of the current valuation method, in which we use a correlated two-process Hull-White model instead of two independent single-factor Hull-White models. This allows us to calibrate the correlation parameter on the current market prices, instead of using historical correlations (Brigo & Mercurio, 2006). Furthermore this specification also enables us to obtain exact solutions for inflation caps and floors, instead of having to rely on simulations (Dodgson & Kainth, 2006).

Because the calibration process of the model parameters is often overlooked, we take a detailed look at the different calibration methods, in particular for capturing the volatility dynamics. In most current literature this is done using a time-dependent volatility parameter, and applying for example a local volatility model (Brigo et al., 2003; Gurrieri et al., 2009). However, this approach loses the analytical tractability, forcing us to calibrate using Monte Carlo simulation. Therefore, in this paper we take a different approach, by using a strike-dependent volatility parameter. We incorporate this strike-dependence in the calibration procedure, rather than the model itself. This allows us to maintain full analytical tractability of the original model by developing a closed-form solution. We find that this approach is able to fully capture the volatility dynamics while still remaining arbitrage-free for non-exotic inflation caps and floors. Moreover, by testing different weight functions we find that a normalized weight gives the best overall performance, while equal weights yield the best performance for the longer tenors, which is preferred as insurance companies often perform hedges over longer time frames of 10-30 years. Based on the out-of-sample performance, this paper finds that the model can be trusted for pricing new products with strike prices up to 6%. However, for tenors of 30 years and higher it is advised to not solely rely on the model for pricing.

Lastly, this paper finds that the correlated two-process HW model performs on average 0.92% better over a range of 108 different products compared to the uncorrelated two-process HW model as currently in use by NN.

2 Literature Review and Theory

For a clearer picture of the problem, we first look at the past developments for pricing of interest-rate derivatives. This way, we can place ourselves in the literature to see where useful contributions can be made. Next, we compare different contender models, in which we consider the advantages and shortcomings of each model, and relate them to our model requirements. We make our definite model choice and look more deeply at the characteristics of our chosen model. Lastly, we go over the basis of swaps and the pricing of caps and floors, which are extensively used in the rest of the report.

2.1 Developments in Pricing Inflation Derivatives

The earliest pricing methods for inflation derivatives mostly relied on using inflation predictions, such as using a VAR (Webb, 1995) or GARCH (Chou, 1988) model, which are then incorporated in a binomial pricing tree or Monte Carlo simulation. However, Barone & Castagna (1997) proposed developing a framework based on the original insight to to view inflation as if it were the price of a foreign currency, that is, the evolution of inflation is interpreted as the 'exchange rate' between the nominal and real economies. In this setting, the valuation of an inflation-indexed payoff becomes equivalent to that of a cross-currency interest rate derivative.

This allows the problem to be rewritten in a familiar framework, which Barone & Castagna (1997) applied for pricing TIPS. A year later Hughston (1998) developed the first closed-form solution for simple interest-rate derivatives based on the same foreign-currency analogy, rewriting the problem as an extension to the HJM framework of Heath et al. (1992). This allowed pricing without using direct inflation predictions and simulations, but instead relies on a calibration of the derivatives market prices to obtain a closed-form solution.

One of the greatest shortcomings of this early model however, is that the HJM framework tends to have infinite dimensions (Ritchken & Sankarasubramanian, 1995), which makes it very difficult and inefficient to accurately compute parameter estimates of the model (Filipović, 2000).

A framework with the same desirable properties, but without the problem of infinite dimensions is the Hull-White (HW) model. However, at the time of Hughston (1998) only closed-form HW solutions of interest rate derivatives had been found (Hull & White, 1990), but closed-form solutions to inflation derivatives did not yet exist.

The first closed-form solution to pricing inflation-linked derivatives was published by Mercurio (2005), based on the Jarrow-Yildirim (JY) model. This paper was breakthrough in the field and allowed for significantly more efficiënt calculation of interest-rate derivatives. However, the formulas in the JY framework were long and complicated, and quickly required numerical simulation when parameters in the model were made to be time-varying. Therefore, Dodgson & Kainth (2006), transformed the formulas of the JY to the HW framework by setting the volatility structure to $\sigma(t,T) = \sigma e^{-a(T-t)}$ without loss of generality. This results in a correlated two-process Hull-White processes for the short interest rate and inflation, which allows more flexibility in choosing the parameter structures, often without requiring numerical simulations.

Apart from the model itself, calibration of the model parameters to market prices is just as important, especially for the two-process case, where at least five parameters need to be estimated. Gurrieri et al. (2009) improve upon the calibration methods of the single process Hull-White model for inflation derivatives, by looking at different optimization criteria, and the effects of making certain parameters time-dependent. However, as of today the literature on the calibration of the two-process HW model is still very limited, especially on the different methods of capturing the volatility dynamics. Therefore, we wish to expand on the papers of Dodgson & Kainth (2006) and Gurrieri et al. (2009) by exploring a new way of capturing the volatility dynamics, in which we make use of a strike-dependent instead of a time-dependent volatility parameter.

2.2 Comparing Different Contender Models

Since we have a term structure available for both the interest rate, as well as for the inflation rates, we want to choose a model which can fully capture these initial term structures to maintain risk-neutrality. This requirement already excludes the Vasicek model. Furthermore the model should be able to handle negative rates, as in the past years, short interest rates have been predominantly negative. This excludes the CIR model. Lastly, we wish to calibrate the model on current market prices, instead of relying on past data, which VAR models. Therefore we are left with three main contenders: the HJM framework, the HW model and market models.

The HJM framemork is based on the assumption that the drifts of the no-arbitrage evolution of certain variables can be expressed as functions of their volatilities and autocorrelations. This approach brings the large advantage that no drift estimation is needed (Heath et al., 1992). The main disadvantage of the HJM framework is that it tends to have infinite dimensions (Ritchken & Sankarasubramanian, 1995). This makes it very difficult to compute, especially for closed-form solutions. There are various expansions to the HJM framework which aim to express it as a finite-state model. However, these models give rise to more complications, such as losing risk-neutrality, and drawing criticism from many academics such as Paul Wilmott, which calls the entire HJM framework "a big rug for mistakes to be swept under"[1].

Market models, such as the LIBOR model have become increasingly popular. These models are based on modelling a whole set of directly observable forward rates, instead of the short rate. This

[1]Source: https://www.newsweek.com/one-math-geeks-plan-reform-wall-street-80135

has the advantage of being able to price all kinds of exotic derivatives such as Bermudian swaptions or barrier options (Rebonato, 2012). However, due to the complex nature of the model, with many parameters and differential equations, accurate pricing is only possible through simulation. This requires a large number of simulated paths, each containing multiple stochastic differential equations with complicated drift terms, which is computationally expensive (Gurrieri et al., 2009).

2.3 Hull-White Model

The HW model is a no-arbitrage affine term structure model (ATSM), that is able to capture the entire initial forward curve. The general model, developed by Hull & White (1990) for modelling future interest rates is still widely used for modelling and pricing various interest-rate derivatives. The model is defined as follows:

$$dr(t) = \alpha(t)[\theta(t) - r(t)]dt + \sigma(t)dW_t, \tag{1}$$

where $\theta(t)$ is the mean process, which is defined such that it perfectly replicates the current term structure of interest rates. Parameter α is the speed of mean-reversion and σ is the volatility, which is able to capture the non-linearity in payouts of options.

In practice, α and σ are often kept constant over time, since this is the only form for which fully closed-form solutions are possible. However, for the valuation of options, modern approaches often allow the volatility to be time-varying, as is done in Gurrieri et al. (2009) and Singor et al. (2013) for a single HW process in combination with a Heston model for the volatility. Although these extensions remove the full analytical tractability, it is often worth the sacrifice in terms of accuracy, as we will see later in this paper.

2.3.1 Basic HW Model Characteristics

To obtain a solid theoretical basis, we first look at the characteristics of the basic short-rate process, which the further models in this paper are built upon. Since the standard expression of the HW model is in differential form, as seen in (1), we take the integral on both sides of short-rate process. Trying the solution $f(r(t)) = r(t)e^{at}$ and applying Itô's lemma, we can write the short-rate evolution as follows

$$r(t) = e^{-\alpha t}r(0) + \int_0^t e^{\alpha(s-t)}\theta(s)ds + \sigma e^{-\alpha t}\int_0^t e^{\alpha u}\, dW(u). \tag{2}$$

Note that we have taken α and σ to be constant, as this assumption does not change the structure of the outcome (Brigo & Mercurio, 2006).

As the Brownian motion is normally distributed, the short-rate process is also is normally distributed, with mean

$$E[r(t)|\mathcal{F}_s] = e^{-\alpha t}r(0) + \int_0^t e^{\alpha(s-t)}\theta(s)ds, \tag{3}$$

and variance

$$Var[r(t)|\mathcal{F}_s] = \frac{\sigma^2}{2\alpha}\left(1 - e^{-2\alpha(t-s)}\right), \tag{4}$$

which is actually the same variance as for the Vasicek model, since $\theta(t)$ is only present in the expression for the mean (3). Therefore we often call $\theta(t)$ the mean parameter of the HW process.

2.3.2 Combining Two HW Processes

What makes the HW model specifically appealing for pricing inflation derivatives is that it allows for simultaneous modelling of multiple correlated processes, while still retaining closed-form solutions. We do this by defining two separate models. One for the short interest rate $r(t)$ and one for the inflation $i(t)$, which we combine using a correlated Brownian motion. We define our two-process HW model under the risk-neutral measure $\mathbb{Q}$ as follows, based on the notation of Hull & White (1993) and Dodgson & Kainth (2006):

$$\begin{aligned} dr(t) &= \alpha[\theta(t) - r(t)]dt + \sigma dW_t^r \\ di(t) &= \alpha_I[\theta_I(t) - i(t)]dt + \sigma_I dW_t^I \\ dW^I &= \rho dW^r + \sqrt{1-\rho^2}dZ_t. \end{aligned} \tag{5}$$

Note that the volatility parameters σ and σ_I are fixed, which is common practice in the original papers. We will later expand on this by allowing these parameters to be strike-dependent. Moreover, we will derive the corresponding closed-form solutions for inflation caps and floors in 4.4.

2.4 Inflation Derivatives

2.4.1 Fundamental Pricing Formula

All derivatives used in this paper can be summarized under a single fundamental pricing equation that captures the risk-neutral payoff structure:

$$p_t = E^{\mathbb{Q}}\left[\int_t^T e^{-\int_t^s r_u du} X_s(\mathcal{F}_s)ds|\mathcal{F}_t\right]. \tag{6}$$

The specific payoff structure of the derivative is denoted by $X_s(\mathcal{F}_s)$, which we will define for swaps, caps and floors in the next section. All available data up to time point s is denoted by $\mathcal{F}_s$, which in our case are the short-rate and inflation. The maturity of the option is denoted by T.

2.4.2 Inflation Swaps

The inflation-linked derivatives market is dominated by zero-coupon inflation swaps (ZCIS) (Fleming & Sporn, 2013). The main advantage of these products, is that we can easily extract very useful information by expressing the ZCIS as a single cash flow at time T, which is done as follows: We denote the strike price by K, the nominal value by N and the current time by T_0, such that the time till expiry is $\tau = T - T_0$. The fixed payment is therefore $N[(1+K)^\tau - 1]$ and the floating amount received is $N\left[\frac{I(T)}{I(T_0)} - 1\right]$. Under the no-arbitrage condition the ZCIS must have zero value when $(1+K)^\tau = \frac{\hat{I}(t,T)}{I(T_0)}$, where $\hat{I}(t,T)$ is projected price index, which can be interpreted as the expected compounded inflation.

This projected price index can be obtained from the difference between the market value of an inflation-linked bond (ILB), denoted by $P_I(t,T)$ and a the market value of a ZCB with the same maturity:

$$\hat{I}(t,T) = \frac{P_I(t,T)}{P(t,T)}. \tag{7}$$

In this way we only keep the inflation-payout part of the ILB, and thus gives a market expectation of the inflation up till time point T. Combining the no-arbitrage condition with (7) we can express the ILB as a function of observed variables from the inflation swaps and ZCB's:

$$P_I(t,T) = I(T_0)(1+K)^T P(t,T). \tag{8}$$

We often evaluate the inflation and term structures at the current time $t = 0$, for which we set normalize the inflation to $I(0) = 1$ without loss of generality.

