



**Valuation of Inflation-Indexed Derivatives Using
Path-Dependent Simulation**

Delft University of Technology
Faculty of Electrical Engineering, Mathematics and Computer Science

Thesis submitted to the
Delft Institute of Applied Mathematics

Thesis committee

dr. ir. J. Bierkens (supervisor)

ir. P. L. Ruitenber (KPMG supervisor)

dr.ir F. Fang

H. Vermeer

**Delft, Nederland
December 2024**

Copyright © 2024 by Hennes Vermeer. All rights reserved.

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Joris Bierkens, for his consistent guidance, insightful feedback, and challenging me throughout the past nine months. The same thanks go to my second supervisor, Patrick Ruitenbergh at KPMG, who provided constant support and for allowing me to help KPMG with this journey. Their expertise, and constructive sessions have been instrumental in shaping this thesis. Much love to my family and girlfriend who have supported me during this period as well.

Abstract

A vital aspect of managing inflation risk is the use of inflation-indexed derivatives. Currently, inflation-indexed bonds and swaps are the primary instruments purchased by institutions. Inflation options (also known as inflation caps/floors) are also available in the market. Risk-neutral pricing of these derivatives is a difficult challenge due to the connection between inflation and interest rates.

In this thesis, the Heston model and its extensions to stochastic interest rates are investigated in the context of inflation-indexed derivatives. First, existing analytical pricing formulas and simulation methods are summarized. Then the multilevel Monte Carlo (MLMC) method is applied as a potent variance reduction technique. For the standard Heston model, the MLMC method reduces the computation costs by a factor of 10 to 50 for short maturities. The Python code implementing the applied methods is also published.

Contents

1	Introduction	1
1.1	Objectives of the thesis	1
1.2	Organisation of the thesis	2
2	Background	4
2.1	Inflation-Indexed derivatives	6
2.2	Pricing Formulas for Inflation-Indexed Derivatives	8
3	Model Overview	10
3.1	Black-Scholes	10
3.2	Jarrow-Yildirim framework	11
3.3	Heston framework	13
3.3.1	Constant interest rate	13
3.3.2	Hull-White model for interest rates	15
3.3.3	Libor Market Models for forward rates	17
3.3.4	Multi-Factor SABR model	20
3.4	Exact Solution Using the Cos Method	22
4	Monte Carlo Simulation	24
4.1	Simulation of the Heston model	25
4.1.1	Euler discretization	25
4.1.2	Exact simulation	26
4.1.3	NCI scheme	28
4.1.4	QE scheme	29
4.2	Simulation of the interest rates	30
4.3	Multilevel Monte Carlo (MLMC)	32
4.4	Convergence properties	35
4.5	MLMC for path-dependent options using the Heston model	35
5	Numerical experiments	38
5.1	Valuing an inflation cap using the Heston model	38
5.2	MLMC vs standard Monte Carlo using the Heston model	39
6	Conclusion & Discussion	42
6.1	Conclusion	42
6.2	Discussion	42
A	Appendix	47
A.1	Proof of Theorem 1	47
A.2	Calibrations performed in Literature	49
A.3	MLMC for the Heston Hull-White model	53

1 Introduction

Inflation-indexed derivatives play a vital role in modern financial markets, offering investors and institutions tools to hedge against inflation risk and manage their portfolios. Since their introduction in the 1980s, following the period of high inflation in the previous decades, these derivatives have become increasingly important as inflation spikes continue to occur in global economies. KPMG has a vested interest in making sure that its clients are sufficiently aware of inflation risk, the ways they can hedge themselves, and the models they can use.

Currently, insurers and pension funds invest most of their money in stocks, bonds, and real estate [40]. Due to volatility in the capital markets, the value of investments is uncertain over time. Inflation rates and interest rates, in particular, significantly impact the ability of pension funds to cover their future liabilities. As a result, pension funds have a high ambition and are required by law to understand the risks associated with inflation. Inflation-indexed derivatives, such as inflation-indexed bonds and options, offer solutions that let them manage inflation risk.

The risk-neutral pricing of inflation-indexed derivatives poses significant challenges due to the complex structure between inflation rates, interest rate, and model parameters. Various models and methodologies have been developed to address these challenges, including the famous Black-Scholes model. Extensions of the Black-Scholes [4] model have become very important in capturing essential properties, such as the Heston model [27] which models the variance process with a Cox-Ingersoll-Ross[15] process. While this model allows for efficient and exact pricing, further extensions in the context of inflation-indexed derivatives require potentially undesirable assumptions to obtain closed-form solutions. Fourier-based methods, such as the COS method [19], are challenging to apply to models like the Heston Hull-White model due to the complex expressions required for its characteristic function.

Many papers have focused on the use of simulation techniques to obtain the solutions instead. For example, Glasserman [24], a classic book on Monte Carlo methods in finance, Broadie and Kaya [9], Lord et al. [34], and Giles [23]. These papers show that certain discretization methods can significantly improve the computation speed of prices of financial derivatives. Furthermore, improved estimators in the general Monte Carlo simulation framework result in even faster algorithms. Although much attention has been paid to applying these techniques to stochastic variance such as the Heston model, little attention has been given to the application of simulation methods to a combination of stochastic variance and stochastic interest rates.

1.1 Objectives of the thesis

The primary objective of this thesis is to develop accurate and efficient methods for pricing inflation-indexed derivatives. Achieving high accuracy in derivative pricing involves minimizing errors associated with the numerical approximation of stochastic models. In this context, accuracy is increased by:

- Reducing the discretization error introduced by discretization of a continu-

ous process. Advanced discretization schemes are explored to achieve better convergence.

- Minimizing the statistical error in the Monte Carlo estimators, either by increasing the number of samples or employing variance reduction techniques.

To start with, the Heston model [27] is a well-researched model for assets and interest rates, on which many of the techniques can be applied in the context of inflation derivatives. Subsequently, more complex models, such as the Heston Hull-White model[43], will be analyzed. Currently, the pricing structures are complicated and highly model dependent. This thesis aims to propose a broadly applicable simulation framework that can be extended to future developments, and different payoff structures.

The objectives of this thesis are as follows:

- Investigate techniques for pricing inflation-indexed derivatives, particularly in the context of multi-factor models.
- Review the state of the art methods for simulating stochastic processes used in derivative pricing.
- Compare the existing discretization schemes, with a focus on balancing computational efficiency and accuracy.
- Demonstrate the applicability of multilevel Monte Carlo[23] as a flexible and efficient method for handling path-dependent payoffs, with an emphasis on reducing statistical error.
- Deliver an extensible Python implementation of the described methods to facilitate reproducibility and further research.

1.2 Organisation of the thesis

Firstly, in Chapter 2 some context on inflation, inflation-indexed derivatives, and the pricing thereof.

Chapter 3 will give a complete overview of existing research on models related to inflation-indexed derivatives. This involves the Black-Scholes model, the Jarrow-Yildirim model, the (extended) Heston model, and LIBOR market models.

In Chapter 4 several simulation methods are discussed, which are applicable to the models observed in Chapter 3. Several of the existing discretization schemes from the literature are introduced. The focus of the chapter is on a variance reduction technique called multilevel Monte Carlo introduced by Giles [23]. Particular attention is given to the application of multilevel Monte Carlo to inflation-indexed derivatives.

Chapter 5 contains the numerical results obtained from the application of the techniques in Chapter 4. The simulation techniques are tested on the Heston model for which an exact solution is known in closed form.

Finally, the last chapter contains the conclusion and discussion.

2 Background

Inflation-indexed derivatives are financial instruments designed to manage inflation risk. These derivatives include various products such as inflation swaps, options, and inflation-indexed bonds. The main function of these instruments is to provide protection against inflation fluctuations, allowing investors and institutions to stabilize their portfolios.

Inflation refers to the rate at which the general level of prices for goods and services rises over time. Central banks, such as the European Central Bank, closely monitor and target inflation rates to maintain economic stability. Continuously high inflation can lead to increased uncertainty in financial markets, influence interest rates, and distort the valuation of assets and portfolios.

An inflation rate is usually measured by a consumer price index (CPI) which reflects the actual price level of a basket of typical consumer goods [8]. In the Eurozone this is a harmonized consumer price index (HCPI) determined by EUROSTAT, the statistical office for the European Union (see Figure 2.1). In the USA, the Bureau of Labor Statistics publishes the consumer price index (CPI-U) which is chosen to represent the expenditure of urban residents.

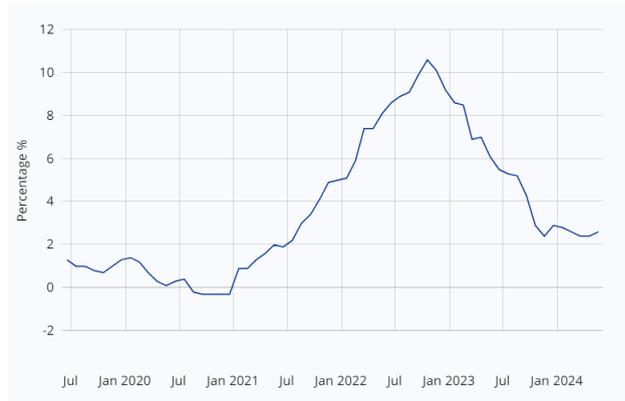


Figure 2.1: Eurozone CPI index percentage change per year in the last 5 years.

With the role central banks play in stabilizing inflation and investors asking for compensation to cover inflation when investing into bonds, there is a clear connection between interest rates and expectations on the future inflation rate. Studies into this relation date back to Fisher in 1930 [21] who described the relation between interest rates and inflation with the so-called Fisher equation

$$(1 + r_n) = (1 + r_r)(1 + I), \quad (2.1)$$

where r_N is the nominal interest rate, r_R the real interest rate, and I the inflation rate. Historically the following equation tends to function as a close approximation,

$$r_n = r_r + I.$$

Since the inflation rate is measured by a change in the consumer price index, the

inflation rate can be represented in terms of the inflation index using the following,

$$\frac{I(T) - I(t)}{I(t)}. \quad (2.2)$$

Many articles on inflation modeling use a so-called foreign currency analogy which uses this relationship between inflation and interest rates [7], [30]. In this analogy, real interest rates are viewed as interest rates in a "foreign" economy, while the inflation index is interpreted as the exchange rate between the nominal (i.e., domestic) and the real (i.e., foreign) economy. For example, if the index increases in a certain time frame, more nominal currency is required to buy the same amount of goods and services. A unit of money at time T is worth $\frac{I(0)}{I(T)}$ in the real currency, with $I(0)$ and $I(T)$ the value of the inflation index at times 0 and T , respectively. The interest rate in the real economy is called the real rate and reflects the true cost of borrowing or investing adjusted for inflation. It is important to note that this real rate is not fixed, it is based on expected inflation index changes and is only known for certain at the end of the time period. Thus, both the real interest rate and the inflation index should be modeled along with the nominal interest rate.

Quantities related to the nominal economy are denoted by subscript n and for the real economy by subscript r . The usual real-world probability space is denoted by $(\Omega, \mathcal{F}, \mathbb{P})$, while \mathbb{Q}_n and \mathbb{Q}_r denote the risk-neutral measures in their respective economies. The risk-neutral measures differ because the no-arbitrage principle leads to a different discount factor in each economy, resulting in a different risk-neutral valuation of derivative payoffs. Both sets of computations occur within the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$, however, under the real risk-neutral measure the cash flows are adjusted for inflation in contrast to the nominal risk-neutral measure. To bring the dynamics of a model under the same measure, the change of measure technique using the Radon-Nikodym derivative can be used. For a broader insight into this concept, see, for example, Brigo and Mercurio [7, p. 45], Michael [38, p. 365] and Björk [3, p. 396].

This section concludes with some definitions and notation used throughout this thesis. The term structures of the nominal and real interest rate curves at time t are denoted by $P_n(t, T)$ and $P_r(t, T)$ respectively. $P_n(t, T)$ and $P_r(t, T)$ can be viewed as the value of a nominal/real zero-coupon bond with maturity T at time t , $t \leq T$. The respective forward interest rates are defined by the constant rate at which an investment of $P_n(t, T)$ or $P_r(t, T)$ units of currency at time t accrues continuously to yield a unit amount of currency at maturity T :

$$f_x(t, T) = -\frac{\partial \ln P_x(t, T)}{\partial T}, \quad x \in \{n, r\}.$$

The nominal and real instantaneous short rates, which represent at each instant the initial point of the forward rates are then denoted by

$$\begin{aligned} r_n(t) &= \lim_{T \rightarrow t^+} f_n(t, T), \\ r_r(t) &= \lim_{T \rightarrow t^+} f_r(t, T). \end{aligned}$$

Throughout financial modeling, both forward rates and short rates are employed to price derivatives.

Given future times T_{i-1} and T_i , the forward London Interbank Offered Rate (LIBOR) rates at time t are defined by

$$F_x(t, T_{i-1}, T_i) = \frac{P_x(t, T_{i-1}) - P_x(t, T_i)}{\tau_i P_x(t, T_i)}, \quad x \in \{n, r\}, \quad (2.3)$$

where τ_i is the year fraction for the interval $[T_{i-1}, T_i]$, which is assumed to be the same for both nominal and real interest rates.

2.1 Inflation-Indexed derivatives

Inflation-indexed derivatives were developed to manage exposure to inflation risk. These derivatives are typically tied to a specific inflation index such as the Consumer Price Index. By linking the derivative's value to the inflation index, investors can hedge against unexpected changes in inflation. The most commonly purchased derivatives are government issued inflation-indexed bonds which have their principal and interest payments adjusted according to an inflation index. There is a varied history of the government issuance of such bonds [17] [10]. Treasury Inflation-Protected Securities (TIPS) linked to the CPI-U were introduced by the US in 1997 (the EU followed closely in 1998), with 300 billion USD outstanding in 2005 with maturities up to thirty years and an estimated 2.8 trillion outstanding worldwide in 2024.

There are many reasons an institution would buy an inflation-indexed bond. Most pensions increase from year to year with inflation, sometimes under statutory requirements. Therefore, a pension fund or a life insurer can match these liabilities by buying these bonds. Furthermore, compared to conventional riskless bonds whose real value declines over time, the effective duration of an inflation-indexed bond is much longer. This is useful for investors who want guaranteed cashflows thirty years in the future.

Similarly to a standard zero-coupon bond (a bond without interest payments) which considers the term structure of interest rates, a zero-coupon inflation-indexed bond pays the value of a price index $I(T)$ at maturity T . At an earlier time t , the bond has an arbitrage-free price of $P_I(t, T)$, with no delay between payment and the publication of the inflation index. In principle, the present value of inflation-indexed bonds in the market can be used to determine the curve of $P_I(t, T)$. Unfortunately, this is not practical as there is only a limited supply of inflation-indexed bonds for different maturities. Currently, the best way to extract the market value of $P_I(t, T)$ is by using the prices of inflation swaps.

The inflation-indexed derivatives market is dominated by zero-coupon inflation-indexed swaps (ZCIIS) due to their simplicity. Inflation-indexed swaps can be tailored to match the requirements of a clients future liabilities more closely than with bonds. These swaps are traded over-the-counter between banks, investors, pension

funds, and insurers to transfer inflation risk. A zero-coupon swap involves a fixed payment of

$$(1 + K)^{T_M} - 1 \tag{2.4}$$

at the final time $T_M = M$ years, in exchange for receiving the relative increase of an index

$$\frac{I(T_M) - I(0)}{I(0)} = \frac{I(T_M)}{I(0)} - 1.$$

The fixed rate K is usually quoted, e.g. on Bloomberg. The value of the zero-coupon swap is zero when

$$(1 + K)^{T_M} - 1 = \frac{I(T_M)}{I(0)} - 1.$$

Another important inflation-indexed swap is the year-on-year swap inflation-indexed (YYIIS) where at each time T_i the fixed amount K is exchanged for the average rate of inflation over the previous year with $T_0 := 0$

$$\frac{I(T_i) - I(T_{i-1})}{I(T_{i-1})} = \frac{I(T_i)}{I(T_{i-1})} - 1.$$

In contrast to the zero-coupon swap, the present value of a year-on-year swap cannot be written in terms of zero-coupon inflation bonds and requires a model to be valued. This is shown in the next section.

In a swap contract, both parties are obligated to fulfill their cashflow obligations as per the terms of the contract. However, similar to assets, it is also possible to define options on the inflation index to provide additional flexibility. An inflation-indexed cap (IIC) is a call option on an inflation index. Analogously, an inflation-indexed floor is a put option on the same rate. The names cap/floor come from the idea of providing an upper/lower limit on in this case the inflation index. The payoff at time T is given by

$$[\omega (I(T) - \kappa)]^+, \tag{2.5}$$

where κ is the strike price, $\omega = 1$ for a cap and $\omega = -1$ for a floor.

More generally, a year-on-year inflation-indexed cap/floor is defined by a series of so-called year-on-year caplets/floorlets which start at T_{i-1} ($0 \leq t \leq T_{i-1}$) and mature at time T_i ($T_{i-1} \leq T_i$). Each caplet/floorlet has the payoff

$$\left[\omega \left(\frac{I(T_i)}{I(T_{i-1})} - K \right) \right]^+. \tag{2.6}$$

where $K = (1 + \kappa)^{T_i - T_{i-1}}$ with κ the quoted strike price, $T_0 = 0$, and $T_n = T$ is the end date. Here, n denotes the number of caplets/floorlets in the cap/floor and is dependent on the time interval which is typically fixed.

Currently this is the extent of the options environment for inflation-indexed derivatives. An open question remains whether exotic derivatives, such as barrier and look-back options, will be introduced in the future.

2.2 Pricing Formulas for Inflation-Indexed Derivatives

Standard no-arbitrage pricing theory implies that at time t , $0 \leq t < T$, the price of the ZCIIS under the nominal risk-neutral measure \mathbb{Q}_n is given by

$$\mathbf{ZCIIS}(t, T, I(0)) = \mathbb{E}^{\mathbb{Q}_n} \left[e^{-\int_t^T r_n(u) du} \left(\frac{I(T)}{I(0)} - 1 \right) \middle| \mathcal{F}_t \right], \quad (2.7)$$

where \mathcal{F}_t denotes the σ -algebra of market information generated by the underlying processes up to time t . The foreign-currency analogy implies that the nominal price of a real zero-coupon bond is equal to the nominal price of the contract paying off one unit of the CPI index at maturity. This means we also have

$$I(t)P_r(t, T) = I(t)\mathbb{E}^{\mathbb{Q}_r} \left[e^{-\int_t^T r_r(u) du} \middle| \mathcal{F}_t \right] = \mathbb{E}^{\mathbb{Q}_n} \left[e^{-\int_t^T r_n(u) du} I(T) \middle| \mathcal{F}_t \right].$$

Therefore, equation (2.7) can be rewritten to

$$\mathbf{ZCIIS}(t, T, I(0)) = \frac{I(t)}{I(0)} (P_r(t, T) - P_n(t, T)), \quad (2.8)$$

which at time $t = 0$ is the current value of the real and nominal zero-coupon bonds

$$\mathbf{ZCIIS}(0, T, I(0)) = (P_r(0, T) - P_n(0, T)), \quad (2.9)$$

Formulas (2.8) and (2.9) are model-independent prices of the zero-coupon swap which are not based on specific assumptions on the interest rate market. This important result makes it possible to strip real zero-coupon bond prices from the quoted prices of zero-coupon inflation-indexed swaps. Let $K(T_M)$ be the fixed rate of the contract for a given maturity T_M . Then the nominal discounted value of the fixed payment (2.4) must be equal to the formula in equation (2.8) which gives the price of a real zero-coupon bond at time $t = 0$ with maturity T_M :

$$P_r(0, T_M) = P_n(0, T_M)(1 + K(T_M))^M.$$

It is apparent that the price of a ZCIIS is not dependent on the choice of model. The same is not true for the YYIIS. Similarly computing the arbitrage-free price of a YYIIS under the nominal risk-neutral measure with $t < T_{i-1}$ yields

$$\mathbf{YYIIS}(t, T_{i-1}, T_i) = \mathbb{E}^{\mathbb{Q}_n} \left(e^{-\int_t^{T_i} r_n(u) du} \left[\frac{I(T_i)}{I(T_{i-1})} - 1 \right] \middle| \mathcal{F}_t \right),$$

The part in the square brackets can be rewritten to

$$\mathbb{E}^{\mathbb{Q}_n} \left(e^{-\int_t^{T_{i-1}} r_n(u) du} \mathbb{E}^{\mathbb{Q}_n} \left[e^{-\int_{T_{i-1}}^{T_i} r_n(u) du} \left(\frac{I(T_i)}{I(T_{i-1})} - 1 \right) \middle| \mathcal{F}_{T_{i-1}} \right] \middle| \mathcal{F}_t \right),$$

and recognizing that the inner expectation is $\mathbf{ZCIIS}(T_{i-1}, T_i, I(T_{i-1}))$, this results in

$$\mathbf{YYIIS}(t, T_{i-1}, T_i) = \mathbb{E}^{\mathbb{Q}_n} \left(e^{-\int_t^{T_{i-1}} r_n(u) du} P_r(T_{i-1}, T_i) \middle| \mathcal{F}_t \right) - P_n(t, T_i).$$

This expectation can be viewed as the nominal price of a derivative that pays off, in nominal units, the real zero-coupon bond price $P_r(T_{i-1}, T_i)$ at time T_{i-1} . If real rates are assumed to be stochastic then the expectation is model-dependent. In Section 3.2 the price of a year-on-year inflation swap is calculated based on the Jarrow-Yildirim model.

For inflation-indexed caplets/floorlets the price is derived again with standard no-arbitrage pricing theory using the payoff (2.6). The value at time $t \leq T_{i-1}$ is:

$$\mathbf{IICplt}(t, T_{i-1}, T_i, \omega) = \mathbb{E}^{\mathbb{Q}_n} \left(e^{-\int_t^{T_i} r_n(u) du} \left[\omega \left(\frac{I(T_i)}{I(T_{i-1})} - K \right) \right]^+ \middle| \mathcal{F}_t \right). \quad (2.10)$$

If the nominal interest rate and inflation index are assumed to be stochastic, then the expressions in the expectation are not independent, and solving this expectation can become quite cumbersome. In many model structures, the pricing formula is calculated under the nominal T -forward measure. This measure is generated by the nominal zero-coupon bond $P_n(t, T)$ so that the "forward" inflation is a martingale, i.e.

$$\mathbb{E} \left[I(T) \frac{P_r(t, T)}{P_n(t, T)} \middle| \mathcal{F}_t \right] = I(t) \frac{P_r(t, T)}{P_n(t, T)}. \quad (2.11)$$

This means under this measure the discount factor can be decoupled from the expectation,

$$\mathbf{IICplt}(t, T_{i-1}, T_i, \omega) = P_n(t, T_i) \mathbb{E}^{\mathbb{Q}_n^{T_i}} \left(\left[\omega \left(\frac{I(T_i)}{I(T_{i-1})} - K \right) \right]^+ \middle| \mathcal{F}_t \right).$$

For models under the T -forward measure containing stochastic interest rates, this greatly simplifies the pricing structure. Pricing formulas are calculated for different model frameworks in the next chapter.

3 Model Overview

This section provides an overview of the models and methodologies used for pricing inflation-indexed derivatives. First, the Black-Scholes model finds its place in inflation modelling as a simple base model. Subsequently, combinations of Hull-White processes and Heston dynamics are illustrated in the form of the famous Jarrow-Yildirim model, the Heston-Hull-White model, and LIBOR market models including the evolved SABR model.

3.1 Black-Scholes

In the pricing of financial derivatives, the Black-Scholes [4] framework is considered one of the most natural starting points. The simplicity of this model makes it widely accessible, and it results in a closed-form solution for the price of an option. It is commonly used as the benchmark for pricing options despite its limitations in capturing stochastic dynamics. Applying this framework to inflation by considering the inflation index as an underlying asset is therefore a natural first step.

Kruse [31], [32] considers the simple inflation model where the inflation index $I(t)$ follows a geometric Brownian motion with (Ω, \mathcal{F}, P) the real-world probability space with filtration $(\mathcal{F}_t)_{t \geq 0}$ of market information.

$$dI(t) = (r_n - r_r)I(t)dt + \sigma_I I(t)dW_I(t),$$

where $I(t)$ is the inflation index, r_n the constant interest rate in the nominal economy, r_r the constant interest rate in the real economy, σ_I the constant volatility, and $dW_I(t)$ a Brownian motion. By calculating the expectation, it can be shown that this model preserves the Fisher equation, so it satisfies the expected relation between inflation and interest rates, for $T > t$,

$$\mathbb{E}_Q \left(\frac{I(T)}{I(t)} \middle| \mathcal{F}_t \right) = e^{(r_n - r_r)(T-t)}.$$

Analogous to the Black-Scholes solution, the price of a plain vanilla call option on the inflation index with strike price K , maturity T and payoff $C(T, I(T)) = (I(T) - K)^+$ at maturity, is at time $t = 0$ given by

$$C(0, I(0)) = I(0)e^{-r_r T} N(d) - Ke^{-r_n T} N(d - \sigma_I \sqrt{T}),$$

with

$$d = \frac{\ln \left(\frac{I(0)}{K} \right) + (r_n - r_r + \frac{1}{2} \sigma_I^2) T}{\sigma_I \sqrt{T}}.$$

Similarly the price of a caplet on the inflation rate over a future time interval from T_{i-1} to T_i with payoff (2.6) is also derived

$$\mathbf{HCplt}(0, T_{i-1}, T_i, K) = e^{-r_r(T_i - T_{i-1})} e^{-r_n T_{i-1}} N(d) - Ke^{-r_n T_i} N(d - \sigma_I \sqrt{T_i - T_{i-1}}),$$

with

$$d = \frac{-\ln(K-1) + (r_n - r_r + \frac{1}{2}\sigma_I^2)(T_i - T_{i-1})}{\sigma_I\sqrt{T_i - T_{i-1}}}.$$

With this pricing formula, the implied volatility surface is derived from data on inflation caps. It is shown that the implied volatilities are far from constant, indicating the deficiencies in the model.

While this model is easy to use, the drawbacks that exist in the Black-Scholes model for options on assets carry over to derivatives on the inflation index. Most notably the assumption of constant volatility and interest rates over time is unrealistic. Just like in the stock or currency option market, the implied Black-Scholes volatilities are not constant for different strike levels and maturities. Likewise, interest rates are observed to have stochastic behavior.

3.2 Jarrow-Yildirim framework

In 1997 the US treasury started issuing inflation-indexed bonds, differing from conventional bonds in that the principal is constantly adjusted for inflation. With the first inflation-indexed derivatives on the market also came the first pricing models for such products. The model published by Jarrow and Yildirim [30] in 2003 is one of the first applications of the foreign-currency analogy to inflation-indexed derivatives, as it is the first model to jointly model interest rates and inflation. Analogously to the well-known Heath-Jarrow-Morton (HJM) framework [16], the forward rates under the nominal and real economies are assumed to follow Hull-White processes. See also Dodgson et al. [18] and Mercurio [35]. Simultaneously the inflation index is assumed to follow a log-normal process correlated to the forward rates,

$$\begin{aligned} df_n(t, T) &= \alpha_n(t, T)dt + \sigma_n(t, T)dW^n(t), \\ df_r(t, T) &= \alpha_r(t, T)dt + \sigma_r(t, T)dW^r(t), \\ \frac{dI(t)}{I(t)} &= \mu_I(t)dt + \sigma_I dW^I(t), \end{aligned}$$

where the Brownian motions $dW^n(t)$, $dW^r(t)$ and $dW^I(t)$ are correlated by $\rho_{i,j}$, $i \in \{n, r, I\}$ and $\sigma_n, \sigma_r, \sigma_I$ are deterministic functions. According to Jarrow and Turnbull [29] the processes are arbitrage free if and only if the processes $\frac{P_n(t, T)}{B_n(t)}$, $\frac{I(t)P_r(t, T)}{B_n(t)}$ and $\frac{I(t)B_r(t)}{B_n(t)}$ are martingales under the nominal risk-neutral measure, with $B_n(t)$ and $B_r(t)$ money market accounts that increase with the nominal and real interest rate, respectively. This implies the following conditions

$$\alpha_n(t, T) = \sigma_n(t, T) \int_t^T \sigma_n(t, s) ds \tag{3.1}$$

$$\alpha_r(t, T) = \sigma_r(t, T) \left(\int_t^T \sigma_r(t, s) ds - \sigma_I(t) \rho_{r, I} \right) \tag{3.2}$$

$$\mu_I(t) = r_n(t) - r_r(t) \tag{3.3}$$

where equation (3.1) refers to the arbitrage-free forward rate drift restriction as in the original HJM model. Equation (3.2) is the analogous restriction for the

real forward rates, and equation (3.3) corresponds to the Fisher equation relating expected inflation with the nominal and real interest rates.

From here Jarrow and Yildirim use general results for Hull-White interest rate processes to show that the nominal and real zero-coupon bonds must have the price processes

$$\begin{aligned}\frac{dP_n(t, T)}{P_n(t, T)} &= r_n(t)dt - \int_t^T \sigma_n(t, s) ds dW^n(t), \\ \frac{dP_r(t, T)}{P_r(t, T)} &= \left(r_r(t) + \rho_{r,I}\sigma_I(t) \int_t^T \sigma_r(t, s) ds \right) dt - \int_t^T \sigma_r(s) ds dW^r(t).\end{aligned}$$

Additionally the dynamics of the nominal and real interest rates are derived under the nominal risk neutral measure using the change-of-measure technique,

$$dr_n(t) = (\theta_n(t) - a_n r_n(t))dt + \sigma_n dW^n(t), \quad (3.4)$$

$$dr_r(t) = (\theta_r(t) - \rho_{r,I}\sigma_I\sigma_r - a_r r_r(t))dt + \sigma_r dW^r(t), \quad (3.5)$$

$$\frac{dI(t)}{I(t)} = (r_n(t) - r_r(t))dt + \sigma_I(t)dW^I(t), \quad (3.6)$$

where $\theta_n(t)$ and $\theta_r(t)$ are deterministic functions used to exactly fit the term structures of nominal and real rates respectively. For the Hull-White model this results in the following formula,

$$\theta_l(t) = \frac{\partial f^l(0, t)}{\partial T} + a_l f^l(0, t) + \frac{\sigma_l}{2a_l}(1 - e^{-2a_l t}), \quad (3.7)$$

with $l \in \{n, r\}$ and $f_l(t, T)$, $t \leq T$ the forward rate at time t for maturity M .