2.4.3 Inflation Caps and Floors

Contrary to swaps, which are linear products, caps and floors have a nonlinear payout structure. First, we define the general inflation option payout structure, which serves as a basis, as we can formally rewrite a caplet as a put option, and a floorlet as a call option, as we will formally show in Section 4.3

$$X_t^{call} = \left[\left(\frac{I(T)}{I(T_0)} - 1\right) - K\right]^+ \tag{9}$$

Note that this is not a typical call option, as the price index $I(T)$ is not a real traded asset. However, as shown in (7) and (8) we can obtain the price index from the inflation swaps, which are actively traded assets[2]. We can write the payoff of a caplet and floorlet in the convenient single equation

[2] Note that we are writing $I(T)$ instead of $\hat{I}(T)$ since we are talking about a realised payout. However, when estimating the payout conditional on $\mathcal{F}_t$ formal notation would be to use $\hat{I}(T)$.

by Mercurio (2005):

$$X_{t_i}^{Cpl/Fl} = N\psi_i\left[\omega\left(\frac{I(T_i)}{I(T_{i-1})} - 1 - K\right)\right]^+, \tag{10}$$

where the time in years between the reset date and payout date is represented by the fraction ψ_i. Since all caps and floors used in this paper are settled on a yearly basis, we have $\psi = 1$. The parameter $\omega = 1$ represents a caplet, and $\omega = -1$ a floorlet. Again, we will set the nominal value $N = 1$ for the remainder of the calculations. The present value of a YoY cap with a duration of n years is then:

$$Cp_t^{YoY} = \sum_{i=1}^{n} P(t,T_i)\left[\frac{I(T_i)}{I(T_{i-1})} - 1 - K\right]^+ \tag{11}$$

Often the simpler and more liquid ZC swap is traded, with a payoff only at maturity, giving a present value of

$$Cp_t^{ZC} = P(t,T)\left[\frac{I(T)}{I(T_0)} - (1+K)^n\right]^+ \tag{12}$$

Note that this is equivalent to substituting $X_t^{ZC} = \left[\frac{I(T)}{I(T_0)} - (1+K)^n\right]^+$ in the general pricing equation (6).

3 Data

3.1 Historical Inflation

European inflation-linked derivatives are usually based on either the Consumer Price Index (CPI) of the respective country, or the Harmonised Index of Consumer Prices (HICP). Both are designed to measure the change in the average level of prices for a large basket of household consumer goods. The HICP is a weighted average of the CPI's of all countries that have adopted the Euro. In order to avoid different data for different countries, we will focus on the HICP products.

The inflation rate is defined as the year-on-year change in the price of the bundle of goods as included in the HICP [3]. From now on we will denote the price of this bundle of goods as the Retail Price Index $I(t)$. The year-on-year inflation, denoted by $i(t)$ is then: $i(t) = I(t)/I(t-1) - 1$.

In Figure 1 the HICP of the past 20 years is plotted, which shows a strong deviation from the target inflation level of 2% in the past year. This has a large effect on derivatives with inflation as underlying. Therefore, we will take a closer look at some of these derivatives in the next sections.

[3] The CPI and HICP data are publicly published and updated monthly by Eurostat.

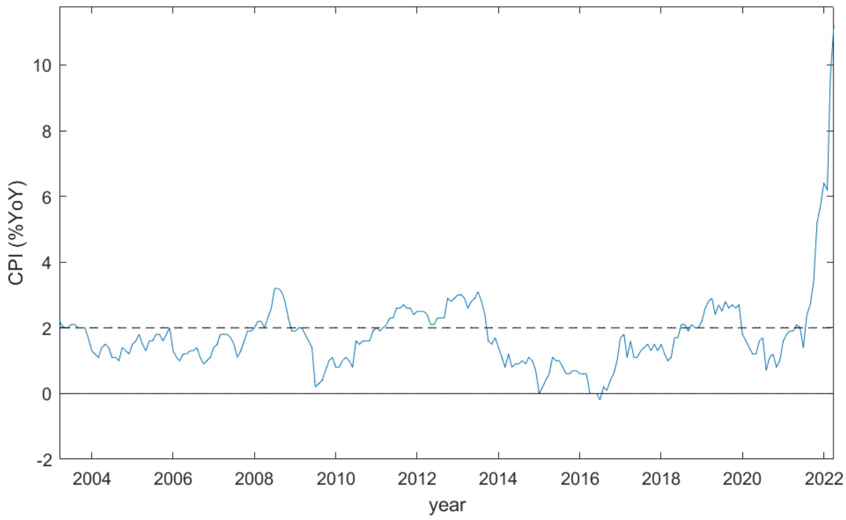

Figure 1: Year-on-year HICP for the period 01-2003 till 06-2022. The dotted line is the ECB inflation target of 2%.

3.2 Interest Rate Swaps

For the discounting of future cash-flows, we wish to obtain insight in the expectations of future interest rates. Therefore, we look at zero-coupon interest-rate swaps, for which the swap rate is the expectation of the interest rate over the corresponding tenor. Since we are concerned with both short-term and long-term forecasts, we wish to capture the entire forward curve. For this, we have chosen to use the Euribor swap rates, as they offer a swap curve for multiple maturities from 6 months up to 30 years. The 6M Euribor is the rate at which European banks lend out their money to other banks. From 1Y and onward the rates represent swap rates against the 6M rate. In Figure 2, the forward Euribor curve as of 1/6/2022 is plotted.

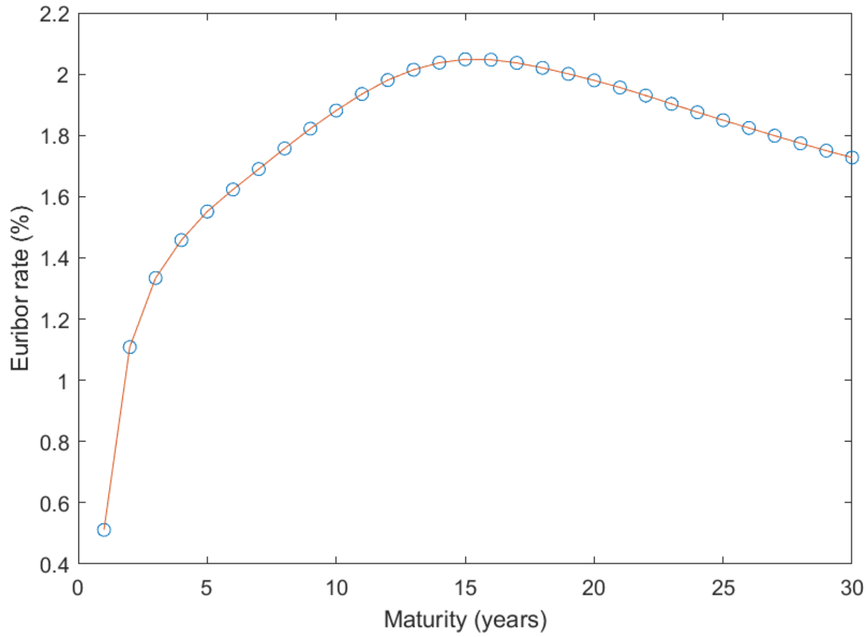

Figure 2: Euribor zero-coupon swap curve for maturities up to 30 years, obtained at 1/6/2022

In order to maintain risk-neutrality, we make the assumption that the swap contracts between two banks are collateralized loans, such that default risks are not considered in the price. This assumption can be done since the counterparty institutions are often very large banks which have a small probability of defaulting, and thus the difference concerns only a few basis points Canabarro & Duffie (2003). Therefore, we can now take the Euribor rates as risk-free rates, which we will denote by $r_{i,t}$ for whicih $i = 1$ is the 6M rate and $i = 10$ is the 30Y rate respectively.

3.3 Inflation Swaps

As an input for pricing inflation-linked derivatives using the HW model, we include the forward inflation curve, for which we use zero coupon HICP ex Tobacco swaps (HICPxT) with a 3M lag from Tullett, available on the Bloomberg Terminal.

Firstly, these swaps use the same underlying as the other inflation derivatives we will be using later, including the same 3M lag. Secondly, these swaps are very liquid, and are traded for many different tenors, of up to 50 years. In order to directly use the quoted inflation levels, we assume that the swap spread is zero, meaning that we ignore counterparty risk, using the same arguments as for the interest rate swaps. As of 1/6/2022 the forward inflation curve looks as follows:

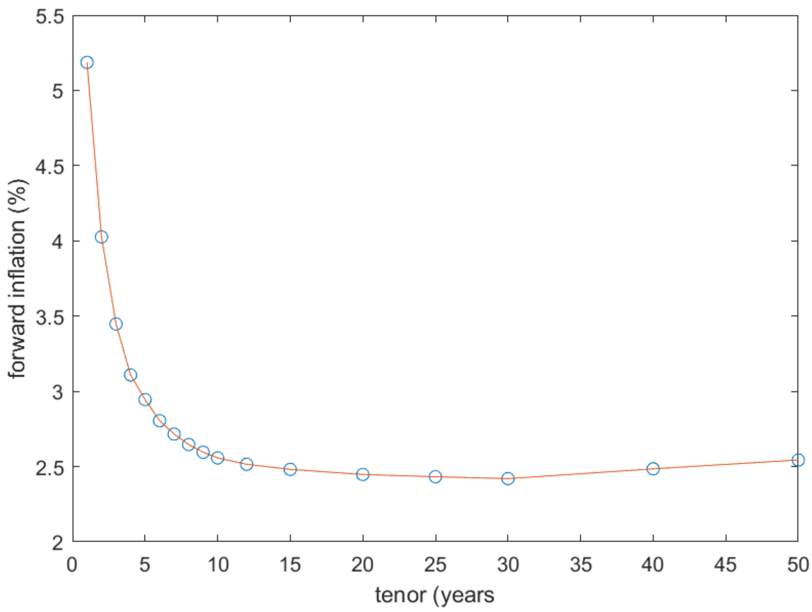

Figure 3: HICPxT swap forward curve with 3M lag as of 1/6/2022. The actual market forward inflation rates are represented by the dots. The smoothed curve is obtained by cubic spline interpolation.

The inflation term structure is inverted compared to the interest-rate term structure. This is because of the current inflationary spike as shown in Figure 1, which causes the one-year inflation expectation to be exceptionally high. Note the steep decline in rates for the longer tenors, implying that inflation is expected to return to levels of around 2.5% in the future.

3.4 Interest Rate Derivatives

In order to calibrate the model on inflation caps and floors, we first need to calibrate on interest rate derivatives with a nonlinear payoff, also called optionality. The available products containing this optionality are caps, floors and swaptions, which can all be used for calibration interchangeably, provided they have the same strikes and tenors[4]. In Table 1 a selection of the YoY Euribor cap and floor data is displayed, and the full data set can be found in Appendix A, from which we can see that especially the caps are very actively traded for many different strike prices. The floors however, are less actively traded, except for the 0% floor (Brigo & Mercurio, 2006).

[4]Due to the put-call parity, which says that for the same strike and maturity, a cap can be written as a floor plus a swap. Furthermore the implied volatility should be the same for all products with a nonlinear payoff structure (Brigo & Mercurio, 2006).

Table 1: Selection of the market prices of YoY 3M Euribor caps and floors from Tullett as of 1/6/2022. The percentages displayed horizontally are the strikes with respect to the Euribor. The years displayed vertically are the tenors. The market prices are given in basis points of the nominal. The full dataset can be found in Appendix A.

Cap	1%	2%	3%	4%	5%
1Y	8.770	0.650	0.100	0.020	0.010
3Y	199.0	91.86	50.41	32.03	22.33
5Y	430.3	232.5	143.3	98.83	73.47
7Y	699.9	408.6	265.0	188.7	143.3
10Y	1147	708.3	467.2	329.7	245.5
15Y	1865	1199	800.2	559.4	408.2
20Y	2412	1584	1076	762.9	562.8

Floor	0%	0.5%	1%	1.5%	2%
1Y	0.800	0.850	7.600	25.00	47.70
3Y	11.30	31.40	48.30	103.4	182.2
5Y	33.80	73.00	114.7	210.2	342.3
7Y	58.30	119.0	185.4	319.8	501.5
10Y	93.50	178.1	278.9	460.5	702.9
15Y	158.0	274.1	434.6	686.8	1018
20Y	270.3	429.9	666.5	1001	1429

The illiquid nature of the other floor strike prices can be seen by simple arbitrage opportunities in the data. For example the 1Y floor with strike 0.5% is 8.5 bps, whereas the quoted market price of a 1Y floor with stike 1% is 7.6 bps, which is impossible under the no-arbitrage assumption. Therefore, for calibration of the interest rate model we will only use the caps data, also because these already cover almost all strike prices.

3.5 Inflation Caps and Floors

In 2012, the OTC inflation swap market was estimated to be worth around 109 billion GBP in the UK alone, with trading volumes in excess of 100 million a day (Fleming & Sporn, 2013). Through Tullett, many OTC trades of inflation caps and floors are published, which are updated on a live basis.

All inflation products used in this paper are based on the HICP ex tobacco (HICPxT) with a 3M lag, which is the same as is used for the inflation swaps. Both ZC and YoY caps/floors are quoted, but since the YoY products are much more liquid, we have chosen to only use the YoY caps and floors for calibration. The cap maturities range from 0.5% to 5% and the floor maturities from 0% till 2%. A snapshot of the YoY HICP cap and floor data is displayed in Table 2.

Table 2: YoY HICP caps and floors market prices from Tullett as of 1/6/2022. The percentages displayed horizontally are the strikes with respect to the HICP index. The years displayed vertically are the tenors. The market prices are given in basis points of the nominal.