This provides all the ingredients required to price inflation-indexed bonds and derivatives with closed-form formulas. Brigo and Mercurio [7] use the Jarrow and Yildirim model to price Year-on-Year swaps and caplets/floorlets For the Year-on-Year swap, the price under the nominal T-forward measure can be written as

$$\mathbf{YYIS}(t, T_{i-1}, T_i) = P_n(t, T_{i-1})\mathbb{E}_n^{T_{i-1}}[P_r(T_{i-1}, T_i)|\mathcal{F}_t] - P_n(t, T_i)$$

where $\mathbb{E}_n^{T_{i-1}}$ is the expectation under the T_{i-1} -forward measure. From the theory behind the Hull-White model the price of the real zero-coupon bond is

$$\begin{aligned}P_r(t, T) &= A_r(t, T)e^{-B_r(t, T)r_r(t)}, \\ B_r(t, T) &= \frac{1}{a_r}(1 - e^{-a_r(T-t)}), \\ A_r(t, T) &= \frac{P_r(0, T)}{P_r(0, t)} \exp \left[B_r(t, T)f_r(0, t) - \frac{\sigma_r^2}{4a_r}(1 - e^{2a_r t})B_r(t, T)^2 \right].\end{aligned}$$

With a change of measure the real interest rate under the T_{i-1} -forward measure follows the dynamics

$$dr_r(t) = [-\rho_{n,r}\sigma_n\sigma_r B_n(t, T_{i-1}) + \theta_r(t) - \rho_{r,I}\sigma_I\sigma_r - a_r r_r(t)]dt + \sigma_r dW^r(t)$$

with $dW^r(t)$ a $\mathbb{Q}^{T_{i-1}}$ -Brownian motion. After some calculations the price for the Year-on-Year inflation indexed swap is:

$$\mathbf{YYIIS}(t, T_{i-1}, T_i) = P_n(t, T_{i-1}) \frac{P_r(t, T_i)}{P_r(t, T_{i-1})} e^{C(t, T_{i-1}, T_i)} - P_n(t, T_i),$$

where

$$C(t, T_{i-1}, T_i) = \sigma_r B_r(T_{i-1}, T_i) \left[B_r(t, T_{i-1}) \left(\rho_{r,I} \sigma_I - \frac{1}{2} \sigma_r B_r(t, T_{i-1}) \right) + \frac{\rho_{n,r} \sigma_n}{a_n + a_r} (1 + a_r B_n(t, T_{i-1})) \right] - \frac{\rho_{n,r} \sigma_n}{a_n + a_r} B_n(t, T_{i-1})$$

Thus under the JY model the expectation of a real zero-coupon bond under the nominal forward measure is equal to the current forward price of the real bond multiplied by a correction factor C , which depends on the volatilities of the nominal rate, real rate and the CPI, and the correlation between the real rate and the CPI. The derivation of the price of a caplet is also long but straightforward, using the idea that the ratio $I(T_i)/I(T_{i-1})$ conditional on \mathcal{F}_t is lognormal under the forward measure.

3.3 Heston framework

As a result of the significant skew/smile observed in the option market, the Heston framework is often preferred in practice. In this case the variance process of the inflation index is modeled by a Cox-Ingersoll-Ross (CIR) process which is considered the gold standard in capturing the smile. For even more realistic dynamics, this can be coupled with a stochastic term structure model for the interest rates such as the Hull-White model [7]. Choosing this structure makes it possible to derive explicit formulas for the prices of options on inflation. For more complicated dynamics some assumptions are required to obtain such formulas.

3.3.1 Constant interest rate

The simple case where the interest rates are assumed to be constant almost exactly overlaps the scenario where Heston's model is applied to an asset. An asset price $A(t)$ which follows the Heston dynamics has the process

$$\frac{dA(t)}{A(t)} = r dt + \sqrt{V(t)} dW^A(t)$$

where r is the constant risk free nominal interest rate r and $V(t)$ follows a mean reverting stochastic process. Kruse [32] assumed the inflation index $I(t)$ follows the same type of process as $A(t)$:

$$\begin{aligned} \frac{dI(t)}{I(t)} &= (r_n - r_r) I(t) dt + \sqrt{V(t)} dW^I(t) \\ dV(t) &= \alpha(\bar{V} - V(t)) dt + \sigma_V \sqrt{V(t)} dW^V(t) \end{aligned} \quad (3.8)$$

with risk neutral long term mean \bar{V} , mean reverting speed $\alpha \geq 0$, σ_V the constant volatility of the variance process, and $W^I(t)$ and $W^V(t)$ are correlated through parameter ρ . Additionally the feller condition which ensures positivity of the variance process is satisfied if $2\alpha\bar{V} > \sigma_V^2$.

The price of a cap on the inflation index at time $t = 0$ with payoff $(I(T) - K)^+$ for some strike K is now the same as the price of a call option with strike K on the asset $A(t)$. In the original paper about his model, Heston[27] constructs an exact solution by deriving the characteristic function,

$$\begin{aligned} \phi_{hes}(u, t) = \exp \left\{ iut(r_n - r_r) + \frac{\alpha\bar{V}}{\sigma_V^2} ((\alpha - i\sigma_V\rho u - d)t - 2 \log \left(\frac{1 - ge^{-dt}}{1 - g} \right)) \right. \\ \left. + \frac{V_0}{\sigma_V} (\alpha - i\sigma_V\rho u - d) \left(\frac{1 - e^{-dt}}{1 - ge^{-dt}} \right) \right\}, \end{aligned} \quad (3.9)$$

where

$$d = \sqrt{(\alpha - i\sigma_V\rho u)^2 + \sigma_V^2(u^2 + iu)}, \quad (3.10)$$

and

$$g = \frac{\alpha - i\sigma_V\rho u - d}{\alpha - i\sigma_V\rho u + d}. \quad (3.11)$$

Heston then inverts the characteristic function and numerically integrates to obtain the exact solution of call options. A more recent and more efficient method that has become popular is the COS method [19], which approximates the integral using Fourier expansions. This method is illustrated in Section 3.4.

For a caplet payoff as in equation (2.6) the result for forward starting options can be followed (see, for example, Kruse and Nögel [33]). The price at time $t = 0$ of a caplet on the inflation index from T_{i-1} to T_i with $0 < T_{i-1} < T_i$ is given by

$$\begin{aligned} \mathbf{HCpl}(0, T_{i-1}, T_i, K) &= e^{-r_r(T_i - T_{i-1})} e^{-r_n T_{i-1}} P_1(0, e^{-r_r(T_i - T_{i-1})}, V(0), 1 + K) \\ &\quad - (1 + K) e^{-r_n T_i} P_2(0, e^{-r_r(T_i - T_{i-1})}, V(0), 1 + K) \end{aligned}$$

where P_1, P_2 are the Heston probabilities given in [32]. The full formula is not presented here, but it involves the standard method of characteristic functions that is widely used for Heston dynamics. As with option pricing for assets, the characteristic function has an exponential affine form for which the terms can be derived by solving ordinary differential equations.

While pricing using this model is relatively straightforward, as it does not rely on specific assumptions on the correlation between inflation and its variance, the assumption of constant interest rates is generally not realistic. The fisher equation has long been accepted as a defined relation between the nominal/real interest rates and the inflation rate. As a result, stochastic interest rates are expected in an accurate stochastic inflation pricing model. Therefore, the natural next step is introducing stochastic models for the interest rate which work in conjunction with the model for inflation.

3.3.2 Hull-White model for interest rates

In the interest of giving interest rates more structure than a constant or deterministic function, much attention has been devoted to combining the Heston dynamics of an asset with a model for the interest rates. Hout et al. [28] and Chval [13] among others extend the Black-Scholes model for an asset with Heston dynamics for the variance and a Hull-White process for the interest rate to price European options. The reason for this popularity is many features in the option market can be captured by this combination of models while analytic pricing formulae exist. This is highly preferred, as it makes the valuation procedure much faster since numerical/simulation techniques can be avoided.

Naturally, this makes the inflation option market another interesting application for this structure. Singor et al. [43][25] have a series of papers that investigate this model structure, including an application to pension funds. The nominal and real interest rates, r_n and r_r follow the one-factor Hull White model under the nominal risk-neutral measure \mathbb{Q}_n just as in equations (3.4) - (3.6):

$$dr_n(t) = (\theta_n(t) - a_n r_n(t))dt + \sigma_n dW^{r_n}(t) \quad (3.12)$$

$$dr_r(t) = (\theta_r(t) - \rho_{I,r} \sigma_r \sqrt{V(t)} - a_r r_r(t))dt + \sigma_r dW^{r_r}(t) \quad (3.13)$$

where a_i are mean-reversion parameters and σ_i are volatility parameters for $i \in \{n, r\}$. The functions $\theta_i(t)$ are determined by the nominal/real initial term structure observed in the market. Meanwhile the CPI $I(t)$ follows the Heston model as in equations (3.8) under the nominal economy measure \mathbb{Q}_n but now with time-dependent interest rates,

$$\begin{aligned} \frac{dI(t)}{I(t)} &= (r_n(t) - r_r(t))I(t)dt + \sqrt{V(t)}dW^I(t), \\ dV(t) &= \alpha(\bar{V} - V(t))dt + \sigma_V \sqrt{V(t)}dW^V(t). \end{aligned}$$

The correlation structure between the Brownian motions $dW^I(t)$, $dW^V(t)$, $dW^{r_n}(t)$, $dW^{r_r}(t)$ is provided through the symmetric positive-definite correlation matrix

$$d\mathbf{W}(t)(d\mathbf{W}(t))^T = \begin{bmatrix} 1 & \rho_{I,V} & \rho_{I,n} & \rho_{I,r} \\ \rho_{I,V} & 1 & \rho_{V,n} & \rho_{V,r} \\ \rho_{I,n} & \rho_{V,n} & 1 & \rho_{n,r} \\ \rho_{I,r} & \rho_{V,r} & \rho_{n,r} & 1 \end{bmatrix} dt. \quad (3.14)$$

From here the dynamics are given under the T -forward measure so that the forward CPI is a martingale. This greatly simplifies the derivation of the characteristic function later. The full derivation is presented by Grzelak and Oosterlee [25]

$$\begin{aligned} dr_n(t) &= (\theta_n(t) - \sigma_n^2 B_n(t, T) - a_n r_n(t))dt + \sigma_n dW^{r_n}(t), \\ dr_r(t) &= \left(\theta_r(t) - \rho_{I,r} \sigma_r \sqrt{V(t)} - \sigma_n \sigma_r \rho_{n,r} B_n(t, T) - a_r r_r(t) \right) dt + \sigma_r dW^{r_r}(t), \\ \frac{dI^T(t)}{I^T(t)} &= \sqrt{V(t)}dW^I(t) + \sigma_n B_n(t, T)dW^{r_n}(t) - \sigma_r B_r(t, T)dW^{r_r}(t), \\ dV(t) &= \left(\alpha(\bar{V} - V(t)) - \sigma_V \sigma_n \rho_{V,n} B_n(t, T) \sqrt{V(t)} \right) dt + \sigma_V \sqrt{V(t)}dW^V(t). \end{aligned}$$

where $B_i(t, T) = \frac{1}{a_i}(1 - e^{-a_i(T-t)})$ for $i \in \{n, r\}$. It is worth noting that the forward CPI does not depend directly on the interest rate processes but only depends on the Brownian motions.

Just as in the previous section, the characteristic function belonging to the inflation cap/floor under this model must be derived to apply the Fourier-based pricing methods. The derivation of this characteristic function is discussed in Grzelak and Oosterlee [25] in the context of foreign exchange markets. By applying the Feynman-Kac theorem [42] to the log inflation process $x = \log I(t)$ they show the characteristic function is the solution of the following Kolmogorov backward partial differential equation:

$$\begin{aligned} -\frac{\partial \phi}{\partial t} = & (\alpha(\bar{V} - V) + \rho_{I,n}\sigma_V\sigma_n\sqrt{V}B_n\frac{\partial \phi^T}{\partial V} + \left(\frac{1}{2}V - \zeta(t, \sqrt{V})\right)\left(\frac{\partial^2 \phi^T}{\partial x^2} - \frac{\partial \phi^T}{\partial x^2}\right) \\ & + (\rho_{I,V}\sigma_VV - \rho_{V,n}\sigma_n\sqrt{V}B_n + \rho_{V,r}\sigma_V\sigma_r\sqrt{V}B_r)\frac{\partial^2 \phi^T}{\partial x \partial V} + \frac{1}{2}\sigma_V^2V\frac{\partial^2 \phi^T}{\partial V^2} \end{aligned} \quad (3.15)$$

with $B_n = B_n(t, T)$, $B_r = B_r(t, T)$ and

$$\zeta(t, \psi(t)) = (-\rho_{I,n}\sigma_nB_n + \rho_{I,r}\sigma_rB_r)\sqrt{V(t)} + \rho_{r,n}\sigma_n\sigma_rB_rB_n - \frac{1}{2}(\sigma_n^2B_n^2 + \sigma_r^2B_r^2). \quad (3.16)$$

By approximating the non-affine \sqrt{V} terms with a linearization technique, and letting $\tau := T - t$, an approximate closed-form solution of the characteristic function is provided,

$$\phi^T(u, \tau) = \exp(A(u, \tau) + B(u, \tau)x^T(t) + C(u, \tau)\sigma(t)). \quad (3.17)$$

with

$$B(u, \tau) = iu, \quad (3.18)$$

C recognized from the Heston model,

$$C(u, \tau) = \frac{1 - e^{-d\tau}}{\sigma_V^2(1 - ge^{-d\tau})}(\alpha - \rho_{I,V}\sigma_Viu - d) \quad (3.19)$$

with

$$d = \sqrt{(\rho_{I,V}\sigma_Viu - \alpha)^2 - \sigma_V^2iu(iu - 1)} \quad (3.20)$$

$$g = \frac{\alpha - \sigma_V\rho_{I,V}iu - d}{\alpha - \sigma_V\rho_{I,V}iu + d}, \quad (3.21)$$

and

$$\begin{aligned} A(u, \tau) = & \int_0^T (\alpha\bar{V} - \rho_{V,n}\sigma_V\sigma_n\psi(s)B_n(s) + \rho_{\sigma,n}\sigma_n\sigma_V\psi(s)B_n(s)iu \\ & - \rho_{V,r}\sigma_V\psi(s)B_r(s)iu)C(s)ds + (u^2 + iu) \int_0^T \zeta(s, \psi(s))ds, \end{aligned} \quad (3.22)$$

where

$$\psi(t) := \mathbb{E}(\sqrt{V(t)}) \quad (3.23)$$

is approximated. Under certain assumptions on the correlation parameters, an analytic solution for $A(u, \tau)$ exists. Otherwise, numerical procedures can be used. Grzelak and Oosterlee [43] further derive an approximate expression for the characteristic function of year-on-year inflation caps/floors which involves $\phi^T(u, \tau)$. This topic, however, falls outside the scope of this thesis.

After a complicated sequence of calculations an inflation cap/floor can now be priced exactly for the Heston Hull-White model using Fourier methods.