Cap	1%	2%	3%	4%	5%	6%
1Y	417.1	317.9	219.7	124.9	45.79	10.33
2Y	603.8	413.2	251.2	134.9	49.47	10.90
5Y	997.5	593.9	333.5	182.9	81.49	26.54
7Y	1248	720.0	402.7	229.4	116.3	50.93
10Y	1625	921.2	517.8	308.0	176.1	98.29
12Y	1870	1047	583.8	350.7	207.3	121.0
15Y	2220	1222	670.7	405.2	246.7	146.7
20Y	2765	1500	812.1	496.1	313.7	193.3
30Y	3772	2038	1095	678.9	448.8	298.5

Floor	-2%	-1%	0%	1%	2%	3%
1Y	0.08	0.14	0.24	0.47	0.99	2.39
2Y	0.08	0.14	0.39	2.38	10.10	46.03
5Y	0.89	4.37	17.93	43.72	123.2	345.6
7Y	3.69	14.01	44.38	88.42	225.1	572.5
10Y	13.08	39.30	93.69	168.9	389.4	909.8
12Y	16.09	48.84	118.2	208.8	473.1	1096
15Y	18.93	58.96	148.0	259.1	579.4	1345
20Y	25.49	79.70	203.5	345.3	751.3	1732
30Y	46.60	135.8	324.2	525.2	1085	2435

Under the no-arbitrage condition, YoY cap and floor prices should be strictly increasing with maturity, as for example the 2Y cap is the sum of the 1Y cap and a 1Y cap bought after a year. Moreover, floor prices should be increasing with the strike price, and caps should be decreasing. We see that all products in Table 2 adhere to these conditions, which is a sign that the data is reasonably liquid. Most notably, the floors are much more liquid compared to those of Table 1, which is no surprise since inflation-indexed floors, especially those with low strike prices are the most actively traded inflation products (Brigo & Mercurio, 2006).

We do have to be careful with the shorter maturities, as more variation in the pricing could be present due to a lower reported price accuracy, and possible market inefficiencies such as transaction costs and service costs. For this reason we decide to exclude the 1Y tenors, as they often negatively impact the estimation accuracy of the model parameters. Moreover, caps and floors with a 1Y tenor are often of small interest for banks and insurers.

4 Methodology

4.1 Replication of the Term Structure

We start by defining the mean processes $\theta(t)$ and $\theta_I(t)$ of the HW models such that the initial term structure can be replicated. Because the initial yield curve is only given for a discrete set of tenors, and since we wish to work in the continuous-time domain in order to obtain closed-form solutions, we approximate a continuous yield curve[5] by performing daily interpolation using cubic splines over the full 30 years, using swap-curve yields as knots. This results in 10,950 data points for both the interest-rate and inflation curves.

Since we are in essence still using discrete programming, we make use of the forward difference. For more convenient notation we define the discrete time index by $t' = t/\Delta t$ where $\Delta t = 1/365$ for daily interpolation. We now compute the instantaneous forward curve as follows:

$$f(0,t) = \frac{\partial}{\partial t}(ty(t')) \approx \frac{(t+\Delta t)y(t'+1) - t\Delta t y(t')}{\Delta t} = (t'+1)y(t'+1) - t'y(t'). \tag{13}$$

In order to perfectly replicate the initial observed market term structure, which can be done for any choice of parameters α_r and σ_r, the mean process for the short rate θ_r is set equal to the following:

$$\theta(t) = \frac{\partial f_r(0,t)}{\partial t} + \alpha_r f_r(0,t) + \frac{\sigma^2}{2\alpha_r}(1 - e^{-2\alpha_r t}). \tag{14}$$

Because we are using cubic spline interpolation, which is twice differentiable, (14) is always defined. The last part of (14) is usually fairly small, meaning that the drift process follows the slope of the forward curve, and mean-reverts at speed α_r. Omitting the last term and substituting back in the short-rate process shows that the drift at time t is $r = \frac{\partial f(0,t)}{\partial t} + \alpha(f(0,t) - r)$, which shows that on average r follows the slope of the initial instantaneous forward curve, with a mean reversion back to α.

For the inflation process, we replicate the initial term structure in a similar fashion, albeit a slightly more complicated expression, since the mean process is dependent on both the short-rate and the inflation rate. In Dodgson & Kainth (2006) we find that we have to set θ_I equal to:

$$\theta_I(t) = f_I(0,t) + \frac{1}{\alpha}\frac{\partial}{\partial t} f_I(0,t) - \frac{\sigma_I^2}{2\alpha_I^2}(1 - 2e^{-2\alpha_I t}) + \frac{\rho\sigma\sigma_I}{\alpha_I}\left[\frac{1}{\alpha}(1 - e^{-\alpha t}) + \frac{1}{\alpha_I}(1 - e^{-\alpha_I t})e^{-\alpha t}\right] \tag{15}$$

Where the instantaneous forward is calculated as $f_I(0,t) = \frac{\partial \ln(\hat{I}(T))}{\partial T}$. We found earlier that at $t = 0$ we can directly find the projected price index as a function of the zero-coupon swap rates: $\hat{I}(T) = (1+K)^T$.

4.2 Risk-Neutral Valuation of a ZCB

Using the HW model it is relatively straight forward, although tedious, to model zero coupon bond prices using the affine term structure model (ATSM).

[5]Note that we make the assumption that in between the points of the initial discrete yield curve no unexpected humps or other dynamics occur.

We start off by writing out the ZCB pricing formula $P(t,T) = E^{\mathbb{Q}}\Big[e^{-\int_t^T r(u)du}|\mathcal{F}_t\Big]$ using the short-rate expression as defined in (2). Due to the Gaussian distribution of $r(T)$ conditional on $\mathcal{F}_t, t < T$, we can show after a rather tedious derivation that $\int_t^T r(u)du$ is also normally distributed (Blanchard, 2014):

$$\int_t^T r(u)du|\mathcal{F}_t \sim N\bigg(B(t,T)[r(t) - \phi(t) \ + \ln\frac{P(0,t)}{P(0,T)} + \frac{1}{2}[V(0,T) - V(0,t)], V(t,T)\bigg).$$

where

$$B(t,T) = \int_t^T e^{-\alpha s}ds = \frac{1}{\alpha}\Big[1 - e^{-\alpha(T-t)}\Big], \tag{16}$$

and

$$V(t,T) = \frac{P(0,T)}{P(0,t)}\exp\bigg[T - t + \frac{2}{\alpha}e^{-\alpha(T-t)} - \frac{1}{2\alpha}e^{-2\alpha(T-t)} - \frac{3}{2\alpha}\bigg], \tag{17}$$

so that we obtain the ATSM by using the log-normality property $E[e^Z] = e^{\mu + \frac{1}{2}\sigma^2}$

$$P(t,T) = A(t,T)e^{-B(t,T)r(t)}, \tag{18}$$

where

$$A(t,T) = \frac{P(0,T)}{P(0,t)}\exp\Big\{B(t,T)f(0,t) - \frac{\sigma^2}{4a}(1 - e^{-2\alpha t}B(t,T)^2)\Big\}.$$

We can now reconstruct the initial yield curve by translating back the initial ZCB prices P(0,T) of (18) by $y(0,T) = -\frac{1}{T}\log P(0,T)$. A detailed derivation of the above formulas can be found in Chapters 3.3.2 and 4.2.2 of Brigo & Mercurio (2006).

4.3 Closed-Form Solution of IR Caps and Floors

In order to calibrate the entire model to inflation products, we first have to calibrate the single-factor HW model to interest rate caps and floors. In this section we will do the derivation for caps, as the process for floors is done in similar fashion.

Because there is no direct closed-form solution for IR caps or floors, we first rewrite the cap as a sum of zero bond put options (ZBP), for which a closed-form solution for the HW model is available. The rewriting of the IR cap in caplets and then in ZBP's is done as follows:
Based on (6) the present risk-neutral value of a ZBP option can be expressed as

$$\mathbf{ZBP}(t,T,S,X) = E\Big[e^{-\int_t^T r^s ds}(X - P(T,S))^+|\mathcal{F}_t\Big].$$

We consider a cap based on the LIBOR rate L with nominal N and strike X, evaluated between payout times t_i and t_{i-1}, with δ being the time between payments, which we can rewrite as follows:

$$\begin{aligned}\mathbf{Cpl}(t,t_{i-1},t_i,N,X) &= E\Big[e^{-\int_t^{t_i} r_s ds} N\delta[L(t_{i-1},t_i)-X]^+|\mathcal{F}_t\Big] \qquad (19)\\ &= NE\Big[e^{-\int_t^{t_{i-1}} r_s ds} P(t_{i-1},t_i)\delta[L(t_{i-1},t_i)-X]^+|\mathcal{F}_t\Big]\\ &= NE\Big[e^{-\int_t^{t_{i-1}} r_s ds} P(t_{i-1},t_i)\Big[\frac{1}{P(t_{i-1},t_i)}-1-X\delta\Big]^+|\mathcal{F}_t\Big]\\ &= NE\Big[e^{-\int_t^{t_{i-1}} r_s ds}[1-(1+X\delta)P(t_{i-1},t_i)]^+|\mathcal{F}_t\Big]\\ &= N(1+X\delta)E\Big[e^{-\int_t^{t_{i-1}} r_s ds}\Big[\frac{1}{1+X\delta}-P(t_{i-1},t_i)\Big]^+|\mathcal{F}_t\Big]\\ &= N'E\Big[e^{-\int_t^{t_{i-1}} r_s ds}[X'-P(t_{i-1},t_i)]^+|\mathcal{F}_t\Big]\\ &= N'\mathbf{ZBP}(t,t_{i-1},t_i,X')\end{aligned}$$

where

$$X' = \frac{1}{1+X\delta}$$
$$N' = N(1+X\delta).$$

Analogously

$$\mathbf{Fll}(t,t_{i-1},t_i,N,X) = N'\mathbf{ZBC}(t,t_{i-1},t_i,X'). \qquad (20)$$

Using this, we can write the cap and floor prices as a sum of zero coupon options:

$$\mathbf{Cp}(t,\mathcal{T},N,X) = \sum_{i=1}^{n}\mathbf{ZBP}(t,t_{i-1},t_i,X') \qquad (21)$$
$$\mathbf{Fl}(t,\mathcal{T},N,X) = \sum_{i=1}^{n}\mathbf{ZBC}(t,t_{i-1},t_i,X')$$

where $\mathcal{T} = \{t_0,t_1,...,t_n\}$ is the set of times between the current date and the settlement date. Important for IR caps and floors to keep in mind is that the payout occursr after the settlement date. Therefore when programming the above formulas, keep in mind to add one year to every element in $\mathcal{T}$. We can now use this to write a closed form expression for HW cap and floor

$$\begin{aligned}\mathbf{Cp}(t,\mathcal{T},N,X) =& N\sum_{i=1}^{n}(1+X\delta)\mathbf{ZBP}\Big(t,t_{i-1},t_i,\frac{1}{1+X\delta}\Big) \qquad (22)\\ =& N\sum_{i=1}^{n}\big[P(t,t_{i-1})\Phi(-h_i+\sigma_p^i)-(1+X\delta)P(t,t_i)\Phi(-h_i)\big]\end{aligned}$$

where

$$\sigma_p^i = \sigma\sqrt{\frac{1-e^{-2a(t_{i-1}-t)}}{2a}}B(t_{i-1},t_i), \qquad (23)$$

$$h_i = \frac{1}{\sigma_p^i}\ln\frac{P(t,t_i)(1+X\delta)}{P(t,t_{i-1})}+\frac{\sigma_p^i}{2} \qquad (24)$$

A detailed derivation of (23) and (24) can be found in Chapter 3 of Brigo & Mercurio (2006).

4.4 Closed-Form Solution of Inflation Caps and Floors

Having obtained the parameters α and σ for the IR part of the HW model, we can now use these to price the inflation-linked caps and floors, using the full model as defined in (5).

First, we need to obtain the market prices of ILB's. However, these cannot be obtained reliably through Tullett. Luckily, in Section 2.4.2 we have seen that we can work around this by rewriting the ILB prices as a function of our liquid inflation swaps together with the ZCB prices, by applying the no-arbitrage condition to (7):

$$P_I(t,T) = (1+K)^T P(t,T), \tag{25}$$

where K is the projected inflation rate as obtained from the HICP swap rates. Again we have set the starting index at $I(T_0) = I(0) = 1$ without loss of generality.

Building further on our general model as derived in Section 4.2 we now include a second correlated HW model for inflation $i(t)$ to our calculations, by which we can express the ZC ILB price as

$$P_I(t,T) = I(t)E\left[e^{\int_t^T [i(t')-r(t')]dt'}\middle| i(t), r(t)\right] \tag{26}$$

where $I(T) = e^{\int_0^T i(t')dt'}$.