3.3.3 Libor Market Models for forward rates

LIBOR (London Interbank Offer Rate) market models, also known as Brace-Gatarek-Musiela (BGM) models, have long been an industry standard in pricing exotic interest rate derivatives by modeling forward rates over short rates. Unlike when modeling short rates like in the Hull-White framework, forward rates are directly observable in the market. Using forward rates eliminates the requirement to calibrate the short-rate process parameters to the prices of market instruments, which would be an extra layer of fitting [2]. Because the discount curve can evolve in a complicated manner, using the forward rate means more focus can be on the dynamics of the asset rather than the fitting of the short rate curve. This motivated Mercurio and Moreni [36] to use the LIBOR market model in pricing Inflation-Indexed derivatives.

The forward rates F_i (see equation (2.3)) are assumed to be lognormally distributed according to a driftless lognormal LIBOR model, see, for example, Brace et al. [6]. Under a given reference measure \mathbb{Q} with tenor dates $0 < T_0 < T_1 < \dots < T_n$ and corresponding forward rates f_0, \dots, f_{n-1} with f_i the forward rate for the period T_i to T_{i+1} , this means

$$\frac{dF_i(t)}{F_i(t)} = \sigma_i^F dW_i^{\mathbb{Q}, F}(t) \quad (3.24)$$

Due to the dependence of forward CPI's on forward rates, Mercurio and Moreni combined this model with a Heston-like model for the forward CPI's \mathcal{I}_i each with the same mean-reverting square-root volatility process. Compared to the CPI, which is simply an exchange rate, forward CPI's are price processes of tradable securities. Furthermore, a forward CPI is a martingale under its corresponding forward measure, so its dynamics are fully defined by its volatility process. Having a single volatility process for each forward CPI is a reasonable assumption since the dynamics for each forward CPI are expected to be similar.

In Section 3.3.4 an expansion of the model is explored where each forward CPI is modeled by a different volatility process at the cost of the requirement to calibrate

many more parameters. The complete model is as follows:

$$\begin{aligned}\frac{dF_i(t)}{F_i(t)} &= \sigma_i^F dW_i^{\mathbb{Q},F}(t) \\ \frac{d\mathcal{I}_i(t)}{\mathcal{I}_i(t)} &= \sigma_i^I \sqrt{V(t)} dW_i^{\mathbb{Q},I}(t), \\ dV(t) &= \alpha(\theta - V(t))dt + \sigma_V \sqrt{V(t)} dW^{\mathbb{Q}}(t),\end{aligned}$$

where $\sigma_i^F, \sigma_i^I, \alpha, \theta, \sigma_V$ are positive constants, and the Brownian motions $dW_i^{\mathbb{Q},F}(t)$, $W_i^{\mathbb{Q},I}(t)$ and $dW^{\mathbb{Q}}(t)$ are correlated. To ensure the variance process is positive, the Feller condition $2\alpha\theta > \sigma_V$ must be satisfied.

Instead of defining the dynamics under a risk-neutral measure, which would hide the market price of volatility risk, or under the terminal forward measure which would use the bond price $P(t, T_M)$ as numeraire and thus depend on the choice of the last maturity, the authors chose the spot LIBOR measure \mathbb{Q}_0 . This measure uses the numeraire

$$N(t) := P(t, \beta(t)) \prod_{l=1}^{\beta(t)} (1 + \tau F_l(t)), \quad \beta(t) = T_j \text{ if } T_{j-1} < t \leq T_j$$

This numeraire involves an initial portfolio of a zero-coupon bond expiring at time T_0 , with the proceeds received upon expiration of every bond being reinvested in bonds expiring at the next tenor date, up to T_n . The use of this numeraire was introduced in 1997 by Jamshidian [20].

The authors' motivation behind using the spot LIBOR measure for this forward CPI model is that it is payoff independent, so it allows the valuation of a single payoff of a caplet 2.6 without needing to consider the number of caplets that follow. However, in the context of interest rate modeling, this numeraire is used because the variance of Monte Carlo simulation is generally lower than with the terminal forward measure [2]. This can be of particular importance when exact solutions cannot be found.

Following a change-of-measure technique as presented by Geman et al. [22], Mercurio and Moreni proceed to define the dynamics under both the spot LIBOR measure and the forward measure. Since caplets and floorlets with payoff at time T_j are derivatives that depend on \mathcal{I}_j and \mathcal{I}_{j-1} it is convenient to use the forward measure, as it means both F_j and \mathcal{I}_j are martingales. The dynamics under the forward

measure with $P(t, T_j)$ as numeraire and $X_j(\cdot) := \ln\left(\frac{\mathcal{I}_j(\cdot)}{\mathcal{I}_{j-1}(\cdot)}\right)$ become

$$\begin{aligned}
\frac{dF_i(t)}{F_i(t)} &= \sigma_i^F dW_i^F(t) \\
\frac{d\mathcal{I}_j(t)}{\mathcal{I}_j(t)} &= \sqrt{V(t)}\sigma_j^I dW_j^I(t), \\
\frac{d\mathcal{I}_{j-1}(t)}{\mathcal{I}_{j-1}(t)} &= \sqrt{V(t)}\sigma_{j-1}^I \left[-\sigma_j^F \rho_{j,j-1}^{F,I} \frac{\tau_j F_j(t)}{1 + \tau_j F_j(t)} dt + dW_{j-1}^I(t) \right], \\
dX_j(t) &= \left[\frac{V(t)}{2}((\sigma_{j-1}^I)^2 - (\sigma_j^I)^2) + \sqrt{V(t)}\sigma_{j-1}^I \sigma_j^F \rho_{j,j-1}^{F,I} \frac{\tau_j F_j(t)}{1 + \tau_j F_j(t)} \right] dt, \\
&\quad + \sqrt{V(t)}(\sigma_j^I dW_j^I(t) - \sigma_{j-1}^I dW_{j-1}^I(t)) \\
dV(t) &= \left[\alpha\theta - \sigma_V \sqrt{V(t)} \sum_{l=\beta(t)+1}^j \frac{\tau_l F_l(t)}{1 + \tau_l F_l(t)} \sigma_l^F \rho_l^{F,V} - \alpha V(t) \right] dt + \sigma_V \sqrt{V(t)} dW(t),
\end{aligned} \tag{3.25}$$

where $\rho_l^{F,V} dt = dW_l^F(t)dW(t)$ for every l , $\rho_{j,l}^I = dW_j^I(t)dW_l^I(t)$ is the correlation between forward CPI's \mathcal{I}_j and \mathcal{I}_l , and $\rho_i^{I,V} = dW_i^I(t)dW(t)$ the correlation between the forward CPI \mathcal{I}_i and the volatility.

Despite the focus on the spot LIBOR measure, which results in more simple dynamics than the forward measure, the pricing formula is only calculated with the forward measure in mind. The method of characteristic functions is used to obtain the caplet price in terms of its Fourier transform

$$\begin{aligned}
\text{Cplt}_j(t, K) &= P(t, T_j) \mathbb{E}_t^j \left(\frac{\mathcal{I}_j(T_j)}{\mathcal{I}_{j-1}(T_{j-1})} - (K + 1) \right) \\
&= P(t, T_j) \frac{e^{-\eta k}}{\pi} \int_0^\infty e^{-iuk} \frac{\phi_t^j(u - (\eta + 1)i)}{(\eta + iu)(\eta + 1 + iu)} du,
\end{aligned}$$

where $k = \ln(K + 1)$ and $\eta \in \mathbb{R}^+$, is used to ensure L^2 -integrability. Further reading on this method can be found in [11]. The only unknown value is the conditional characteristic function

$$\phi_t^j(u) = \mathbb{E}_t^j \left(e^{iu \ln\left(\frac{\mathcal{I}_j(T_j)}{\mathcal{I}_{j-1}(T_{j-1})}\right)} \right)$$

which can be written as the solution to a partial differential equation that can be found using Feynmann-Kac's theorem (by definition of the characteristic function and the Markov property).

Due to the general dynamics in equations (3.25) it may not be explicitly solvable. Instead, results are derived separately for two specific assumptions. First, by assuming $\rho_{i,l}^{F,I} = \rho_i^{F,V} = 0$, many terms in the dynamics vanish, so finding $\phi_t^j(u)$ reduces to solving a Heston-like partial differential equation for which an explicit solution exists. This assumption is reasonable for basic derivatives which are not sensitive to the correlation between forward rates and forward CPI's. When pricing more exotic

derivatives on interest rate and inflation, approximated dynamics can be derived instead by freezing the drift terms to their value at time 0. This also provides a closed-form solution by solving a Heston-like partial differential equation.

3.3.4 Multi-Factor SABR model

By choosing a single volatility process for every forward CPI as in the previous section, the model is sensitive to calibration errors, especially when considering long-term maturities. Mercurio and Moreni [37] therefore expanded their model by defining a separate volatility process for each forward CPI. The resulting caplet price can be valued with a SABR formula which has been researched by Hagan et al. [26] for the pricing of European options.

In the expanded model, the forward rates follow the same process as in Equation (3.24). Next the forward CPI's \mathcal{I}_i evolve under the associated forward measure \mathcal{Q}^{T_i} according to

$$\frac{d\mathcal{I}_i(t)}{\mathcal{I}_i(t)} = \sum_{j=\beta(t)}^i V_j(t) dW_j^i(t)$$

where $\beta(t)$ is the index of the first tenor date T_i strictly larger than t . This can be interpreted as the volatility of the forward CPI to T_i is a sum of the volatility processes for each tenor date up to T_i .

The M volatility processes are modeled by driftless geometric Brownian motions under their respective forward measure,

$$\frac{dV_i(t)}{V_i(t)} = \nu_i dW_i(t)$$

In order to provide structure to the Brownian motion structure, $M-1$ M -dimensional Brownian motions are recursively defined

$$\begin{aligned} W^1 &:= \{W_1^1, W_2^1, \dots, W_M^1\} \\ W^2 &:= \{W_1^2, W_2^2, \dots, W_M^2\} \\ &\vdots \\ W^{M-1} &:= \{W_1^{M-1}, W_2^{M-1}, \dots, W_M^{M-1}\} \end{aligned}$$

by the rule

$$dW_j^{i-1}(t) = dW_j^i(t) - \frac{\tau_i \sigma_i^F(t) F_i(t)}{1 + \tau_i F_i(t)} \rho_{i,j}^{F,W} dt \quad (3.26)$$

where $j = 1, \dots, M$ and $i = 2, \dots, M$ and $\rho_{i,j} := dW_j^i(t) dW_i^F(t) / dt = dW_j^{i-1}(t) dW_i^F(t) / dt$. This recursive structure comes from the change-of-measure technique, when moving from \mathcal{Q}^{T_i} to $\mathcal{Q}^{T_{i-1}}$ dynamics of W_j^i contain a drift term equal to the final term stated in equation (3.26). By taking this drift term away, the recursively defined Brownian motions are indeed Brownian motions.

Similar to the previous section, the dynamics of $\mathcal{I}_{i-1}(t)$ are derived under the forward measure, this time with a sum over the volatility processes. Using equation (3.26),

$$\frac{d\mathcal{I}_{i-1}(t)}{\mathcal{I}_{i-1}(t)} = -\frac{\tau_i \sigma_i^F(t) F_i(t)}{1 + \tau_i F_i(t)} \sum_{j=\beta(t)}^{i-1} V_j(t) \rho_{i,j}^{F,W} dt + \sum_{j=\beta(t)}^{i-1} V_j(t) dW_j^i(t).$$

Next with the goal of producing SABR dynamics in mind, the dynamics of the forward inflation rate Y_i defined by $Y_i(t) := \frac{\mathcal{I}_i(t)}{\mathcal{I}_{i-1}(t)} - 1$ are derived using Itô's lemma,

$$dY_i(t) = (1 + Y_i(t)) \left[\sum_{j=\beta(t)}^{i-1} V_j(t) \left(\frac{\tau_i \sigma_i^F(t) F_i(t)}{1 + \tau_i F_i(t)} \rho_{i,j}^{F,W} - V_i(t) \rho_{i,j}^W \right) dt + V_i(t) dW_i^i(t) \right]$$

To produce analytically tractable dynamics the forward rates $F_i(t)$ and volatilities $V_j(t)$ in the drift term are frozen to their value at time 0. Furthermore, setting $\bar{Y}_i(t) := 1 + Y_i(t)$ produces

$$\begin{aligned} \frac{d\bar{Y}_i(t)}{\bar{Y}_i(t)} &= D_i(t) dt + V_i(t) dW_i^i(t) \\ dV_i(t) &= \nu V_i(t) dZ_i(t), V_i(0) = \alpha_i \end{aligned}$$

with

$$D_i(t) := \sum_{j=\beta(t)}^{i-1} V_j(0) \left(\frac{\tau_i \sigma_i^F(t) F_i(0)}{1 + \tau_i F_i(0)} \rho_{i,j}^{F,W} - V_i(0) \rho_{i,j}^W \right)$$

which is only dependent on the forward rate volatility's and the correlations.

Finally the SABR dynamics are stated by noticing that $\bar{Y}_i(T_i) = \tilde{Y}_i(T_i)$ where the process $\tilde{Y}_i(T_i)$ is defined by

$$\begin{aligned} d\tilde{Y}_i(T_i)(t) &= \tilde{Y}_i(T_i) V_i(t) dW_i^i(t), \tilde{Y}_i(0) = \bar{Y}_i(0) e^{\int_0^{T_i} D_i(t) dt} \\ dV_i(t) &= \nu V_i(t) dZ_i(t), V_i(0) = \alpha_i. \end{aligned}$$

The caplet price can now be written directly from the SABR lognormal formula with $\beta = 1$ with ρ_i the correlation between \tilde{Y}_i and V_i (see equations 2.16a-2.17c in [26]):

$$\begin{aligned} \mathbf{HCpl}(t, T_{i-1}, T_i, K) &= P(t, T_i) \mathbb{E}^{T_i} [(\bar{Y}_i(T_i) - K)^+ | \mathcal{F}_t] \\ &= P(t, T_i) \mathbb{E}^{T_i} [(\tilde{Y}_i(T_i) - K)^+ | \mathcal{F}_t] \\ &= P(t, T_i) [\tilde{Y}_i(t) \Phi(d_+) - K \Phi(d_-)] \end{aligned}$$

with

$$\begin{aligned} d_{\pm} &= \frac{\ln\left(\frac{\tilde{Y}_i(t)}{K}\right) \pm \frac{1}{2} \sigma^2(K) (T_i - t)}{\sigma(K) \sqrt{T_i - t}} \\ \sigma(K) &= \alpha_i \frac{z}{x(z)} \left[1 + \left(\frac{\rho_i \nu_i \alpha_i}{4} + \nu_i^2 \frac{2 - 3\rho_i^2}{24} \right) (T_i - t) \right] \\ z &:= \frac{\nu_i}{\alpha_i} \ln \left(\frac{\tilde{Y}_i(t)}{K} \right) \\ x(z) &:= \ln \left(\frac{\sqrt{1 - 2\rho_i z + z^2} + z - \rho_i}{1 - \rho_i} \right) \end{aligned}$$

The resulting pricing formula means only the SABR parameters α_i, ρ_i and ν_i have to be calibrated when pricing Year-on-Year caplets. For the analytical pricing of caps/floors and year-on-year swaps on the CPI, a separate approximation procedure is derived, which results in a closed-form SABR formula.

3.4 Exact Solution Using the Cos Method

A commonly used method for the analytical pricing of financial derivatives under Heston models with known characteristic functions is the COS method introduced by Fang and Oosterlee [19]. This method is applicable to the Heston model when this characteristic function can be derived in closed form such as the Heston model with constant interest rates. The Fourier methods involved can also be extended to the Heston-Hull-White model[43]. In these cases, the exact solution can be used for direct comparison with the results using Monte Carlo simulation.