The assumption for this model is that the projected price index $\hat{I}(t,T)$ is modelled as a log-normal process, which naturally includes a drift term. However, since we are interested in risk-neutral derivatives pricing, we wish to transform this to a martingale process by removing this drift. This is done by choosing the ZCB $P(t,T)$ as numeraire, such that $\hat{I}(t,T) = P_I(t,T)/P(t,T)$ is a martingale under $\mathbb{Q}$. This is also known as the T-forward measure, for which we apply a Radon-Nikodym derivative $\frac{d\mathbb{Q}^{\mathbb{T}}}{d\mathbb{Q}} = \frac{e^{\int_0^T i(t')dt'}}{P(0,T)}$ such that the projected price index now follows the following process

$$\frac{d\hat{I}(0,T)}{\hat{I}(0,T)} = \sigma_I dW_I \tag{27}$$

since we start at $T_0 = 0$. The inflation part of the model develops as follows:

$$i(t) = i(s)e^{-\alpha_I(t-s)} + \alpha_I \int_s^t dt'\theta_I(t')e^{-\alpha_I(t-t')} + \sigma_I \int_s^t dW_I(t')e^{-\alpha_I(t-t')} \tag{28}$$

We can now work out the integral in (26) which we again show only for the inflation part, as the

IR part is similar.

$$\begin{aligned}\int_{T_1}^{T_2} i(t')dt' &= i(t)e^{-\alpha_I(T_1-t)}B(\alpha_I, T_2-T_1)\\ &+\alpha_I\int_t^{T_2} dt'\theta_I(t')^{-\alpha_I(T_1-t')}B(\alpha_I, T2-T1)+\alpha_I\int_{T_1}^{T_2} dt'\theta_I(t')B(\alpha_I, T2-t')\\ &+\sigma_I\int_t^{T_2} dW_I(t')e^{-\alpha_I(T_1-t')}B(\alpha_I, T2-T1)+\sigma_I\int_{T_1}^{T_2} dW_I(t')B(\alpha_I, T2-t'),\end{aligned} \tag{29}$$

where $B(\alpha,\tau) = (1-e^{-\alpha\tau})/\alpha$, which is an expression that pops up frequently in closed-form expressions of the HW model.

We will work out the derivation of the closed-form solution for the cap, as the floor follows in a similar matter. From (10) and (12) we see that the YoY cap is a sum of caplets which can be written as

$$\mathbf{Cpl} = E\left[e^{-\int_t^{T_2} r(t')dt'}\left(\frac{I(T_2)}{I(T_1)}-1-K\right)^+\right], \tag{30}$$

note that we have taken a nominal of $N=1$ as Tullet reports the market prices in basis points of the nominal. However, to convert these basis points to a fraction, we divide all basis points by 100^2. Combining (25) and (26), and substituting these in (30) gives

$$\mathbf{Cpl} = E\left[\left(e^{-\int_t^{T_2} r(t')dt'+\int_{T_1}^{T_2} i(t')dt'} - e^{-\int_t^{T_2} r(t')dt'}(1+K)\right)^+\right], \tag{31}$$

since $r(t)$ and $i(t)$ are normally distributed as can be seen from their respective expressions as found in (2) and (28), we can use the following lemma to rewrite the integrals as ZCB's and ILB's

$$\begin{aligned}E\left[e^{-\int_t^{T_2} r(t')dt'+\int_{T_1}^{T_2} i(t')dt'}\right] &= \frac{E\left[e^{-\int_t^{T_1} r(t')dt'}\right]E\left[e^{\int_t^{T_2}[i(t')-r(t')]dt'}\right]}{E\left[e^{\int_t^{T_1}[i(t')-r(t')]dt'}\right]}e^{Cov\left[\int_t^{T1} i,\int_{T_1}^{T_2}(r-i)\right]}\\ &= P(t,T_1)\frac{P_I(t,T_2)}{P_I(t,T_1)}e^{Cov\left[\int_t^{T1} i,\int_{T_1}^{T_2}(r-i)\right]},\end{aligned} \tag{32}$$

where evaluating the integrals in the covariance we find after some tedious algebra that

$$\begin{aligned}C(t,T_1,T_2) = Cov\left[\int_t^{T1} i,\int_{T_1}^{T_2}(r-i)\right] &= -\frac{\sigma_I^2}{2}B^2(\alpha_I, T_1-t)B(\alpha_I, T_2-T_1)\\ &-\frac{\rho\sigma_I\sigma}{\alpha_I}B(\alpha, T_2-T_1)[B(\alpha_I+\alpha, T_1-t)-B(\alpha, T_1-t)].\end{aligned} \tag{33}$$

We call this term the convexity correction, which arises after taking the expected value of a nonlinear product. The correction follows from Jensen's inequality $E[f(X)] \geq f(E[X])$, and compensates for the correlation between the two integrals, the use of which is essential for nonlinear products such as caps and floors. Applying the lemma of (32) to our caplet expression of (31) together with the normal properties of (30) we obtain

$$\mathbf{Cpl} = P(t,T_2)\Big[\frac{\hat{I}(t,T_2)e^{C(t,T_1,T_2)}}{\hat{I}(t,T_1)}\Phi(d_+) - (1+K)\Phi(d_-)\Big], \tag{34}$$

with

$$d_\pm = \frac{1}{\sqrt{V_I(t,T_1,T_2)}}\Big[\ln\Big(\frac{\hat{I}(t,T_2)e^{C(t,T_1,T_2)}}{\hat{I}(t,T_1)(1+K)}\Big) \pm \frac{V_I(t,T_1,T_2)}{2}\Big] \tag{35}$$

and variance of the inflation process as

$$\begin{aligned} V_I(t,T_1,T_2) = \mathrm{V}\Big[\int_{T_1}^{T_2} i(t')dt'|i(t)\Big] &= \frac{\sigma_I^2}{\alpha_I^2}\Big[T_2 - T_1 - B(\alpha_I, T_2-T_1) - \frac{\alpha_I}{2}B(\alpha_I,T_2-T_1)^2\Big] \\ &\quad + \sigma_I^2 B(2\alpha_I, T_1-t)B(\alpha_I,T_2-T_1)^2. \end{aligned} \tag{36}$$

Obtaining these exact expressions requires a lot of tedious but straightforward algebra. Even though the expressions are a bit more convoluted, we can clearly see the familiar Black-Scholes form in (34) and (35). Therefore, for a more intuitive expression we rewrite the above equations in the Black form of a call $c^{Black} = P(T)[S\Phi(d_+) - K\Phi(d_-)]$:

$$\mathbf{Cpl}(t,T_1,T_2,K) = c^{Black}\Big(\frac{\hat{I}(t,T_2)e^{C(t,T_1,T_2)}}{\hat{I}(t,T_1)}, (1+K), \sqrt{V_I(t,T_1,T_2)}\Big), \tag{37}$$

and in the same way

$$\mathbf{Fll}(t,T_1,T_2,K) = p^{Black}\Big(\frac{\hat{I}(t,T_2)e^{C(t,T_1,T_2)}}{\hat{I}(t,T_1)}, (1+K), \sqrt{V_I(t,T_1,T_2)}\Big). \tag{38}$$

As a last step we sum the caplets and floorlets to obtain the analytical model prices of the caps and floors based on the correlated two-processs HW model.

4.5 Parameter Calibration

4.5.1 Calibration of the IR Process

Before directly calibrating on inflation derivatives, we first calibrate the IR part of the two-process HW model on interest rate derivatives to capture the discounting dynamics given by parameters α and σ. However, as previously discussed, the IR floors are fairly illiquid and contain arbitrage opportunities, causing the calibration to become inaccurate. Since the caps already cover every strike except 0%, we choose to calibrate only on IR caps. We could consider cleaning the IR floors data by removing the outliers and interpolating for the missing values. However, since there is so little floors data already, this becomes a dangerous practice.

We have N_T different maturities, for each of which we have N_K strikes. The market prices of the IR caps are denoted by $C_{t,k} = C^{market}(\tau_t, K_k)$. For all of these caps we compute the corresponding model price denoted by $C_{t,k}(\alpha,\sigma) = C^{model}(\tau_t, K_k, \alpha, \sigma)$. We find the optimal IR model parameters by minimizing the squared error between the market prices and the computed model prices.

$$\hat{\alpha}, \hat{\sigma} = \arg\min_{\alpha,\sigma} \left\{ \sum_{t=1}^{N_T} \sum_{k=1}^{N_K} w_{t,k} (C_{t,k} - C_{t,k}(\alpha,\sigma))^2 \right\} \tag{39}$$

Here $w_{t,k}$ is the weight parameter. It is common to set which we set $w_{t,k} = 1/C_{t,k}^2$, such that we optimize over the normalized weight $\frac{C_{t,k} - C_{t,k}^{\Lambda}}{C_{t,k}}$ for each product. Since we are mostly interested in the pricing of longer tenors, we could alter the weight function to for example $w_{t,k} = 1/C_{t,k}^{3/2}$ or $w_{t,k} = 1/C_{t,k}$ giving more weight to the larger tenors as the price of the products is strictly increasing with maturity, which we also saw in Table 1. Another option is to include the tenor in the weight function itself: $w_{t,k} = T_t/C_{t,k}^2$, or a combination of both approaches: $w_{t,k} = T_t/C_{t,k}$, which we will later investigate in Section 5.4. For now, we stick to the normalized weights $w_{t,k} = 1/C_{t,k}^2$.

4.5.2 Calibration of the Inflation Parameters for Constant Volatility

Using the parameters $\hat{\alpha}$ and $\hat{\sigma}$ as calibrated from the IR process, we can now perform calibration on the inflation products, for which we optimize over $\Lambda = \{\alpha, \alpha_I, \sigma, \sigma_I, \rho\}$, keeping the previously found α and σ fixed. In this case we can calibrate over both caps and floors, as we have liquid market quotes for both products. Because both our caps and floors have the same set of maturities N_T, the optimization looks as follows

$$\hat{\Lambda} = \arg\min_{\Lambda} \left\{ \sum_{t=1}^{N_T} \left[\sum_{k=1}^{N_K} w_{t,k} \left(C_{t,k} - C_{t,k}^{\Lambda}\right)^2 + \sum_{j=1}^{N_K} w_{t,j} \left(F_{t,j} - F_{t,j}^{\Lambda}\right)^2 \right] \right\}. \tag{40}$$

The main concern with this objective function is the existence of local minima. Therefore, it is advised to start the numerical optimization with a larger step-size and to use different sets of starting parameters. In this paper a constrained interior-point algorithm is used with constraints: $-1 \leq \alpha_I \leq 1, \quad -1 \leq \rho \leq 1$ and $\sigma_I \geq 0$.

4.5.3 Capturing the Volatility Dynamics

As mentioned in many papers such as Sircar et al. (1999) and Backus et al. (2004), the biggest modelling inaccuracy often comes from the assumption of constant volatility, since in reality we often see a clear volatility smile when pricing options. In order to capture this smile, we allow our volatility parameter to be strike-dependent during the optimization process, while retaining our closed-form solutions obtained in Section 4.4. This is done by treating the problem as N_K

individual optimizations of $\sigma_I(K)$ within the general optimization for α_I and ρ, such that we now obtain a parameter set $\Lambda = \{\alpha, \alpha_I, \sigma, \sigma_I(K_1), ..., \sigma_I(K_{N_K}), \rho\}$ consisting of $N_K + 4$ parameters. The most intuitive way is to view this as a $(N_K + 1)$-process HW model of different volatilities, where each inflation process has the same mean-reversion α_I and is correlated to the single IR process by the same correlation ρ, as expressed in 41.

$$\begin{aligned} dr(t) &= \alpha[\theta(t) - r(t)]dt + \sigma dW_t^r \\ di_1(t) &= \alpha_I[\theta_I(t) - i_1(t)]dt + \sigma_I(K_1)dW_t^I \\ &\vdots \\ di_{N_K}(t) &= \alpha_I[\theta_I(t) - i_{N_K}(t)]dt + \sigma_I(K_{N_K})dW_t^I \\ dW_t^I &= \rho dW_t^r + \sqrt{1-\rho^2}dZ_t. \end{aligned} \tag{41}$$

For plain vanilla caps and floors the model is still arbitrage-free as we take a single parameter $\sigma_I(K)$ corresponding to the strike of the cap/floor, reducing the system back to a two-process HW. However, when pricing exotic caps and floors which for example depend on two strike prices, we lose the no-arbitrage conditions as we now need to use a combination of the inflation processes corresponding to the different strike prices, allowing us to create arbitrage opportunities.

The objective function based on (41) looks as follows

$$\hat{\Lambda} = \arg\min_{\Lambda} \left\{ \sum_{t=1}^{N_T} \left[\sum_{k=1}^{N_K} w_{t,k} \left[C_{t,k} - C_t^{\Lambda}\left(\sigma_I(K_k)\right)\right]^2 + \sum_{j=1}^{N_K} w_{t,j} \left[F_{t,j} - F_t^{\Lambda}\left(\sigma_I(K_j)\right)\right]^2 \right] \right\}. \tag{42}$$

In the case of our combined inflation cap and floor data we have $N_K = 9$ so that we have 13 total parameters, used for modelling 108 caps and floors. We could go even further and also relax the volatility of the IR model to $\sigma = \sigma(K)$ resulting in least $2N_K + 2$ parameters, but we have to be careful for over-parameterization.