For the Heston model with constant interest rates, the risk-neutral value of an option can be written as

$$V(x, t_0) = e^{-r\Delta t} \mathbb{E}_{\mathbb{Q}}(V(y, T)|x) = e^{-r\Delta t} \int_{\mathbb{R}} V(y, T) f(y|x) dy \quad (3.27)$$

where V is the option value, T the maturity date, t_0 the initial date, $\Delta t = T - t_0$. x and y are state variables representing the payoff function at time t_0 and T respectively, and $f(y|x)$ the probability density of y given x , and r is the risk-neutral interest rate.

By truncating the infinite integration range and replacing the density with its cosine expansion, Fang and Oosterlee obtain the approximation

$$\tilde{V}(x, t_0) = e^{-r\Delta t} \int_b^a V(y, T) \sum_{k=0}^{\infty} A_k(x) \cos\left(k\pi \frac{y-a}{b-a}\right) dy, \quad (3.28)$$

with

$$A_k(x) := \frac{2}{b-a} \int_a^b f(y|x) \cos\left(k\pi \frac{y-a}{b-a}\right) dy. \quad (3.29)$$

By using the definition

$$U_k := \frac{2}{b-a} \int_a^b V(y, T) \cos\left(k\pi \frac{y-a}{b-a}\right) dy, \quad (3.30)$$

the value of the option is written in terms of the Fourier-cosine series of $f(y|x)$ and $V(y, T)$,

$$\tilde{V}(x, t_0) = \frac{1}{2}(b-a)e^{-r\Delta t} \sum_{k=0}^{\infty} A_k(x) U_k. \quad (3.31)$$

Due to the periodic nature of the cosine series, the coefficients decay quickly, so only a finite number of terms can be used to obtain sufficiently accurate results. Additionally the coefficients A_k can be approximated by

$$A_k \approx F_k \equiv \frac{2}{b-a} \operatorname{Re} \left(\phi \left(\frac{k\pi}{b-a} \right) \exp \left(-i \frac{ka\pi}{b-a} \right) \right) \quad (3.32)$$

with ϕ the characteristic function. This combined results in the following approximation

$$\tilde{V}(x, t_0) = e^{-r\Delta t} \sum_{k=0}^N \operatorname{Re} \left(\phi \left(\frac{k\pi}{b-a}; x \right) e^{-ik\pi \frac{a}{b-a}} \right) U_k. \quad (3.33)$$

Equation (3.33) is referred to as the COS formula for general underlying processes. The coefficients U_k are obtained analytically for vanilla options

$$U_k^{call} = \frac{2}{b-a} K(\chi_k(a, b, 0, b) - \psi_k(a, b, 0, b)), \quad (3.34)$$

with

$$\begin{aligned} \chi_k(a, b, c, d) = & \frac{1}{1 + \left(\frac{k\pi}{b-a}\right)^2} \left[\cos \left(k\pi \frac{d-a}{b-a} \right) e^d - \cos \left(k\pi \frac{c-a}{b-a} \right) e^c \right. \\ & \left. + \frac{k\pi}{b-a} \sin \left(k\pi \frac{d-a}{b-a} \right) e^d - \frac{k\pi}{b-a} \sin \left(k\pi \frac{c-a}{b-a} \right) e^c \right], \end{aligned} \quad (3.35)$$

and

$$\psi_k(a, b, c, d) = \begin{cases} \frac{b-a}{k\pi} (\sin(k\pi \frac{d-a}{b-a}) - \sin(k\pi \frac{c-a}{b-a})) & k \neq 0, \\ d - c & k = 0. \end{cases} \quad (3.36)$$

To find the exact solution only the characteristic function ϕ , a truncation range for the integral, and a truncation value N is required. However, as the models include more realistic processes for the interest rate, the characteristic function is only known in closed form for specific assumptions.

4 Monte Carlo Simulation

As observed in the previous sections not all model setups admit analytical solutions for all parameter combinations, and those that do can rely on undesirable assumptions. Monte Carlo methods, widely used in the pricing of financial derivatives, have efficiently addressed this problem [24][5][12]. These methods use stochastic simulations of future paths of underlying processes and the strong law of large numbers to estimate expectations of future values. This enables relatively easy evaluation of complex pricing formulas where analytical solutions may not be feasible. Monte Carlo simulations have been employed to verify or compare results from analytical pricing formulas for many types of derivative structures. Furthermore, Monte Carlo methods offer great flexibility in pricing more complex derivatives. Such options do not exist yet for inflation, but may be introduced for institutions looking to manage inflation risk in a flexible manner.

At its simplest, Monte Carlo can be used to estimate an expectation $\mathbb{E}[P]$. A standard Monte Carlo estimate \hat{P}_{MC} is an average of values $P(\omega)$ for N independent samples ω from a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$,

$$\hat{P}_{MC} = N^{-1} \sum_{n=1}^N P(\omega^{(n)}), \quad (4.1)$$

where $P(\omega^{(n)})$ is the n 'th sample. Due to the Strong Law of Large Numbers this estimator converges to the average $\mathbb{E}[P]$ as $N \rightarrow \infty$. With the Central Limit Theorem it can be shown that the convergence rate is of order $\frac{1}{\sqrt{N}}$. Additionally a $1 - \alpha$ confidence interval is provided:

$$\left[\hat{P}_{MC} - \phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\frac{\mathbb{V}(P(\omega))}{N}}, \hat{P}_{MC} + \phi^{-1} \left(1 - \frac{\alpha}{2} \right) \sqrt{\frac{\mathbb{V}(P(\omega))}{N}} \right], \quad (4.2)$$

where ϕ^{-1} is the inverse of the cumulative distribution function of the standard normal distribution.

There exist two sources of error in the estimator, the sampling error related to estimating the expected value with a finite sample average, and the approximation of P by $P(\omega)$ which is related to the discretization of the continuous model. The contribution of each becomes clear when the mean square error (MSE) is expanded,

$$\begin{aligned} MSE(\hat{P}_{MC}) &= \mathbb{E} \left[\left(\hat{P}_{MC} - \mathbb{E}(P) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{P}_{MC} - \mathbb{E}(\hat{P}_{MC}) + \mathbb{E}(\hat{P}_{MC}) - \mathbb{E}(P) \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{P}_{MC} - \mathbb{E}(\hat{P}_{MC}) \right)^2 \right] + \left[\mathbb{E}(\hat{P}_{MC}) - \mathbb{E}(P) \right]^2 \\ &= \mathbb{V}(\hat{P}_{MC}) + \left[\mathbb{E}(\hat{P}_{MC}) - \mathbb{E}(P) \right]^2 \\ &= \frac{\mathbb{V}(P(\omega))}{N} + (\mathbb{E}[P(\omega) - P])^2. \end{aligned} \quad (4.3)$$

Hence the first term in the MSE is the variance of the Monte Carlo estimator which represents the sampling error. The second term is the square of the error in the mean between the samples and the exact value. From this it is apparent the root mean square error is $O(1/\sqrt{N})$. This means the computational cost can be very high for a given accuracy, especially when each sample $P(\omega)$ requires an approximate solution, or the computation of many time steps.

Several approaches have been considered to address the high computational cost. One class of methods is called the control variate approach which reduce the variance using prior knowledge such as the exact solution of a value correlated to P . Multi-level Monte Carlo is a method which constructs its own control variate in the process by simulating at a coarse and fine level of time steps simultaneously. Before Monte Carlo can be used, the continuous model in question must be discretized.

4.1 Simulation of the Heston model

For the valuation of inflation caps/floors using simulation, a numerical discretization scheme is required. Since the stochastic model is continuous it cannot be used directly and needs to be approximated. For the Heston model in particular much attention has been devoted to constructing efficient simulation methods. Consider an arbitrary set of discrete times $\mathcal{T} = \{t_i\}_{i=1}^M$. Now consider the problem of generating random paths of $(I(t), V(t))$ for all $t \in \mathcal{T}$. The interest lies in generating random variables $(I(t+\Delta t), V(t+\Delta t))$ conditional on $(I(t), V(t))$, with an arbitrary time step Δt . By repeating the generation of these random variables a full path is produced $(I(t), V(t))_{t \in \mathcal{T}}$. This is notably required for path-dependent options where the payoff depends on multiple values of the inflation index over time. Throughout this thesis, approximations of $I(t)$ and $V(t)$ are denoted by $\hat{I}(t)$ and $\hat{V}(t)$.

4.1.1 Euler discretization

A general and well studied approach to simulate the Heston model is using the Euler discretization scheme. Applying the Euler scheme to the Heston model with constant nominal and real interest rates results in

$$\hat{I}(t + \Delta t) = I(t) + (r_N - r_R)I(t)\Delta t + I(t)\sqrt{V(t)}\Delta W^I(t), \quad (4.4)$$

$$\hat{V}(t + \Delta t) = V(t) + \alpha(\bar{V} - V(t))\Delta t + \sigma_V\sqrt{V(t)}\Delta W^V(t). \quad (4.5)$$

Using this scheme for the variance process leads to a positive probability of the variance becoming negative,

$$\mathbb{P}(\hat{V}(t + \Delta t) < 0 | \hat{V}(t) > 0) = 1 - \Phi\left(\frac{(1 - \alpha\Delta t)\hat{V}(t) + \alpha\bar{V}\Delta t}{\sigma_V\sqrt{\hat{V}(t)\Delta t}}\right). \quad (4.6)$$

This means the scheme can break down since negative values appear in the square root term, especially when the Feller condition $2\alpha\theta > \sigma_V$ is not satisfied so that a high probability mass of the variance is concentrated around the boundary. Lord et al.[34] discuss several methods to approach this issue. They recommend combining

several fixes into a single framework called the full truncated (FT) scheme. For the variance process this is given by,

$$\hat{V}(t + \Delta t) = \max((\hat{V}(t) + \alpha(\bar{V} - ((\hat{V}(t))^+))\Delta t + \sigma_V \sqrt{(\hat{V}(t))^+} \Delta W_V(t)), 0), \quad (4.7)$$

guaranteeing the scheme cannot break down due to taking the square root of a negative value. However, even when the Feller condition is fulfilled this scheme still introduces bias due to the correction for negative values.

Since the log-inflation index is log-normally distributed, usually the log-normal Heston model is preferred for numerical practice. This results in the following simulation for the inflation index,

$$\hat{I}(t + \Delta t) = \hat{I}(t) e^{(r_N - r_R - \frac{1}{2}\hat{V}(t))\Delta t + \sqrt{\hat{V}(t)\Delta W^I(t)}} \quad (4.8)$$

The correlated Brownian motions can be implemented using Cholesky decomposition.

4.1.2 Exact simulation

An alternative popular approach for simulating the Heston model is using exact simulation, first introduced by Broadie and Kaya [9]. They use the property that the variance process at time $t + \Delta t$ conditioned on time t , $0 < \Delta t$, follows a scaled noncentral chi-square distribution, a property which was first observed by Cox et al. [15],

$$V_{t+\Delta t} \stackrel{d}{=} \frac{\sigma_V^2(1 - e^{-\alpha(t-u)})}{4\alpha} \chi_d^2 \left(\frac{4\alpha e^{-\alpha(t-u)}}{\sigma_V^2(1 - e^{-\alpha(t-u)})} V_t \right), \quad (4.9)$$

where $\chi_d^2(\lambda)$ denotes a noncentral chi-squared random variable with $d := \frac{4\bar{V}\alpha}{\sigma_V^2}$ and non-centrality parameter λ .

For simulation purposes, two representations of the noncentral $\chi_d^2(\lambda)$ distribution are mentioned which can be used:

$$\chi_d^2(\lambda) \stackrel{d}{=} \begin{cases} (Z + \sqrt{\lambda})^2 + \chi_{d-1}^2 & \text{for } d > 1, \\ \chi_{d+2N}^2 & \text{for } d > 0, \end{cases} \quad (4.10)$$

where $Z \sim N(0, 1)$, χ_v^2 is an ordinary chi-squared distribution with v degrees of freedom, and N is Poisson distributed with mean $\mu := \frac{1}{2}\lambda$. In most practical applications in finance $d \ll 1$ so one is forced to work with the second representation.

A bias-free simulation scheme for the inflation index is obtained by using a different representation of (4.5) and (4.8). The value of the inflation at time $t + \Delta t$, given the value $\hat{I}(t)$ and the sampled $\hat{V}(t)$, can be written as

$$\begin{aligned} \hat{I}(t + \Delta t) = \hat{I}(t) \exp & \left((r_n - r_r)\Delta t - \frac{1}{2} \int_t^{t+\Delta t} \hat{V}(s) ds + \rho \int_t^{t+\Delta t} \sqrt{\hat{V}(s)} dW^V(s) \right. \\ & \left. + \sqrt{1 - \rho^2} \int_t^{t+\Delta t} \sqrt{\hat{V}(s)} dW^2(s) \right). \end{aligned} \quad (4.11)$$

where $dW^V(t)$ and $dW^2(t)$ are independent Brownian motions and the Cholesky decomposition $dW^I(t) := \rho dW^V(t) + \sqrt{1 - \rho^2} dW^2(t)$ is used to represent the correlated Brownian motion.

Furthermore, by integrating the SDE for the variance process (4.5) and rearranging,

$$\int_t^{t+\Delta t} \sqrt{V(s)} dW^V(s) = \frac{1}{\sigma} \left(V(t + \Delta t) - V(t) - \alpha \bar{V} \Delta t + \alpha \int_t^{t+\Delta t} V(s) ds \right), \quad (4.12)$$

which when substituting into equation (4.11) leads to

$$\begin{aligned} \hat{I}(t + \Delta t) &= \hat{I}(t) \exp \left((r_n - r_r) \Delta t - \frac{1}{2} \int_t^{t+\Delta t} \hat{V}(s) ds \right. \\ &\quad \left. + \frac{\rho}{\sigma_V} \left(\hat{V}(t + \Delta t) - \hat{V}(t) - \alpha \bar{V} \Delta t + \alpha \int_t^{t+\Delta t} \hat{V}(s) ds \right) \right. \\ &\quad \left. + \sqrt{1 - \rho^2} \int_t^{t+\Delta t} \sqrt{\hat{V}(s)} dW^2(s) \right) \\ &\stackrel{d}{=} \hat{I}(t) \exp \left((r_n - r_r) \Delta t - \frac{1}{2} \int_t^{t+\Delta t} \hat{V}(s) ds \right. \\ &\quad \left. + \frac{\rho}{\sigma_V} \left(\hat{V}(t + \Delta t) - \hat{V}(t) - \alpha \bar{V} \Delta t + \alpha \int_t^{t+\Delta t} \hat{V}(s) ds \right) \right. \\ &\quad \left. + \sqrt{1 - \rho^2} \sqrt{\int_t^{t+\Delta t} \hat{V}(s) ds} N \right). \end{aligned} \quad (4.13)$$

where N is a standard normal variable. The final step is due to V being independent of W^2 , which means given $\int_t^{t+\Delta t} V(s) ds$, $\int_t^{t+\Delta t} \sqrt{\hat{V}(s)} dW^2(s)$ is normally distributed with mean 0 and variance $\int_t^{t+\Delta t} V(s) ds$. Exact simulation of the Heston model therefore comes down to

1. Sampling $V(t + \Delta t)$ using its distributional properties (4.10).
2. Sampling the integrated variance process $\int_t^{t+\Delta t} V(s) ds$ conditional on $V(t)$ and $V(t + \Delta t)$.

Broadie and Kaya [9] suggest sampling $V(t + \Delta t)$ using an acceptance and rejection method to generate gamma variates. They then sample the integrated variance process by deriving the characteristic function of its distribution and numerically inverting it using a root-finding procedure. This requires a time-consuming iterative Fourier inversion method because the characteristic function involves Bessel functions with complex arguments. While this is theoretically an exact procedure, all the steps combined require immense computational effort and implementation has to be done with care due to the involved procedures. Especially for the purposes of simulation of complete paths with many time steps and long maturities, it is computationally inconvenient.

A straightforward approximation for a sample of the conditional integrated variance process is a simple drift interpolation. That is, the integral is approximated by

$$\int_t^{t+\Delta t} V(s)ds := (c_1V(t) + c_2V(t + \Delta t)) \Delta t \approx \int_t^{t+\Delta t} V(s)ds. \quad (4.14)$$

for some constants c_1, c_2 . Due to the difficulty in sampling the integrated variance process exactly, van Haastrecht [41] mentions that it is computationally more efficient to use this approximation. Anderson also [1] uses this as an alternative to the Fourier methods. Using $c_1 = c_2 = \frac{1}{2}$ corresponds to the trapezoidal method, which is used by Zheng [44] with the purpose of applying the multilevel Monte Carlo method discussed in the next chapter.

4.1.3 NCI scheme

For the Broadie and Kaya scheme in the previous section, the sampling of the non-central $\chi_d^2(\lambda)$ distribution is done by first conditioning on a Poisson variate followed by consecutively generating a sample from a chi-squared or gamma distribution ($\chi_v^2 \stackrel{d}{=} \text{Gamma}(\frac{v}{2}, 2)$). Broadie and Kaya suggest the use of an acceptance and rejection method to generate gamma variates. Van Haastrecht and Pelsser [41] mention this limits the practical use of their scheme because the number of samples depends on the specific model parameters. They suggest what is now considered an "almost exact" scheme as an alternative, where almost exact indicates any desired accuracy can be achieved, and the effect of the time step size on the approximation error is negligible.

Let Q_{max} be a positive integer, $\mathcal{Q} := \{0, 1, 2, \dots, Q_{max}\}$ be a set of Poisson values, and $\mathcal{U} := \{0, \dots, 1 - \delta\}$ be an equidistant grid with $\delta < 10^{-15}$ (the reason for this is that the quantile function evaluated in 1 is infinity). The NCI method starts with a precomputation of the inverse of the chi-square distributions on this grid, i.e.

$$H_Q^{-1}(U) := G_{\chi_{d+2Q}^2}^{-1}(U), \text{ for all } Q \in \mathcal{Q}, U \in \mathcal{U}. \quad (4.15)$$

Remember that the variance process V follows equation (4.9). The method continues by generating one sample Q from the Poisson distribution with mean $\frac{\lambda}{2}$, and another sample U_1 from the uniform distribution on $(0, 1)$. A sample of $V(t + \Delta t)$ is then generated using an interpolation of the precomputed grid,

$$F_Q^{-1}(U) := \begin{cases} C_0 J(U), & \text{for } Q \leq Q_{max} \\ C_0 F_{\chi_{d+2Q}^2}^{-1}(U) & \text{for } Q > Q_{max} \end{cases} \quad (4.16)$$

where J is an interpolation rule based on the precomputed values $H_Q^{-1}(\cdot)$, and $F_{\chi_{d+2Q}^2}^{-1}(\cdot)$ represents the unlikely but possible event the Poisson sample $Q > Q_{max}$. In the latter case $F_Q^{-1}(U)$ is generated using other methods, such as generating gamma variates. Using representation (4.10) the sample for $V(t + \Delta t)$ can be calculated. Van Haastrecht and Pelsser [41] suggest two interpolation rules for J , linear interpolation which is fast to execute, and monotone cubic Hermite spline which might

require less points to be cached for the same accuracy. Because of its simplicity the results in this thesis use the former. Additionally, for high precision Q_{max} is set to 40, and $|\mathcal{U}|$ to 10000.

4.1.4 QE scheme

As an approximation to sampling from the noncentral chi-squared distribution, Anderson [1] suggests drawing from a related distribution and moment matching with the first two moments of the noncentral chi-squared distribution. The choice of distribution is decided in two parts.

1. For a moderate noncentrality parameter, Anderson states that the noncentral chi-squared distribution can be represented by a power function applied to a Gaussian variable. For sufficiently high values of $V(t)$ a sample of $V(t + \Delta t)$ can be generated by

$$V(t + \Delta t) = a(b + Z)^2, \quad (4.17)$$

where Z is a standard normal variable, and a and b are constants.

2. For sufficiently low values of $V(t)$, the density of $V(t + \Delta t)$ can be approximated by

$$\mathbb{P}(V(t + \Delta t) \in [x, x + dx]) \approx (p\delta(0) + \beta(1 - p)e^{-\beta x})dx, \quad (4.18)$$

where δ is the Dirac delta function, and p and β are non negative constants. Sampling from this distribution is done by inverting the distribution function. The distribution function can be obtained by integrating the probability density function and is given by:

$$L^{-1}(u) = \begin{cases} 0 & \text{for } 0 \leq u \leq p \\ \frac{1}{\beta} \log\left(\frac{1-p}{1-u}\right) & \text{for } p < u \leq 1 \end{cases} \quad (4.19)$$

The values of the constants a, b, p and β are determined by moment-matching. What remains is defining a rule that states when to switch between using the first method and the second. Anderson bases the rule on the value of $\psi = \frac{s^2}{m^2}$ where m is the conditional mean and s^2 the conditional variance of the variance process, which are given by

$$m := \bar{V} + (V(t) - \bar{V})e^{-\alpha\Delta t}, \quad (4.20)$$

$$s^2 := \frac{V(t)\sigma_V e^{-\alpha\Delta t}}{\alpha}(1 - e^{-\alpha\Delta t}) + \frac{\bar{V}\sigma_V^2}{2\alpha}(1 - e^{-\alpha\Delta t})^2. \quad (4.21)$$

It can be shown that for $\psi \leq 2$, the quadratic scheme (4.17) can be moment matched with the exact distribution, and for $\psi \geq 1$ the exponential scheme (4.19) can be moment matched. Thus, a value $\psi_c \in [1, 2]$ can be chosen where (4.17) is used if $\psi \leq \psi_c$ and (4.19) when $\psi \geq \psi_c$. Anderson notes that the exact choice of ψ_c does not have a major impact on the quality of the simulation scheme and suggests using $\psi_c = 1.5$.

Moment-matching the constants a, b, p and β results in the following.

For $\psi \leq 2$

$$b^2 = \frac{2}{\psi} - 1 + \sqrt{\frac{2}{\psi} - 1} \sqrt{\frac{2}{\psi} - 1} \quad (4.22)$$

$$a = \frac{m}{1 + b^2} \quad (4.23)$$

For $\psi \geq 1$,

$$p = \frac{\psi - 1}{\psi + 1} \quad (4.24)$$

$$\beta = \frac{1 - p}{m} = \frac{2}{m(\psi + 1)} \quad (4.25)$$

As a result, the QE algorithm is as follows

1. Given $V(t)$, calculate m , s^2 and $\psi = \frac{m^2}{s^2}$.
2. If $\psi \leq \psi_c$:
 - Compute a and b .
 - Generate a sample Z from the standard normal distribution.
 - Set $V(t + \Delta t) = a(b + Z)^2$.
3. If $\psi > \psi_c$:
 - Compute β and p .
 - Generate a uniform random number U .
 - Set $V(t + \Delta t) = L^{-1}(U)$.

This is the extent of the simulation schemes considered for the variance process in this thesis. An additional moment matching scheme is proposed by Anderson [1] called the Truncated Gaussian scheme and van Haastrecht [41] also proposed a combination of the QE and NCI schemes. Furthermore, martingale corrections are also discussed, since a discretization scheme may not always satisfy the martingale assumptions. Both authors state the practical relevance of a martingale correction is often minor and trying to implementing this may not result in more accurate option prices for the same computation time. This could be an interesting point of further research, however, it is not considered for the results of this thesis.

4.2 Simulation of the interest rates

So far the simulation schemes have involved constant interest rates due to the application to the Heston model. For the Heston Hull-White model, the interest rates also have distributional properties which enable exact simulation. Assuming the

nominal interest rate has Hull-White dynamics, following Brigo and Mercurio [7, p.73] and Section 3.2, using the integrating factor results in

$$r_n(t) = r_n(s)e^{-a_n(t-s)} + \int_s^t e^{-a_n(t-u)}\theta(u)du + \sigma_n \int_s^t e^{-a_n(t-u)}dW^{r_n}(u), \quad (4.26)$$

$$= r_n(s)e^{-a_n(t-s)} + \alpha_n(t) - \alpha_n(s)e^{-a_n(t-s)} + \sigma_n \int_s^t e^{-a_n(t-u)}dW^{r_n}(u) \quad (4.27)$$

where $s < t$ and

$$\alpha_n(t) = f_n(0, t) + \frac{\sigma_n^2}{2a_n^2}(1 - e^{-a_n t})^2. \quad (4.28)$$

Under the nominal risk-neutral measure, the real interest rate similarly satisfies,

$$\begin{aligned} r_r(t) &= r_r(s)e^{-a_r(t-s)} + \int_s^t e^{-a_r(t-u)}\theta(u)du + \sigma_r \int_s^t e^{-a_r(t-u)}dW^{r_r}(u) \\ &\quad - \int_s^t e^{-a_r(t-u)}\rho_{I,r}\sigma_r\sqrt{V(u)}du. \end{aligned} \quad (4.29)$$

By approximating the integrated square root variance with the trapezoidal rule described in Section 4.1.2,

$$\int_s^t e^{-a_r(t-u)}\sqrt{V(u)}du := \frac{\sqrt{V(t)} + e^{-a_r(t-s)}\sqrt{V(s)}}{2}(t-s), \quad (4.30)$$

the real interest rate at time t can be written as:

$$\begin{aligned} r_r(t) &= r_r(s)e^{-a_r(t-s)} + \alpha(t) - \alpha(s)e^{-a_r(t-s)} + \sigma_r \int_s^t e^{-a_r(t-u)}dW^{r_r}(u) \\ &\quad - \rho_{I,r}\sigma_r \frac{\sqrt{V(t)} + e^{-a_r(t-s)}\sqrt{V(s)}}{2}(t-s). \end{aligned} \quad (4.31)$$

Therefore, using that a deterministic integrand with respect to a Wiener process is normally distributed, $r_n(t)$ and $r_r(t)$ conditional on \mathcal{F}_s and $V(t)$ are normally distributed. The means and variances are given respectively by,

$$\begin{aligned} \mathbb{E}(r_n(t)|\mathcal{F}_s) &= r_n(s)e^{-a_n(t-s)} + \alpha_n(t) - \alpha_n(s)e^{-a_n(t-s)}, \\ \mathbb{E}(r_r(t)|\mathcal{F}_s) &= r_r(s)e^{-a_r(t-s)} + \alpha_r(t) - \alpha_r(s)e^{-a_r(t-s)} \\ &\quad - \rho_{I,r}\sigma_r \frac{\sqrt{V(t)} + e^{-a_r(t-s)}\sqrt{V(s)}}{2}(t-s), \end{aligned} \quad (4.32)$$

$$\mathbb{V}(r_n(t)|\mathcal{F}_s) = \frac{\sigma_n^2}{2a_n} (1 - e^{-2a_n(t-s)}),$$

$$\mathbb{V}(r_r(t)|\mathcal{F}_s) = \frac{\sigma_r^2}{2a_r} (1 - e^{-2a_r(t-s)}).$$

From here a sample of the interest rates conditional on $r_n(t), r_r(t)$ can be obtained by generating a sample from the correlated standard normal distributions and computing,

$$\hat{r}_n(t + \Delta t) = \mathbb{E}(r_n(t)|\mathcal{F}_t) + \sqrt{\mathbb{V}(r_n(t)|\mathcal{F}_s)}Z_n, \quad (4.33)$$

$$\hat{r}_r(t + \Delta t) = \mathbb{E}(r_r(t)|\mathcal{F}_t) + \sqrt{\mathbb{V}(r_r(t)|\mathcal{F}_s)}Z_r. \quad (4.34)$$

Singor et al. [43] briefly mention the use of exact simulation methods for the variance and the interest rate. Since multiple methods exist it is unclear how the integrated square root variance term is sampled or approximated. Here, for simplicity, the integrated term is approximated using the drift interpolation used in the previous section. Exact simulation of the real interest rate may be possible but is left for future studies.

4.3 Multilevel Monte Carlo (MLMC)

The multilevel Monte Carlo (MLMC) method was introduced by Giles [23] as a variance reduction technique for the standard Monte Carlo method. Given a sequence P_0, \dots, P_{L-1} which approximates P_L with increasing accuracy and cost, then the expected value $\mathbb{E}(\hat{P}_L)$ can be written as

$$\mathbb{E}(\hat{P}_L) = \mathbb{E}(\hat{P}_0) + \sum_{l=1}^L \mathbb{E}(\hat{P}_l - \hat{P}_{l-1}), \quad (4.35)$$

where each expectation is estimated independently. As a result we can use the following unbiased estimator for $\mathbb{E}(\hat{P}_L)$,

$$\frac{1}{N_0} \sum_{i=1}^{N_0} P_0^{0,i} + \sum_{l=1}^L \left[\frac{1}{N_l} \sum_{i=1}^{N_l} (P_l^{\ell,i} - P_{l-1}^{\ell,i}) \right], \quad (4.36)$$

where N_ℓ is the number of samples at each level, and the superscript (ℓ, i) indicates the i 'th independent sample at each level.

This is the general MLMC method in which the output $\mathbb{E}(P_L)$ corresponds to the quantity of interest. However, as mentioned by Giles, in many applications involving the simulation of SDEs, the output P_ℓ at level ℓ is an approximation to the random variable P , which cannot be simulated exactly. Let Y be the multilevel estimator

$$Y = \sum_{l=0}^L Y_\ell \quad (4.37)$$

where Y_ℓ are the estimators for each level

$$Y_\ell = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (P_\ell^{\ell,i} - P_{l-1}^{\ell,i}), \quad (4.38)$$

with $P_{-1} \equiv 0$ making Y_0 a standard Monte Carlo estimator. Because the expectations $\mathbb{E}(Y_\ell)$ are estimated independently, this results in

$$\mathbb{E}[Y] = \mathbb{E}[P_L], \quad (4.39)$$

$$\mathbb{V}[Y] = \sum_{l=0}^L \frac{V_\ell}{N_\ell}, \quad (4.40)$$

$$(4.41)$$

where $\mathbb{V}[\cdot]$ is the variance and V_ℓ defined as

$$V_\ell \equiv \mathbb{V}[P_\ell - P_{\ell-1}] \quad (4.42)$$

This shows Y is an approximation to $\mathbb{E}(P)$. Expanding the mean square error as in equation (4.3) results in

$$\text{MSE}(Y) := \mathbb{E}[(Y - \mathbb{E}[P])^2] \quad (4.43)$$

$$= \mathbb{V}[Y] + (\mathbb{E}[Y] - \mathbb{E}[P])^2, \quad (4.44)$$

$$= \sum_{l=0}^L \frac{V_\ell}{N_\ell} + (\mathbb{E}[Y] - \mathbb{E}[P])^2. \quad (4.45)$$

Just as in the standard Monte Carlo case, the MSE consists of discretization error, which is exactly the same, and the sampling error. To achieve a mean square error of ϵ^2 both terms have to be less than $\epsilon^2/2$. Intuitively, the computational cost of MLMC should be cheaper than standard Monte Carlo because of the following:

- If Y converges to $\mathbb{E}(P)$ in mean square then $\mathbb{V}(Y_\ell) \rightarrow 0$ as $\ell \rightarrow \infty$, indicating fewer samples are required to estimate $\mathbb{E}(Y)$ as the levels become finer.
- The coarse level $\ell = 0$ can be kept fixed for all ϵ so the cost per sample does not increase as $\epsilon \rightarrow 0$.