4.5.4 Further Extensions for the Calibration Process

Sometimes during the optimization the parameters can vary wildly, as is often the case for the correlation parameter ρ. We can try to avoid this by imposing regularization, where we can set a baseline set of parameters θ_0. We could for example obtain a baseline for ρ by looking at the past correlation between inflation and interest rates. The optimization problem then becomes

$$\hat{\Lambda} = \arg\min_{\Lambda} \left\{ \sum_{t=1}^{N_T} \left[\sum_{k=1}^{N_K} w_{t,k} (C_{t,k} - C_{t,k}^{\Lambda})^2 + \sum_{j=1}^{N_K} w_{t,j} (F_{t,j} - F_{t,j}^{\Lambda})^2 \right] \right\} + \gamma(\theta - \theta_0)'(\theta - \theta_0), \tag{43}$$

where γ is the hyperparameter which determines how heavy deviations from θ_0 are penalized. It is often set by hand, but can also be taken in the optimization by for example a grid search. Note that we do not have to perform regularization over all parameters that are in $\hat{\Lambda}$, such that we can leave out $\sigma(K)$ as we will see in the results section.

For comparison of the goodness-of-fit from the different calibration strategies, we introduce the weighted RMSE (WRMSE) which takes the weighted average error of all modelled products $\boldsymbol{X}^{\Lambda}$ to their market prices $\boldsymbol{X}$ which can be caps, floors or both.

$$\text{WRMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(\frac{X_i - X_i^{\Lambda}}{X_I}\right)^2}, \tag{44}$$

where N is the total amount of included products. The advantage of using this over the standard RMSE is that now every product has the same impact, whereas otherwise the errors of the longer tenors would dominate the RMSE, giving a distorted picture.

5 Results

5.1 Replicating the Yield Curve by ZCB Pricing

In order to test the basis of our model, we first look at the ability of the HW model to perfectly replicate the initial yield curve. This is done using the ATSM expression of Chapter 4.2, calibrating on the ZCB prices. Since in theory we should obtain an exact replication of the initial pricing curve, this not only tests the correctness of our ATSM, but also whether our numerical optimization method is implemented correctly. The ZCB model fits after calibration, together with the initial pricing curve are plotted in Figure 4.

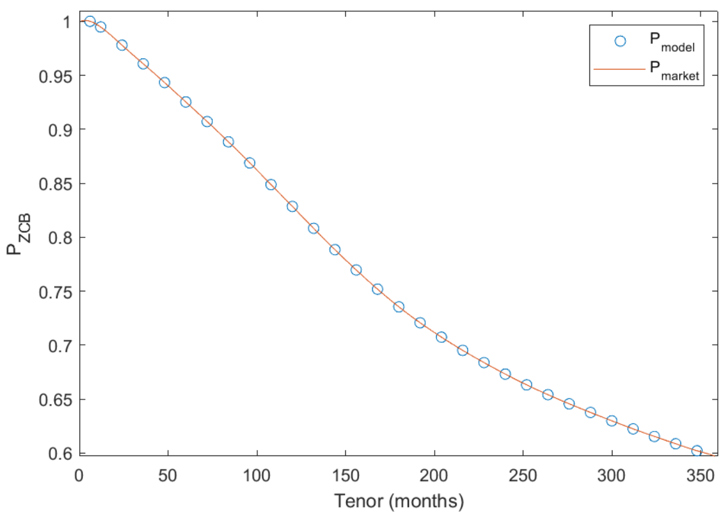

Figure 4: Initial pricing curve based on Euribor of 1/6/2022. The dots are the zero coupon bonds as priced by the HW ATSM, which perfectly lie on the intial pricing curve.

We see that the ZCB prices as computed by the ATSM lie exactly on the initial pricing curve, confirming that our ATSM implementation of the HW model from Section 4.1 is correct, as accidental overfitting is excluded due to the fact that we are only using two parameters for the entire model. This pricing curve can easily be rewritten to the yield curve by $y(T) = -\frac{1}{T}\log P(T)$. Therefore by the exact fitting of ZCB's we have obtained an exact replication of the initial yield curve.

5.2 IR Caps Pricing

Using the Euribor forward curve obtained from the swaps, together with all available IR caps, excluding the 1Y tenors, we calibrate the Hull-White model for constant volatility, by which we obtain the following parameters: $\hat{\alpha} = 0.0441$, $\hat{\sigma} = 0.0098$. Using these two parameters to compute the model prices of all 117 IR caps, we obtain the following fit:

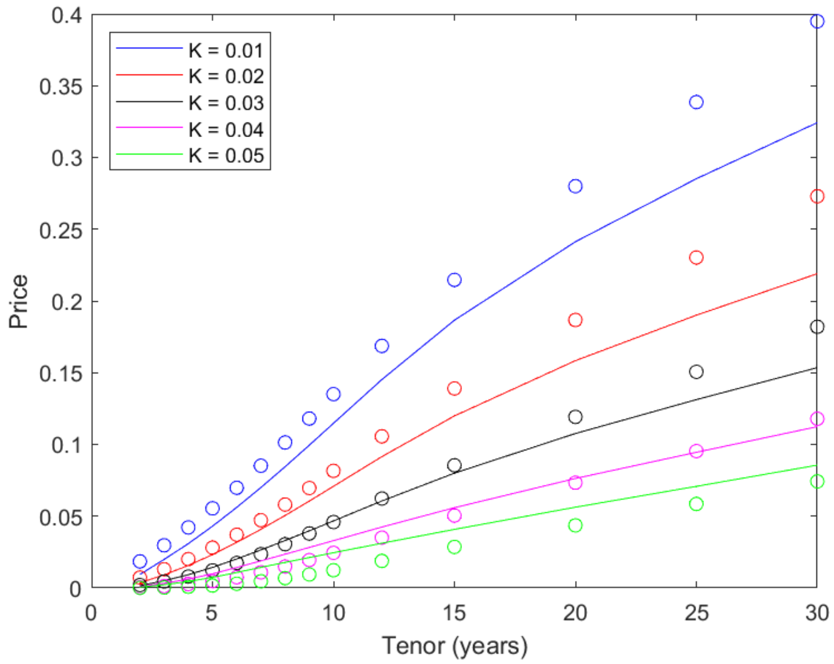

Figure 5: YoY Euribor cap pricing by HW1F using constant σ, calibrated on all available caps, plotted for strikes 1% till 5%. The lines represent the market prices, the dots are the computed model prices.

From this plot we see that for the strikes around 2.5% the caps are modelled fairly accurately, especially considering only two model parameters are used for the pricing of all products. However the higher strike prices show a noticeable downward bias compared to the market prices, and for the lower strikes we see an upward bias. This is most likely due to the volatility smile, as the products that are further ITM (the lower strike prices) are using a value for σ that is too high, whereas the far OTM products are using a value for σ that is too low.

To check whether this is correct, we allow for strike-dependent $\sigma(K)$ and run the optimization process again, by which we obtain the following parameter estimates

Table 3: Parameter estimates for calibration on IR caps using a stike-dependent volatlility $\sigma(K)$.

	$\hat{\alpha}$	$\hat{\sigma}(0.5\%)$	$\hat{\sigma}(1\%)$	$\hat{\sigma}(1.5\%)$	$\hat{\sigma}(2\%)$	$\hat{\sigma}(2.5\%)$	$\hat{\sigma}(3\%)$	$\hat{\sigma}(3.5\%)$	$\hat{\sigma}(4\%)$	$\hat{\sigma}(5\%)$
Coefficient	0.0440	0.0057	0.0058	0.0070	0.0084	0.0097	0.0109	0.0120	0.0130	0.0148

We see that the volatility does not show a smile, but rather a smirk. This is most likely because we are only using caps for the calibration process, which have higher strikes compared to the floors, so that we are only looking at the right side of the volatility smile. Moreover, institutions generally prefer to write calls over puts, making the effect of a skew even stronger (Hoffmann & Fischer, 2012).

Using the parameter estimates for the strike-dependent volatilities from Table 3 we obtain the following model fits:

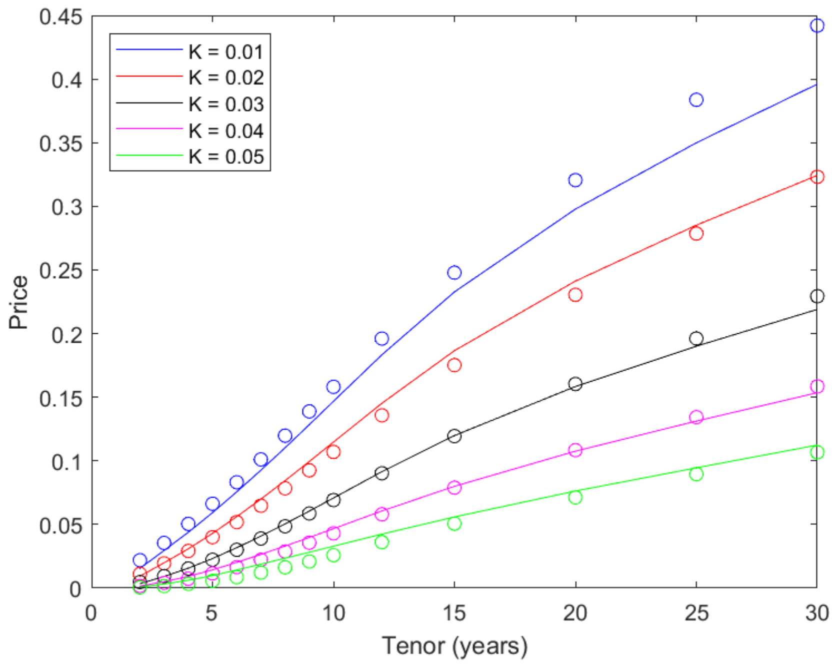

Figure 6: YoY Euribor cap pricing by HW1F using constant α and strike-dependent $\sigma(K)$, calibrated on all available tenors. The lines represent the market prices, the dots are the computed model prices.

We see that the far OTM and ITM products are now better priced compared to the constant volatility case, confirming that including the volatility dynamics positively impacts the estimation.

To quantify this, we look at the weighted RMSE which has decreased from 0.4016 for the constant volatility case to 0.3516 for the strike-dependent volatility calibration, which is a reduction of 12.45%. However, for the amount of extra parameters, which has gone from 2 to 10, this is not as big of a difference as would be preferred, since the addition of parameters brings the risk of over-parameterization. Therefore we choose to stick to the constant volatility for the next steps of the calculation. In addition to this, the interest rates as of 1/6/22 are relatively close to zero, while the inflation is at a 40-year high, meaning that the inflation process plays a much larger role than the discounting process in the valuation of inflation derivatives. However, when interest rates may become significantly different from zero in the future, we could reevaluate our decision to keep the amount of parameters for the IR process at only 2.

5.3 Inflation Caps and Floors Pricing

Using the parameters $\hat{\alpha} = 0.0441$, $\hat{\sigma} = 0.0098$ found by calibration on the IR caps, we now perform a calibration of the inflation process parameters using the inflation caps and floors.

5.3.1 Calibration for Constant Volatility

We start off by calibrating our two-process HW model using constant parameters, thus not yet accounting for the volatility smile. We find the following parameters:

Table 4: Constant parameters from calibration of the inflation HW process for constant volatility.

	$\hat{\alpha}_I$	$\hat{\sigma}_I$	$\hat{\rho}$
Coefficient	0.1149	0.0120	1.000

We see that the correlation parameter reaches its upper bound of 1. Removing the parameter bounds we obtain an even higher correlation of $\hat{\rho} = 1.0281$. While (33) allows for correlations larger than 1, it does not make sense from a logical perspective. Moreover it would give a complex numbered correlation between the Brownian motions of the two processes. We therefore try a different optimization method in which we apply regularization to $\rho_0 = 0.3193$. This value is found after incorporating the full volatility dynamics, as we will see in the next section. Our goal here is to reduce the bias observed in Figure 7 without having to incorporate the volatility dynamics. The results of the regularization are shown in Table 5.

Table 5: Regularization of $\hat{\rho}$ to $\rho_0 = 0.3193$ which is the correlation found after incorporating the full volatility dynamics, shown in Section 5.3.2. The WRMSE is computed for different values of the hyperparameter γ.

	$\gamma = 0.1$	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$	$\gamma = 50$	$\gamma = 100$
$\hat{\rho}$	0.9999	0.9612	0.9317	0.5350	0.4352	0.3438	0.3316
WRMSE	0.3904	0.3909	0.3914	0.4002	0.4030	0.4057	0.4061

As we converge towards ρ_0 the WRMSE gets slightly larger, meaning that even though we obtain a more reasonable correlation parameter, this does not translate to a better fit to the market prices. A possible explanation could be that ρ tries to 'compensate' for any mispricing bias caused by enforcing constant volatility, as we will see in Figure 7. For a more comprehensive picture of the sensitivity of the model fit to the correlation parameter, we refer to Appendix B.