The computational cost C of the estimator is

$$C(Y) = \sum_{\ell=0}^L N_\ell C_\ell, \quad (4.46)$$

where $C_\ell = C(Y_\ell)$ represents the cost of a sample of Y_ℓ . The following theorem which is a generalization of the theorem in Giles [23] presents an upper bound for C . The proof is provided by Cliffe et al. [14].

Theorem 1. *Let P denote a function of the solution of the stochastic differential equation for a given Brownian path, and let \hat{P}_ℓ be the approximation corresponding to a numerical discretization with time step $h_\ell = M^{-l}T$. If there exist independent estimators Y_ℓ based on N_ℓ Monte Carlo samples, each with expected cost C_ℓ and variance V_ℓ , and positive constants $\alpha, \beta, c_1, c_2, c_3$ such that $\alpha \geq \frac{1}{2} \min\{\beta, \gamma\}$ and*

$$1. \mathbb{E}(\hat{P}_\ell - P) \leq c_1 h_\ell^\alpha$$

$$2. \mathbb{E}(Y_\ell) = \begin{cases} \mathbb{E}(\hat{P}_0), & l = 0 \\ \mathbb{E}(\hat{P}_\ell - \hat{P}_{\ell-1}), & l > 0 \end{cases}$$

$$3. \mathbb{V}(Y_\ell) \leq c_2 N_\ell^{-1} h_\ell^\beta$$

$$4. C_\ell \leq c_3 N_\ell h_\ell^\gamma,$$

then there exists a positive constant c_4 such that for any $\epsilon < e^{-1}$ there are values L and N_ℓ for which the multilevel estimator (4.37) has a bounded MSE

$$MSE = \mathbb{E}[(Y - \mathbb{E}(P))^2] < \epsilon^2 \quad (4.47)$$

with a computational complexity C bounded by

$$C \leq \begin{cases} c_4 \epsilon^{-2}, & \beta > \gamma \\ c_4 \epsilon^{-2} (\ln \epsilon)^2, & \beta = \gamma \\ c_4 \epsilon^{-2 - (1-\beta)/\alpha}, & \beta < \gamma. \end{cases} \quad (4.48)$$

Proof. See appendix A.1. □

In the MLMC algorithm the values of L and N_ℓ can be calculated adaptively. If C_0 and V_0 are the cost and variance of one sample of P_0 , and C_ℓ and V_ℓ the cost and variance of one sample of $P_\ell - P_{\ell-1}$, then the total cost and variance of the multilevel estimator are $\sum_{\ell=0}^L N_\ell C_\ell$ and $\sum_{\ell=0}^L \frac{V_\ell}{N_\ell}$ respectively. For a fixed variance, the cost is minimized by choosing N_ℓ to minimize

$$\sum_{\ell=0}^L \left(N_\ell C_\ell + \mu^2 \frac{V_\ell}{N_\ell} \right) \quad (4.49)$$

for some value μ^2 known as the Lagrange multiplier. This results in $N_\ell = \mu \sqrt{V_\ell / C_\ell}$. To obtain the variance $\frac{1}{2} \epsilon^2$ then requires $\mu = \frac{1}{\epsilon^2} \sum_{\ell=0}^L \sqrt{V_\ell C_\ell}$ which results in

$$N_\ell = \left\lceil \frac{2}{\epsilon^2} \sqrt{V_\ell h_\ell} \left(\sum_{\ell=0}^L \sqrt{\frac{V_\ell}{h_\ell}} \right) \right\rceil. \quad (4.50)$$

The MLMC algorithm calculates the optimal value N_ℓ adaptively using the (unbiased) estimates for the variances from the simulation, and the number of levels L is increased until a MSE less than ϵ^2 is achieved. The test for convergence to achieve this MSE tries to make sure that $|\mathbb{E}[P_L - P]| < \frac{\epsilon}{\sqrt{2}}$. If $\mathbb{E}[P_\ell - P_{\ell-1}] \propto h_\ell^\alpha$ as the theorem suggests, then the remaining error is

$$\mathbb{E}[P_L - P] = \sum_{\ell=L+1}^{\infty} \mathbb{E}[P_\ell - P_{\ell-1}] = \frac{\mathbb{E}[P_L - P_{L-1}]}{M^\alpha - 1}, \quad (4.51)$$

which means a convergence test is

$$\frac{|\mathbb{E}[P_L - P_{L-1}]|}{M^\alpha - 1} < \frac{\epsilon}{\sqrt{2}} \quad (4.52)$$

In Giles original paper[23] a short analysis about the optimal value of M is provided. For the numerical results in this thesis M is set to 4 as suggested.

Combining the ideas from this section, the MLMC algorithm given in algorithm 1 is implemented using the following algorithm.

Algorithm 1 Multilevel Monte Carlo algorithm

Start with $L = 2$.
Construct an initial $N_l = 10000$ samples of P_ℓ and $P_{\ell-1}$ for $\ell = 0, 1, 2$.
Calculate Y_ℓ from equation (4.37).
Estimate \hat{V}_ℓ .
Calculate the optimal number of samples N_l using equation (4.50).
while extra samples required **do**
 Compute extra samples for each level until the optimal number is reached.
 Update \hat{V}_ℓ and N_l .
 Test for convergence using (4.52).
 If not converged, set $L := L + 1$.
end while

4.4 Convergence properties

As mentioned at the start of this chapter, it is important to distinguish between the convergence of an unbiased estimator and of a discretization scheme since both contribute to the convergence of the MSE. The importance of both is shown in Theorem 1 in the form of parameters β and α . In practice, for condition 1. the value of α is often established from previous research into the weak convergence of a discretization scheme, and also might depend on the value M .

The main challenge lies in determining the correct value for β , and whether it is possible to develop estimators with a larger value of β . For options with bounded and Lipschitz continuous payoffs often the order of convergence of the variance of the estimator can be derived. In the original paper, Giles [23] shows that for the Euler discretization scheme applied to geometric Brownian motion with $M = 4$, $\alpha = \beta = 2$. However, in the case of the Heston model the volatility does not satisfy a global Lipschitz condition and at the time of this paper no theory existed to predict the order of convergence. Only recently, Mickel and Neuenkirch [39] showed with the FT scheme applied to the Heston model that when the Feller condition is satisfied, $\alpha = 1$ and when it is not satisfied $\alpha = 1 - a$ with a arbitrarily small. Zheng [44] proves that for the path-dependent exact simulation scheme that $\beta = 1 - \epsilon$ for all parameters of the Heston model.

The theorems and proofs on (weak) convergence of the simulation schemes are relatively complex and so for the sake of brevity are not discussed in detail in this thesis.

4.5 MLMC for path-dependent options using the Heston model

To capture the full flexibility of possible option payoffs the full sample paths must be simulated. The key step in using MLMC is to establish a connection between the estimated inflation index at the coarse level and the estimated inflation index at the fine level. This is achieved by specifying the same Brownian path for the coarse and fine approximation at each level.

Consider the Heston model with constant interest rates. Let $\hat{I}^f(t)$ and $\hat{V}^f(t)$, $t = 0, \Delta t, 2\Delta t, \dots, T$ be an approximation of the path of I and V respectively, at the fine level with step size Δt . Let $\hat{I}^c(t)$ and $\hat{V}^c(t)$, $t = 0, M\Delta t, 2M\Delta t, \dots, T$ be an approximation of I and V respectively, at the coarse level with step size $M\Delta t$. Additionally, $\hat{I}^c(0) = \hat{I}^f(0) = I(0)$ and $\hat{V}^f(0) = \hat{V}^c(0) = V(0)$. Using the full truncated Euler discretization scheme for the log-normal Heston model (see Section 4.1.1) results in the following for the fine level

$$\begin{aligned} \hat{I}^f(t + i\Delta t) = & \hat{I}(t + (i-1)\Delta t) \exp \left[(r_N - r_R - \frac{1}{2}\hat{V}(t + (i-1)\Delta t))\Delta t \right. \\ & \left. + \sqrt{\hat{V}(t + (i-1)\Delta t)\Delta t} N_i^I \right] \end{aligned} \quad (4.53)$$

$$\begin{aligned} \hat{V}^f(t + i\Delta t) = & \left[\hat{V}(t + (i-1)\Delta t) + \alpha(\bar{V} - (\hat{V}(t + (i-1)\Delta t))^+)\Delta t \right. \\ & \left. + \sigma_V \sqrt{(\hat{V}(t + (i-1)\Delta t))^+} \sqrt{\Delta t} N_i^V \right]^+ \end{aligned} \quad (4.54)$$

where $i = 1, \dots, M$, $(\cdot)^+ = \max\{\cdot, 0\}$ and N_i^I, N_j^V are standard normal variables which are correlated for $i = j$ and independent for $i \neq j$. For the coarse level,

$$\hat{I}^c(t + M\Delta t) = \hat{I}(t) \exp \left[(r_n - r_r - \frac{1}{2}\hat{V}(t))M\Delta t + \sqrt{\hat{V}(t)M\Delta t} N_I \right], \quad (4.55)$$

$$\hat{V}^c(t + M\Delta t) = \left[\hat{V}(t) + \alpha(\bar{V} - (\hat{V}(t))^+)M\Delta t + \sigma_V \sqrt{(\hat{V}(t))^+} \sqrt{M\Delta t} N_V \right]^+, \quad (4.56)$$

where $N_k := \frac{1}{\sqrt{M}}(N_1^k + N_2^k + \dots + N_M^k)$, $k \in \{I, V\}$ establishes the same Brownian path between the fine and coarse level.

As discussed in Section 4.1.1 this structure is particularly sensitive to the parameters of the Heston model. The unrealistic dynamics produced, especially when the Feller condition is not satisfied, can cause poor convergence in standard Monte Carlo. Even though the variance is reduced using MLMC, the bias from the discretization is still present and MLMC is not designed to improve this.

One way to avoid this is to use the exact simulation schemes presented previously. That is, by simulating the variance exactly and using the trapezoidal scheme to approximate the integrated variance process. This results in the following schemes for the coarse and the fine level in MLMC,

$$\begin{aligned} \hat{I}^f(t + i\Delta t) = & \hat{I}(t + (i-1)\Delta t) \exp \left[\left(r_n - r_r - \frac{\rho\alpha\bar{V}}{\sigma_V} \right) \Delta t \right. \\ & + \left(\frac{\rho\alpha}{\sigma_V} - \frac{1}{2} \right) \frac{\hat{V}(t + (i-1)\Delta t) + \hat{V}(t + i\Delta t)}{2} \Delta t \\ & + \frac{\rho}{\sigma_V} \left(\hat{V}(t + i\Delta t) - \hat{V}(t + (i-1)\Delta t) \right) \\ & \left. + \sqrt{1 - \rho^2} \sqrt{\frac{\hat{V}(t + (i-1)\Delta t) + \hat{V}(t + i\Delta t)}{2} \Delta t} N_i \right] \end{aligned} \quad (4.57)$$

$$\begin{aligned}
\hat{I}^c(t + M\Delta t) = & \hat{I}(t) \exp \left[\left(r_n - r_r - \frac{\rho\alpha\bar{V}}{\sigma_V} \right) M\Delta t \right. \\
& + \left(\frac{\rho\alpha}{\sigma_V} - \frac{1}{2} \right) \frac{\hat{V}(t + M\Delta t) + \hat{V}(t)}{2} M\Delta t \\
& + \frac{\rho}{\sigma_V} \left(\hat{V}(t + M\Delta t) - \hat{V}(t) \right) \\
& \left. + \sqrt{1 - \rho^2} \sqrt{\frac{\hat{V}(t + M\Delta t) + \hat{V}(t)}{2} M\Delta t N} \right]
\end{aligned} \tag{4.58}$$

Since the variance process is now simulated exactly, only a single Brownian path needs to be specified, $N := \frac{1}{\sqrt{M}}(N_1 + N_2 + \dots + N_M)$.

All the ingredients are now available to apply multilevel Monte Carlo to the Heston model. By specifying a simulation scheme for the inflation index and variance process, deciding how to deal with the integrated variance, and implementing the Brownian paths, MLMC can be compared with standard Monte Carlo simulation. Furthermore, MLMC can in theory be implemented with any of the exact schemes described in this chapter.

5 Numerical experiments

To test the simulation schemes, the pricing of European call options on the inflation index is considered for the Heston model. Since inflation caps using the Heston model can be priced analytically using the COS method, this is a standard test case on which to benchmark the implemented schemes. Firstly, the Euler, Euler-FT, NCI and QE schemes are compared. Based on the results from this, multilevel Monte Carlo is compared to standard Monte Carlo for the Heston model.

All the experiments are conducted in Python, the code can be accessed using the link below. Note that the implementation for the Heston Hull-White class is still a work in progress.

<https://github.com/Wobbuffet334/MLMC->

5.1 Valuing an inflation cap using the Heston model

To test the simulation schemes, the value of an inflation cap (see, Equation (2.5)) is estimated and compared with the exact value. For the parameter settings of the Heston model, Anderson [1] provides a test case designed for long-dated foreign exchange options. This parameter set is given in Table 1.

Table 1: Heston model parameters for long dated foreign exchange options.

T	r	$I(0)$	$V(0)$	α	\bar{V}	σ_V	ρ
10	0	100	0.04	0.5	0.04	1.0	-0.9

The value $\frac{2\alpha\bar{V}}{\sigma_V} = 0.04 \ll 1$ shows the Feller condition is not satisfied by a large margin. Having a high volatility of the variance process combined with a low mean-reversion means the variance process has a high probability of being pushed towards zero. This means for this parameter setting, some simulation schemes will have difficulty obtaining exact prices especially for large time steps.

For an inflation cap with maturity T and strike K , the exact price is denoted by P and the price calculated using Monte Carlo is denoted by \hat{P} and calculated using equation (4.1). Since the estimated values are in general not equal to the theoretically exact value, the difference between the estimation and the exact value, defined as the bias, is

$$b := P - \hat{P}. \quad (5.1)$$

For the different simulation schemes the bias is estimated for several values of the time step, strike $K = 100$ and 10^5 samples. The Exact-Gamma scheme refers to sampling using the noncentral chi-squared distribution using the built in function from the Scipy package, which samples directly from the gamma distribution. The results are provided in Table 2. Unsurprisingly the standard Euler method with the variance forced to zero when it becomes negative, exhibits terrible performance with a significant bias even at 4096 time steps. When the number of time steps is 1 the bias is more reasonable than for 4 time steps because the correction for

negative variances does not take place. For reasonable sizes of the time step, the exact schemes produce substantially lower biases than the Euler-FT scheme.