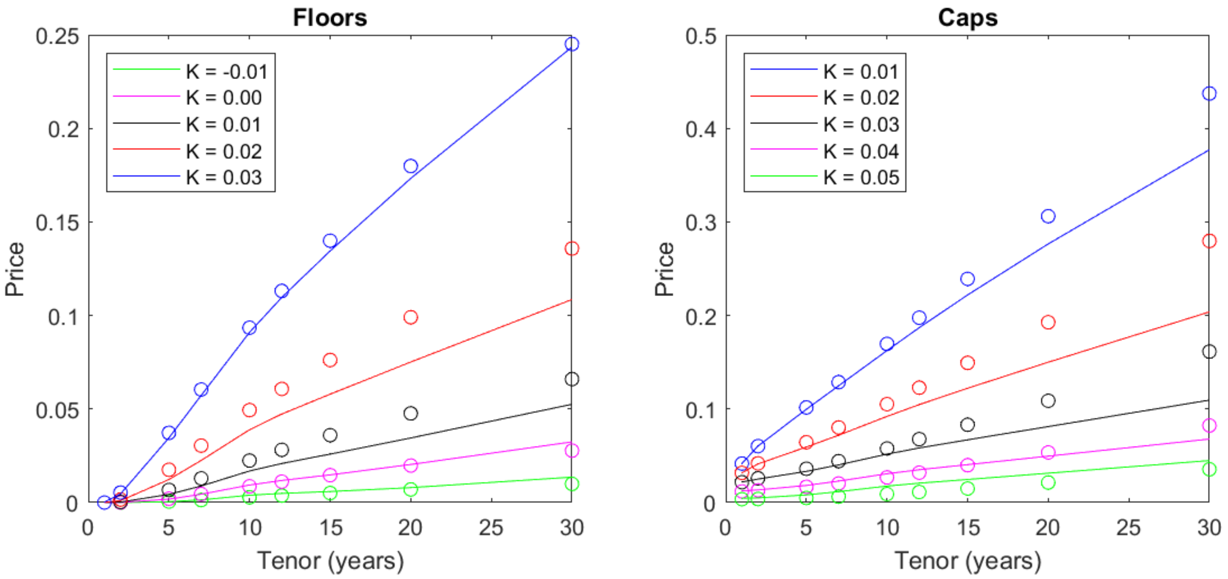

Figure 7: Model prices for YoY HICPxT caps and floors for five different strikes K, as computed by the two-process Hull-White model using constant parameters and no regularization. The lines represent the market prices, the dots are the computed model prices.

For both the caps and the floors there is already a reasonable fit for some of the strike prices. The floors are most accurately priced for the strike around $K = -1\%$ and the caps around $K = 4\%$. This may seem random at first, but in the next section we will see that these results exactly match with our hypothesis of a volatility smile.

To quantify the results above in order to compare them against other calibration methods, we calculate the weighted RMSE, as displayed in Table 6.

Table 6: Weighted RMSE between the market prices and model prices for HICPxT caps and floors, calibrated on constant parameters.

	Floors	**Caps**	**Total**
WRMSE	0.4318	0.3522	0.3904

5.3.2 Capturing the Volatility Smile

For the calibration on constant volatility we saw that for some strike prices there is a clear bias in the pricing. Our hypothesis is that this happens because of the presence of a volatility smile, which we do not take into account. Furthermore, in the sensitivity analysis in Appendix B we see that the model is most sensitive to changes in the volatility parameter. Therefore, we redo the calibration, but now accounting for the volatility smile by allowing for a strike-dependent volatility. The parameters of this calibration are displayed in Table 7.

Table 7: Parameter estimates for calibration on HICPxT caps and floors using a strike-dependent volatility parameter $\sigma_I(K)$, constant α_I and ρ and no regularization.

	$\hat{\alpha}_I$	$\hat{\rho}$	$\hat{\sigma}_I(-1\%)$	$\hat{\sigma}_I(0\%)$	$\hat{\sigma}_I(1\%)$	$\hat{\sigma}_I(2\%)$	$\hat{\sigma}_I(3\%)$	$\hat{\sigma}_I(4\%)$	$\hat{\sigma}_I(5\%)$	$\hat{\sigma}_I(6\%)$
Coeff.	0.1437	0.3193	0.0120	0.0117	0.0102	0.0092	0.0103	0.0128	0.0152	0.0165

Immediately we observe that using this calibration we get a more logical value for the correlation compared to the previous calibration, thus removing the need of regularization. Secondly, from these parameters we see that the average volatility of $\sigma_I(K)$ is 0.0123, which is almost exactly the value of $\sigma_I = 0.0120$ we found when calibrating for constant volatility, confirming that both calibration methods are congruent.

For a clearer picture of the apparent smile, we plot the volatilities of Table 7 against their strike prices on the x-axis, together with the parameter σ_I found in the constant volatiltiy calibration of the previous section, as given in Figure 8 below.

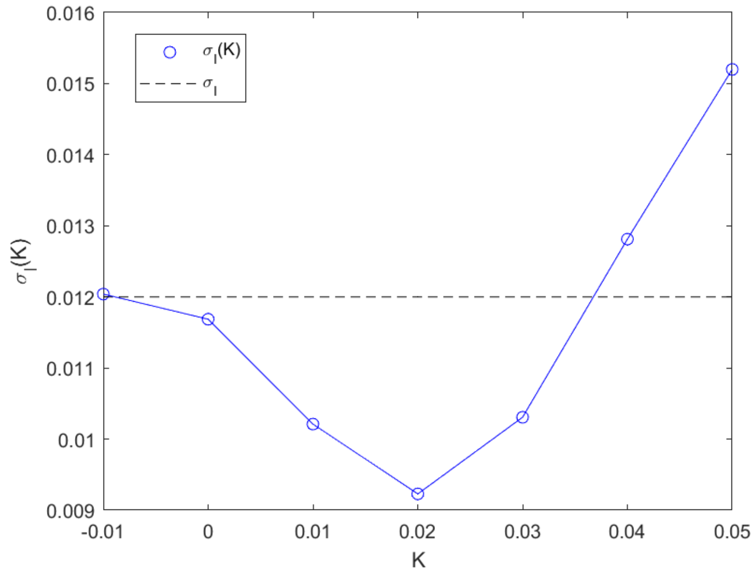

Figure 8: Volatility $\sigma(K)$ calibrated on all HICPxT caps and floors, together with the constant volatility σ_I as obtained from the calibration in section 5.3.1.

Note that the constant volatility obtained in section 5.3.1 intersects ${}_I(K)$ at strikes $K \approx -1\%$ and $K \approx 4\%$. This exactly corresponds with our earlier observations for the constant volatility, were we saw a bias in the model prices except for the floor at -1% and the cap at 4%. To test whether our results have improved, we compute the model prices for all inflation products and compare them to the market prices, as plotted in figure 9

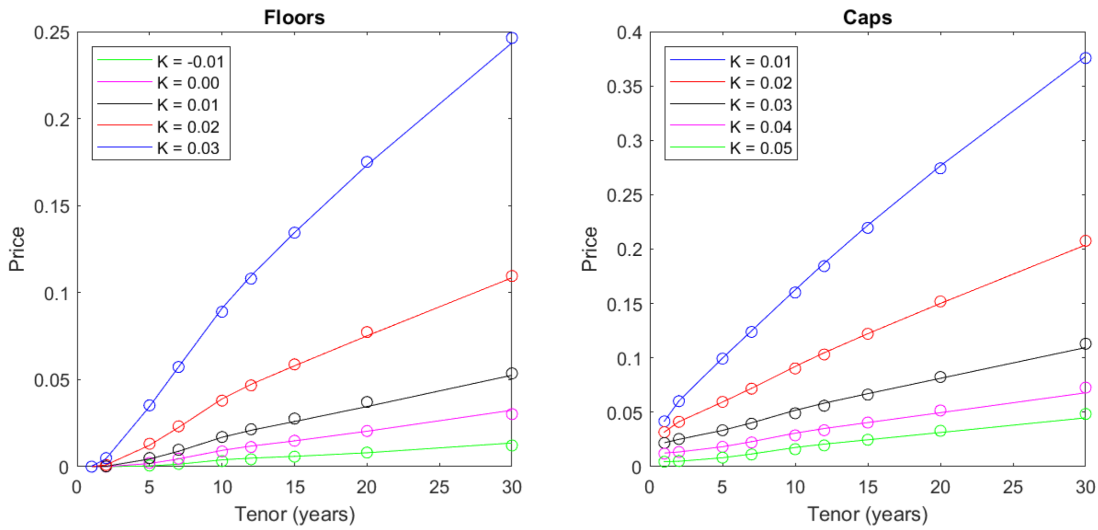

Figure 9: Model prices for YoY HICPxT caps and floors for five different strikes K, as computed by the two-process Hull-White model using a strike-dependent volatility $\sigma_I(K)$. The lines represent the market prices, the dots are the computed model prices.

After incorporating this volatility smile, we immediately see a very clear improvement, with no perceivable bias in any of the products. To confirm this, we again compute the weighted RMSE's as listed in Table 8

Table 8: Weighted RMSE between the market prices and model prices for HICPxT caps and floors, calibrated using a correlated two-process HW model with strike-dependent volatility.

	Floors	**Caps**	**Total**
WRMSE	0.3667	0.0575	0.2508

We see that compared to the constant volatility results from Table 6 all RMSE's have improved quite dramatically, especially for the caps, where the WRMSE has gone down from 0.3522 to 0.0575, which is a 83.7% improvement. The total weighted RMSE has gone down from 0.3904 to 0.2508, which is a 35.8% improvement.

5.3.3 Comparison to the Uncorrelated HW model

Lastly, we compare the performance of our correlated two-process HW model to the uncorrelated two-process HW model as currently in use by NN. First we compare the performances for the model without including the volatility dynamics, of which the performance in terms of the weighted RMSE is given in Table 9.

Table 9: Weighted RMSE between the market prices and model prices for HICPxT caps and floors, using an uncorrelated two-process HW model with constant volatility.

	Floors	**Caps**	**Total**
WRMSE	0.4524	0.3847	0.4168

Comparing these results to the correlated processes of Table 6, we see that based on the WRMSE all 108 products, the modelling performance has decreased by 4.55% for the floors, 8.45% for the caps and 6.33% in total. Thus the correlated two-process HW model is significantly more accurate compared to the uncorrelated model when we use a constant volatility parameter.

Looking at the uncorrelated two-process HW model with strike-dependent volatility we observe the following weighted errors in Table 10

Table 10: Weighted RMSE between the market prices and model prices for HICPxT caps and floors, using an uncorrelated two-process HW model with strike-dependent volatility.

	Floors	**Caps**	**Total**
WRMSE	0.3668	0.0724	0.2531

Comparing these results to the correlated process with strike-dependent volatility in Table 8, we find that that based on the weighted average the uncorrelated model performs the same for the floors, but 25.9% worse for the caps and 0.92% worse in total. Thus overall, the correlated model with strike-dependent volatility performs 0.92% better compared to the uncorrelated model which is currently is use. A percentage point improvement seems insignificant, but for companies dealing with transactions worth hundreds of millions, this can make a large difference.

5.4 Improving the Longer Tenors by Different Weight Functions

Up until now we have exclusively used the normalized weight function for the calibration to market prices as it gives equal weight to all products. In this section we investigate whether we can further improve our estimation by using a different weight function. Since insurance companies are mostly concerned with the valuation of long-term contracts, we specifically look at weight functions that try to improve upon the calibration of the longer tenors. Therefore, we have included six alternative weight functions in Table 11. In order to specifically look at the performance for the different tenors, we split the evaluation for the low (1-7Y) and high (10-30Y) tenors. To avoid a biased outcome in which the evaluation metric is the same as the objective function, we also report the RMSE and mean absolute error (MAE).

Table 11: WRMSE, RMSE and MAE values for different weight functions in the optimization for HICPxT caps and floors using strike-dependent volatility. For simplicity, all products are denoted by f. T is the tenor of the products. The High WRMSE consists of the tenors 10-30Y, the Low are the remaining tenors of 1-7Y. * original weight function for normalization over all products. ** unweighted objective function.

	WRMSE			**RMSE**			**MAE**		
	Low	High	Total	Low	High	Total	Low	High	Total
$w = 1/f^2$ *	**0.3704**	0.0593	**0.2508**	**0.0385**	0.1828	0.1387	**0.0279**	0.1484	0.0948
$w = 1$ **	0.4395	**0.0407**	0.2946	0.0712	0.1217	**0.1024**	0.0520	0.1031	**0.0803**
$w = 1/f$	0.3769	0.0446	0.2535	0.0419	0.1754	0.1337	0.0312	0.1381	0.0906
$w = T/f^2$	0.3743	0.0520	0.2526	0.0464	0.2762	0.2082	0.0296	0.1884	0.1178
$w = T/f$	0.3882	0.0411	0.2606	0.0472	0.1575	0.1216	0.0360	0.1226	0.0841
$w = T^2/f^2$	0.3751	0.0511	0.2530	0.0496	0.2739	0.2068	0.0316	0.1861	0.1174
$w = T$	0.4765	0.0463	0.3195	0.0825	**0.1194**	0.1046	0.0605	**0.0997**	0.0822

Indeed, we see that the total errors favor their equivalent objective function in the first two weight functions. However, when looking specifically at the higher tenors we do not observe this effect as strongly. Looking at the average of three evaluation metrics, $w = 1$ performs the best on the higher tenors, with an average improvement of 31.8%. This is closely followed by $w = T$, with an average improvement of 30.8%. Taking into account the losses in performance for the lower tenors, the unweighted function $w = 1$ performs on average 36,0% worse compared to the original weight function $w = 1/f^2$, and $w = T$ performs 43.1% worse. For both weight functions this is a significant trade-off to be made. Thus, for a balanced performance over all tenors we still recommend using the normalized weight function $w = 1/f^2$. However, when prioritizing the longer tenors, we recommend using the unweighted function $w = 1$, of which modelling performance against the market prices is plotted in Appendix C.

5.5 Out-of-sample Performance

To further evaluate the robustness of the model, we examine the out-of-sample performance of the full model, and compare it to the uncorrelated model. It often occurs that the insurance company wishes to price a product with a certain strike price or tenor, for which no market prices exist. Therefore, we will explore two common situations.

First, we look at the scenario in which the insurance company wishes to price a 6% cap while only having data for strike prices up to 5%. To recreate this scenario we remove the 6% cap from the optimization process and compare the outcomes to the prices of the model calibrated on the full data set in Figure 10.

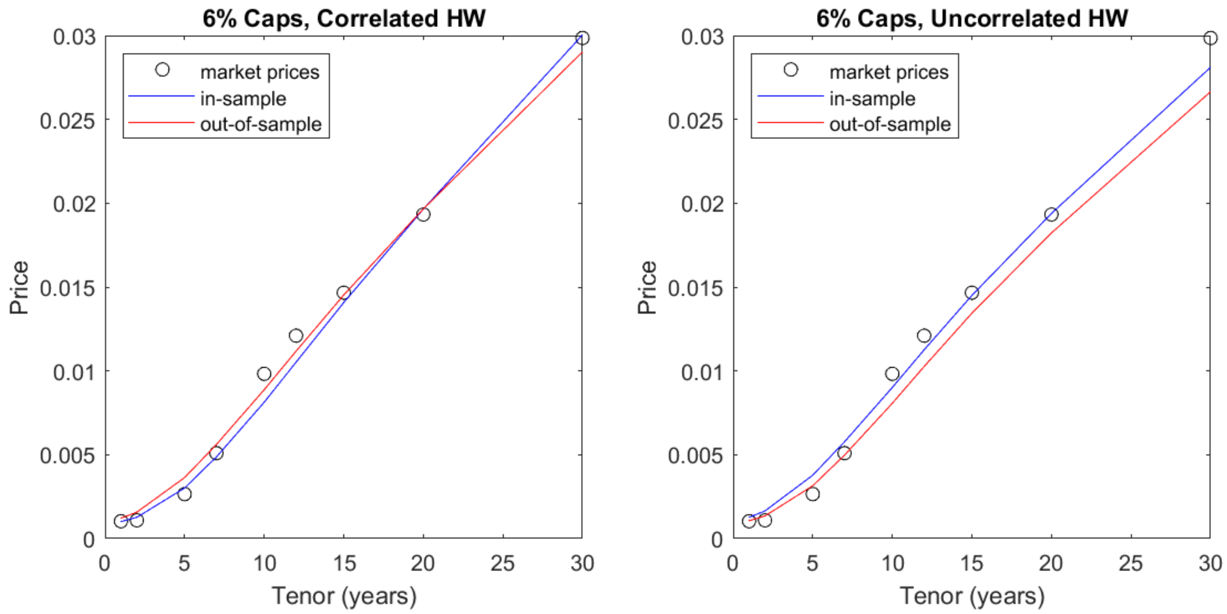

Figure 10: Comparison of in-sample and out-of-sample performance on the model prices of all 6% caps, as computed by the correlated two-process Hull-White model with strike-dependent volatility.

We see that the out-of-sample performance of the correlated HW model is surprisingly accurate, given that the model has no prior information about strike prices above 5%. For the in-sample caps we have an average deviation of 5.40% from the model prices, and 5.61% for the out-of-sample caps. This marginal difference can partly be explained by the hump around 10Y, which the out-of-sample model accidentally prices better due to the upwards bias at the beginning. The model also performs significantly better compared to the uncorrelated model, which has an average in-sample error of 6.52% and an average out-of-sample error of 10.47%.

Second, we look at the situation in which the insurance company wishes to price a product with a tenor of 30 years, but has market information available for tenors up to 20 years. To recreate this scenario we exclude all 30Y tenor caps and floors from our optimization process. In Figure 11 and compare the out-of sample performance to model calibrated on the full data set.

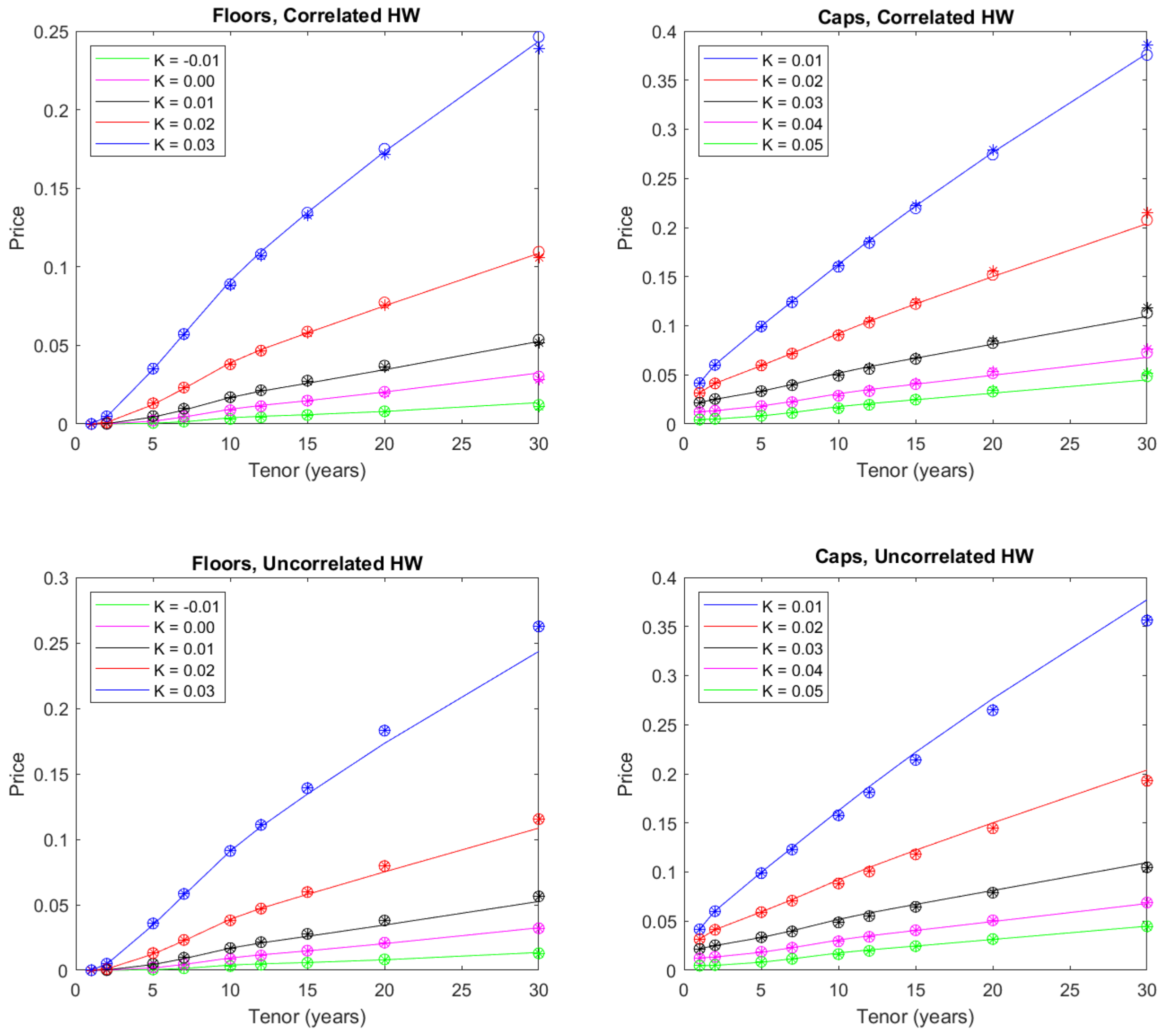

Figure 11: Comparison of in-sample and out-of-sample performance with and without including the 30Y tenors. The circles represent the market prices as calibrated on the full data set. The asterisks represent the out-of-sample fit and the lines represent the market prices.

Looking specifically at the 30Y products from the correlated model, we obtain an average absolute error of 1.91% for the in-sample floors and 2.14% for the in-sample caps. For the out-of-sample fit we obtain an average absolute error of 3.13% for the floors and 5.27% for the caps. Although these out-of-sample errors are still relatively small, it is advised for large companies to be careful, as the potential for mispricing is significantly greater. Compared to the uncorrelated model we obtain an in-sample error of 6.81% for the floors and 4.73% for the caps. For the out-of-sample caps and floors we obtain an average error of 6.99% and 4.68% respectively. This is both significantly higher than for the correlated model.

6 Conclusion

During the writing of this thesis, inflation has continued to rise, forcing insurance companies to further increase their exposure in inflation-linked caps and floors, underlining the importance of an accurate pricing model for these products. To improve the current pricing, and especially the calibration methods, we have implemented a correlated two-process Hull-White model using different optimization techniques.

For capturing the interest-rate dynamics, we first calibrated the IR process to Euribor caps, floors and swaps, using a standard closed-form solution. After analyzing the outcomes of two different optimization techniques, we decided not to incorporate the volatility smile, as this would increase the amount of parameters from 2 to 10, while only yielding a marginal improvement in modelling accuracy.

Using the two parameters from the IR process, we obtained a full closed-form solution for the pricing of inflation-linked caps and floors, which we were able to express in a familiar Black-Scholes form after some tedious algebra. First calibrating this process on market data using constant parameters and normalized weights, we find a reasonable fit for the ATM products, but observe a clear bias in the far ITM and OTM products. We also observe that our correlation parameter converges to the upper limit of 1, implying a perfect correlation between the inflation and interest rate process, which is very unlikely. Using regularization we can shrink the correlation to a more logical level, at the cost of a slightly worse model fit.

After a sensitivity analysis we find that the largest improvements in performance can be obtained by accurate estimation of the volatility parameter. Therefore we have chosen to focus on incorporating the volatility dynamics. We chose to use a strike-dependent volatility parameter instead of the more standard time-dependent volatility approach. By doing so, we obtain a near-perfect fit to the market prices. The advantage of this approach is that the solution is still fully analytical, allowing for efficient computing while still capturing the full volatility dynamics, and remaining arbitrage-free for normal caps and floors. Only when pricing more exotic products which depend on two or more strike prices, we are not able to guarantee risk-neutral pricing.

By testing different weight functions in the optimization process, we find that we can significantly improve the model accuracy for the long tenors, while only sacrificing a few percentage points for the lower tenors. This can be very useful for insurance companies, which are often only interested in the longer tenors. However, when overall performance is the main objective, the standard normalized weight function is still the preferred choice. The model also performs very well on the out-of-sample modelling of higher tenors and strike prices. However, some caution is still advised, as the potential for mispricing is significantly greater when modelling out-of-sample.

Lastly, we find that based on the weighted average, our model performs 0.92% percent better than the uncorrelated two-process HW model, which NN currently uses. Moreover it also performs significantly better for out-of-sample pricing. Therefore, we recommend using our correlated two-process HW model with strike-dependent volatility over the current model, as even a single percentage point difference can have a significant impact when dealing with large transactions.

There are still things that could be improved upon. Next to the strike-dependent volatility model, one could also implement a time-dependent volatility model by for example using a local volatility model together with Monte Carlo simulation. This way we can make a true quantitative comparison between the two methods in terms of model accuracy and computation times.

Another interesting area of research could be to quantify the exact trade-off between the model accuracy and the increasing number of parameters when incorporating the volatility smile. This could for example be done by deriving a likelihood function, such that we can compute the AIC and BIC. However, this would require time series data for all 108 products, which is very difficult to obtain for OTC products.

7 References

Backus, D. K., Foresi, S., & Wu, L. (2004). Accounting for biases in black-scholes. *Available at SSRN 585623*.

Barone, E., & Castagna, A. (1997). The information content of tips. *Available at SSRN 2170675*.

Blanchard, A. (2014). The two-factor hull-white model: pricing and calibration of interest rates derivatives. *KTH Royal Insitute of Technology*.

Brigo, D., & Mercurio, F. (2006). *Interest rate models-theory and practice: with smile, inflation and credit* (Vol. 2). Springer.

Brigo, D., Mercurio, F., & Sartorelli, G. (2003). Alternative asset-price dynamics and volatility smile. *Quantitative Finance*, *3*(3), 173.

Canabarro, E., & Duffie, D. (2003). Measuring and marking counterparty risk. *Asset/Liability Management for Financial Institutions, Institutional Investor Books*.

Chou, R. Y. (1988). Volatility persistence and stock valuations: Some empirical evidence using garch. *Journal of applied econometrics*, 279–294.

David, A. (2008). Inflation uncertainty, asset valuations, and the credit spreads puzzle. *The Review of Financial Studies*, *21*(6), 2487–2534.

Dodgson, M., & Kainth, D. (2006). Inflation-linked derivatives. *Royal Bank of Scotland Risk Training Course, Market Risk Group*.

Filipović, D. (2000). *Consistency problems for hjm interest rate models* (Unpublished doctoral dissertation). ETH Zurich.

Fleming, M. J., & Sporn, J. (2013). Trading activity and price transparency in the inflation swap market. *Economic Policy Review*, *19*(1).

Gurrieri, S., Nakabayashi, M., & Wong, T. (2009). Calibration methods of hull-white model. *Available at SSRN 1514192*.

Heath, D., Jarrow, R., & Morton, A. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, 77–105.

Hoffmann, A. O., & Fischer, E. T. S. (2012). Behavioral aspects of covered call writing: an empirical investigation. *Journal of Behavioral Finance*, *13*(1), 66–79.

Hughston, L. (1998). *Inflation derivatives* (Tech. Rep.). Working Paper, Merrill.

Hull, J., & White, A. (1990). Pricing interest-rate-derivative securities. *The review of financial studies*, *3*(4), 573–592.

Hull, J., & White, A. (1993). One-factor interest-rate models and the valuation of interest-rate derivative securities. *Journal of financial and quantitative analysis*, *28*(2), 235–254.

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

Rebonato, R. (2012). Modern pricing of interest-rate derivatives. In *Modern pricing of interest-rate derivatives.* Princeton University Press.

Ritchken, P., & Sankarasubramanian, L. (1995). Volatility structures of forward rates and the dynamics of the term structure 1. *Mathematical finance*, *5*(1), 55–72.

Singor, S. N., Grzelak, L. A., Van Bragt, D. D., & Oosterlee, C. W. (2013). Pricing inflation products with stochastic volatility and stochastic interest rates. *Insurance: Mathematics and Economics*, *52*(2), 286–299.

Sircar, K. R., Papanicolaou, G. C., et al. (1999). Stochastic volatility, smile & asymptotics. *Applied Mathematical Finance*, *6*, 107–145.

Webb, R. H. (1995). Forecasts of inflation from var models. *Journal of Forecasting*, *14*(3), 267–285.

Appendices

A Full Euribor Caps and Floors Data

Table 12: YoY Euribor caps and floors market prices from Tullett as of 1/6/2022. The percentages displayed horizontally are the strikes with respect to the Euribor. The years displayed vertically are the tenors. The market prices are given in basis points of the nominal.

	Cap									**Floor**				
	0.5%	1%	1.5%	2%	2.5%	3%	3.5%	4%	5%	0%	0.5%	1%	1.5%	2%
1Y	28.68	8.770	2.210	0.650	0.230	0.100	0.050	0.020	0.010	0.800	8.500	7.600	25.00	47.70
2Y	155.6	93.87	55.93	34.80	23.03	16.15	11.88	9.070	5.760	4.710	16.30	22.80	58.20	108.9
3Y	295.5	199.0	132.9	91.86	66.56	50.41	39.58	32.03	22.33	11.30	31.40	48.30	103.4	182.2
4Y	438.6	308.8	216.7	156.8	118.2	92.43	74.59	61.72	44.74	22.00	51.30	80.10	155.9	262.0
5Y	592.9	430.3	312.0	232.5	179.5	143.3	117.7	98.83	73.47	33.80	73.00	114.7	210.2	342.3
6Y	756.2	561.3	416.9	317.3	249.4	202.1	168.0	142.6	108.0	45.70	96.00	150.0	265.8	423.8
7Y	926.5	699.9	529.1	408.6	324.6	265.0	221.5	188.7	143.4	58.30	119.0	185.4	319.8	501.5
8Y	1103	844.6	647.2	504.9	403.7	330.6	276.4	235.3	178.1	70.30	139.8	218.0	369.3	572.6
9Y	1284	994.7	770.1	605.4	485.9	398.2	332.5	282.3	211.9	82.00	159.4	249.0	415.9	639.3
10Y	1468	1147	895.9	708.3	570.1	467.2	389.5	329.7	245.5	93.50	178.1	278.9	460.5	702.9
12Y	1832	1452	1147	915.2	739.5	606.2	503.9	424.4	311.7	116.2	212.8	336.0	545.1	822.7
15Y	2327	1865	1491	1199	973.8	800.2	665.2	559.4	408.2	158.0	274.1	434.6	686.8	1018
20Y	2978	2412	1949	1584	1298	1076	901.4	762.9	562.8	270.3	429.9	666.5	1001	1429

B Model Sensitivity to Individual Parameters

From the regularization of the correlation parameter ρ to the empirical correlation $\rho_0 = 0.3193$ we found that the WRMSE increases as ρ approaches 1. However, this increase in WRMSE seems to be relatively small. Therefore we investigate the sensitivity of the model when varying the correlation parameter between its bounds of -1 and 1, while keeping all other parameters from Table 4 constant, as is displayed in Figure 12 below

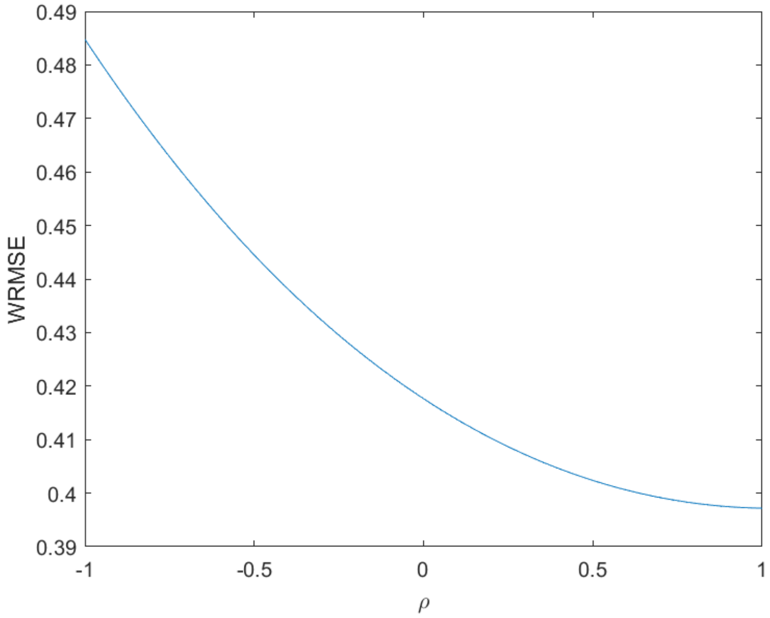

Figure 12: Sensitivity analysis based on the weighted RMSE for the correlation parameter ρ over its whole range from -1 to 1. The parameters α_I and σ_I are taken from Table 4.

From this plot we confirm that the WRMSE is a clearly decreasing function around $\rho_0 = 0.3193$, thus confirming our observations that applying regularization does not improve performance.

To compare the sensitivity of rho to the other parameters, we make the same plot for α_I and σ_I, as plotted in Figures 13 and 14.

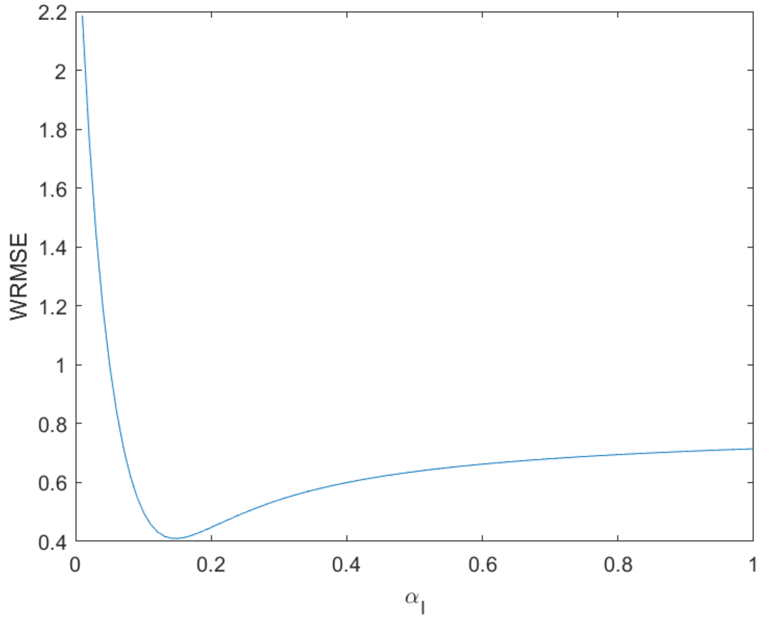

Figure 13: Sensitivity analysis based on the weighted RMSE for the correlation parameter α_I over a selected range from 0 to 1. The parameters σ_I and ρ are taken from Table 4.

We see a clear minimum at 0.147, which corresponds to our findings in Table 4. Moreover we see that the accuracy of the model is highly sensitive to the left of this optimum.

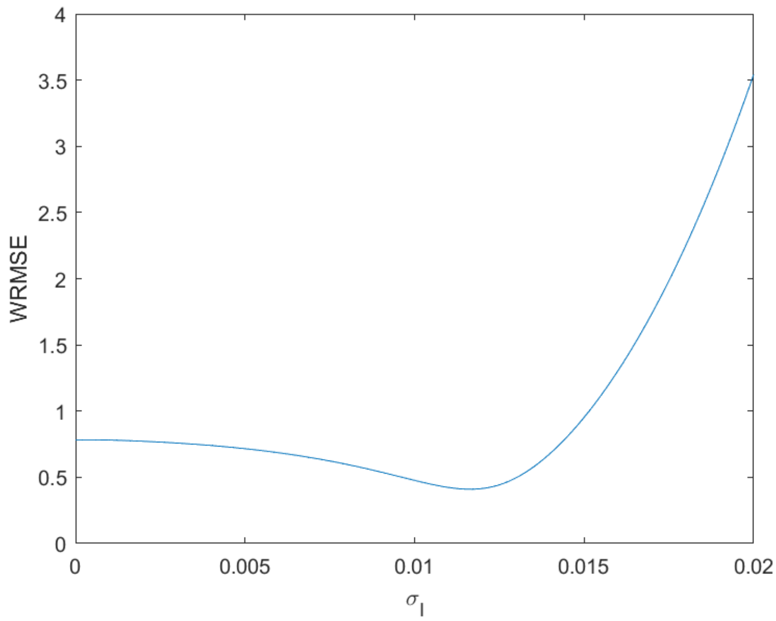

Figure 14: Sensitivity analysis based on the weighted RMSE for the correlation parameter σ_I over a small selected range from 0 to 0.02. The parameters α_I and ρ are taken from Table 4.

From these three figures we see that the model is the least sensitive to the correlation parameter and most sensitive to the volatility parameter, as the average weighted error increases rapidly for small increases in the volatility. This confirms our concerns of properly modelling the volatility and focusing less on α_I or the correlation parameter.

C Model prices for the unweighted optimization

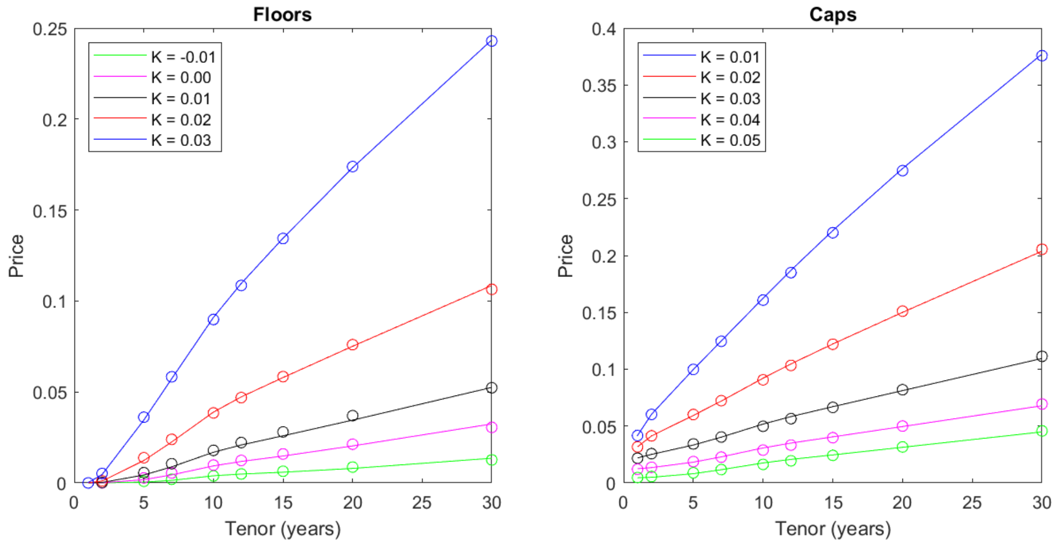

Figure 15: Model prices for YoY HICPxT caps and floors for five different strikes K, as computed by the unweighted two-process Hull-White model using a strike-dependent volatility $\sigma_I(K)$. The lines represent the market prices, the dots are the computed model prices.

We see that both caps and floors are very accurately priced, especially for the higher tenors.