Another observation is the high computation times for the exact QE and NCI schemes. In the literature they are considered state of the art when it comes to combining speed and accuracy. However, in the implementation for this thesis, the python package for sampling from the noncentral chi-squared distribution is significantly faster. This is likely due to the code currently running the QE and NCI algorithms element-wise.

Table 2: Estimated inflation cap prices biases b for different simulation schemes and time steps. Number of samples was 10^5 . The bias is the left value, the standard error is given in the parenthesis, and the computation time (in seconds) is on the right hand side of the standard error.

No. of time steps	Euler		Euler-FT		Exact-Gamma		Exact-QE		Exact-NCI	
1	-11.323(±0.171)	0.01	-11.616(±0.173)	0.50	3.827(±0.046)	0.44	3.327(±0.046)	1.14	3.853(±0.046)	2.21
4	-502.933(±34.812)	0.03	-11.844(±0.172)	0.03	-0.308(±0.042)	0.063	-1.701(±0.041)	2.43	-0.313(±0.042)	8.42
16	-39.867(±0.393)	0.18	-4.491(±0.073)	0.21	-0.111(±0.042)	0.29	-0.488(±0.042)	9.64	-0.075(±0.042)	33.92
64	-13.420(±0.089)	0.70	-1.344(±0.0489)	0.77	0.055(±0.042)	1.22	0.044(±0.042)	39.54	0.009(±0.042)	134.02
256	-4.254(±0.055)	2.97	-0.303(±0.043)	3.08	-0.001(±0.042)	4.64	0.029(±0.042)	157.31		
1024	-1.147(±0.045)	13.73	-0.125(±0.042)	13.39	0.008(±0.042)	18.40	-0.060(±0.042)	633.51		
4096	-0.336(±0.0432)	212.27	0.025(±0.042)	218.26	0.019(±0.042)	190.85	0.044(±0.042)	2681.31		

5.2 MLMC vs standard Monte Carlo using the Heston model

Another set of parameters, used specifically for inflation caps is suggested by Singor et. al. [43] and shown in Table 3.

Table 3: Heston model parameters for short maturity inflation caps.

T	r	$I(0)$	$V(0)$	α	\bar{V}	σ_V	ρ
1	0.05	100	0.04	0.3	0.04	0.6	-0.7

This is again an extreme case where the Feller condition ($\frac{2\alpha\bar{V}}{\sigma_V} > 1$) is not satisfied. Applying the exact simulation method with the drift interpolation scheme described in Section 4.1.2, the MLMC mean and variance $\mathbb{V}(\hat{P}_\ell - \hat{P}_{\ell-1})$ are plotted in Figure 5.2 to observe the convergence to zero as the level ℓ increases. \hat{P}_ℓ is the approximated price of a cap on the inflation index with step size $M^{-\ell}$ with M set to 4. According to Theorem 1, the logarithm base M of the variance should approximately correspond to a slope of -1 . This is indeed the case verifying V_ℓ is approximately $O(h)$.

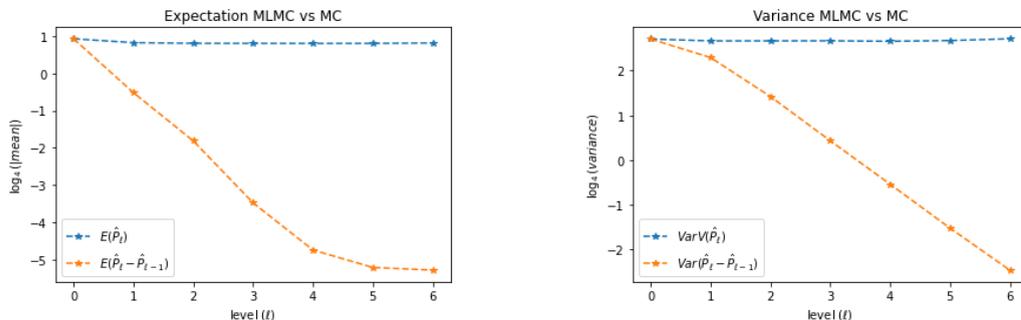


Figure 5.1: Comparison of the expectation and variance between the standard Monte Carlo estimator (blue) and the MLMC estimator (yellow) with $\epsilon = 0.005$.

The computational cost of MLMC is defined as the total number of time steps performed on all levels,

$$C_{MLMC} = N_0 + \sum_{\ell=1}^L N_{\ell}(M^{\ell} + M^{\ell-1}). \quad (5.2)$$

where the term $M^{\ell} + M^{\ell-1}$ reflects that each sample at level $\ell > 0$ requires one fine path with M^{ℓ} time steps, and one coarse path with $M^{\ell-1}$ time steps. For comparison purposes the cost to carry out one time step is not required. The computational cost is compared to the cost of the standard Monte Carlo method which is calculated as

$$C_{MC} = \sum_{\ell=0}^L N_{\ell}^{MC} M^{\ell} \quad (5.3)$$

where $N_{\ell}^{MC} = 2\epsilon^{-2}\mathbb{V}(P_{\ell})$ so that the variance of this estimator is also $\frac{1}{2}\epsilon^2$. Giles mentions the summation over the levels corresponds to an application of the standard Monte Carlo method on each level to enable the estimation of the error in order to apply the same termination criterion as the multilevel method. It is also possible to define $C_{MC} = N_L^{MC} M^L$ since this will be of a similar magnitude.

The number of samples and computation costs for the parameters mentioned previously are shown in Figure 5.2. The plot for the computational costs is ϵ^2 against ϵ because it is expected that $\epsilon^2 C$ is only weakly dependent on ϵ for the multilevel. This is also visible in the figure. For the standard Monte Carlo method, theory suggests that $\epsilon^2 C$ is proportional to the number of time steps on the finest level, which is roughly proportional to $1/\epsilon$ due to the weak convergence. This is also visible in the figure because the number of levels for $\epsilon = 0.005$ and $\epsilon = 0.01$ is higher which translates to the computation costs figure.

The computational savings from using MLMC compared to standard Monte Carlo are significant, ranging from a factor of 10 to 50.

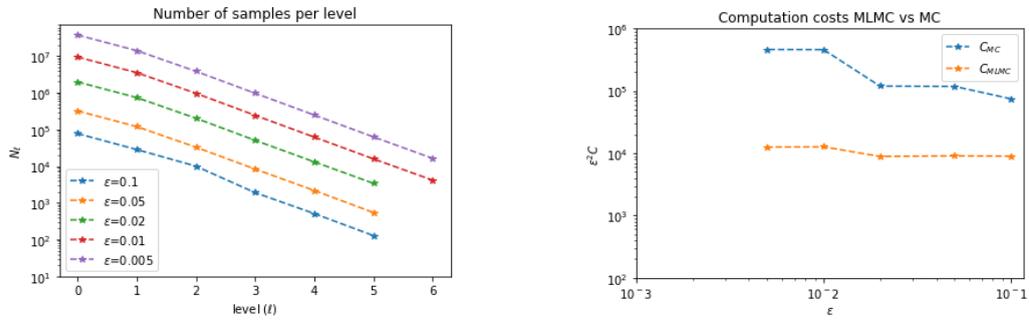


Figure 5.2: Comparison of the optimal number of samples per level and computation costs for different values of ϵ .

To test the longer maturity performance, MLMC is applied to the parameter setting in Table 1 where the maturity is $T = 10$ years. Additionally, as a result from the numerical results in the previous section, the time steps in the first level is set to 16 to ensure a reasonable initial estimate in the MLMC algorithm. In Table 5.2 the difference between the estimated price and the exact price is compared against ϵ . The resulting number of samples and computation costs are given in Figure 5.2. In this case the savings in computation time range between a factor 3 and 10.

Table 4: Difference between MLMC prices and the exact price per MSE (ϵ). Exact value is 13.085.

ϵ	0.1	0.05	0.02	0.01	0.005
b	-0.0431	-0.0021	0.0078	-0.0088	-0.0040

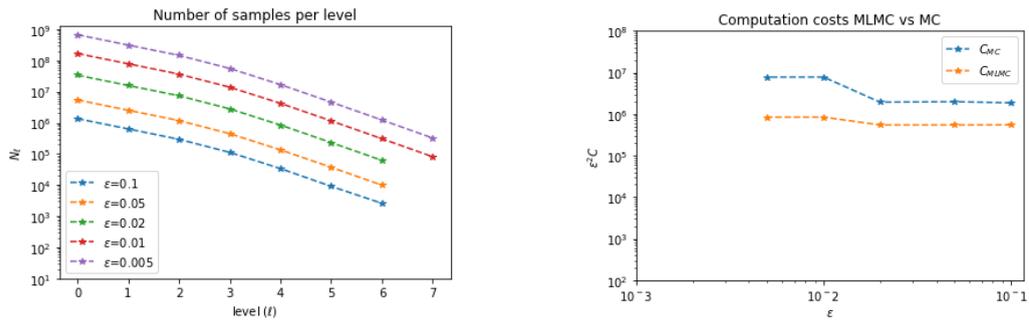


Figure 5.3: Comparison of the optimal number of samples per level and computation costs for different values of ϵ , with the parameter setting from Table 1

6 Conclusion & Discussion

In this chapter the conclusions of the research carried out in this thesis are discussed. In the last section some final thoughts and recommendations for future research are discussed.

6.1 Conclusion

In this thesis the models and methods for pricing inflation-indexed derivatives were analyzed. Particular attention was given to the challenges associated with finding analytical solutions and the simulation techniques which are available as an alternative. This thesis is particularly interesting for financial institutions who deal with inflation risk and require risk-neutral valuation of options on inflation.

To start with, the shortcomings of the well known Black-Scholes and Jarrow-Yildirim frameworks where the volatility is deterministic are discussed. Addressing this, the Heston model models the volatility of the inflation index as a Cox-Ingersoll-Ross process while keeping the interest rates constant. The COS method provides a fast method for calculating the exact solution of the price of a call option to the Heston model using its characteristic function. Several extensions of the Heston model involving stochastic interest rates, such as the Hull-White model, are also discussed.

Considering the complicated pricing and correlation structure using Heston Hull-White, Monte Carlo simulation offers an alternative approach. The distributional properties of the variance process leads to several exact simulation schemes which can sample from the noncentral chi-squared distribution. Comparing the Euler schemes with the exact simulation schemes, it is apparent that exact simulation using a drift interpolation for the integrated variance process is the most accurate method.

One of the main obstacles to using Monte Carlo is the high computational cost associated with its low convergence rate. Multilevel Monte Carlo (MLMC) uses a sequence of estimators with increasing step-size to optimize the computational cost around a large number of samples at a large time step, and a small number of samples at a small time step. Applying the MLMC method to the Heston model with constant interest rates simulating the variance exactly results in a marked decrease in computation times. Compared to the standard Monte Carlo method a decrease up to a factor of 50 was observed for a set of parameters where the Feller condition was not satisfied. This makes MLMC a powerful tool for the path-dependent simulation of inflation indexed-derivatives. However, it is reliant on the underlying simulation scheme used

6.2 Discussion

Considering the complicated approximations required to derive analytical solutions, Monte Carlo simulation is an appealing alternative. In Section 4.1 several simulation schemes are discussed for the Heston model. Many papers have been devoted to the exact simulation of the variance process using its distributional properties. However, little discussion exists about the use of Monte Carlo for inflation-indexed derivatives.

An open question remains whether simulation techniques can bridge the gap in complexity between the complex multi factor models and the use of these models by financial institutions. Currently the inflation-indexed products in the market are limited in structure. American, Asian, lookback, barrier options, to name a few, are not publicly available. Finding an exact pricing formula for such derivatives using the Heston Hull-White model in the context of inflation may prove challenging if not impossible. The same can be said when using other model dynamics, like including jumps or seasonality in inflation, or a different model for the interest rates. Monte Carlo simulation however, is much more flexible and can still be used in such cases.

Aside from the different payoff structures, it would be interesting to see the simulation techniques found in this thesis applied to the LIBOR market models. Furthermore, a calibration to market data more recent than 12 years could not be found, a brief summary of calibrations performed in the literature is provided in Appendix A.2.

The next step for this research would be applying MLMC to the Heston Hull-White model and comparing to the analytical solution found by using the characteristic function in Section 3.3.2. Currently the implemented algorithm does not converge as expected when applying the results from Section 4.2 and Appendix A.3. Another point of interest would be combining antithetic sampling and MLMC to further reduce the variance of the estimator.

References

1. Andersen, L. Efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance* **11** (Mar. 2008).
2. Beveridge, C., Denson, N. & Joshi, M. Comparing Discretisations of the LIBOR Market Model in the Spot Measure. *Australian Actuarial Journal* **15**, 231–253. <https://search.informit.org/doi/10.3316/ielapa.940319561883556> (2009).
3. Björk, T. *Arbitrage Theory in Continuous Time* ISBN: 9780198851615 (Oxford University Press, Dec. 2019).
4. Black, F. & Scholes, M. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* **81**, 637–654. ISSN: 00223808, 1537534X. <http://www.jstor.org/stable/1831029> (2024) (1973).
5. Bormetti, G., Callegaro, G., Livieri, G. & Pallavicini, A. A Backward Monte Carlo Approach to Exotic Option Pricing. *SSRN Electronic Journal* (Nov. 2015).
6. Brace, A., Gatarek, D. & Musiela, M. The Market Model of Interest Rate Dynamics. *Mathematical Finance* **7**, 127–155 (1997).
7. Brigo, D. & Mercurio, F. *Interest Rate Models — Theory and Practice: With Smile, Inflation and Credit* ISBN: 978-3-540-22149-4 (Springer, Jan. 2006).
8. Britannica, T. E. o. E. Consumer Price Index Summary. *Encyclopedia Britannica* (May 2020).
9. Broadie, M. & Kaya, Ö. Exact Simulation of Stochastic Volatility and Other Affine Jump Diffusion Processes. *Operations Research* **54**, 217–231 (Apr. 2006).
10. Campbell, J. Y., Shiller, R. & Viceira, L. M. Understanding Inflation-Indexed Bond Markets. *National Bureau of Economic Research* (2009).
11. Carr, P., Stanley, M. & Madan, D. Option Valuation Using the Fast Fourier Transform. *J. Comput. Finance* **2** (Mar. 2001).
12. Chen, B., Grzelak, L. & Oosterlee, C. Calibration and Monte Carlo Pricing of the SABR-Hull-White Model for Long-Maturity Equity Derivatives. *The Journal of Computational Finance* (Feb. 2011).
13. Chval, D. in *Regulation of Finance and Accounting* (ed Procházka, D.) 85–94 (Springer, Dec. 2022).
14. Cliffe, K., Giles, M. B., Scheichl, R. & Teckentrup, A. L. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science* **14** (Aug. 2011).
15. Cox, J., Ingersoll, J. & Ross, S. A Theory of the Term Structure of Interest Rates. *Econometrica* **53**, 385–407 (Feb. 1985).
16. David, H., Robert, J. & Andrew, M. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica* **60**, 77–105. (2024) (1992).

17. Deacon, M., Derry, A. & Mirfendereski, D. *Inflation-indexed securities* 2nd ed. (Wiley, 2004).
18. Dodgson, M. J. W. & Kainth, D. *Inflation-Linked Derivatives* in (2006).
19. Fang, F. & Oosterlee, C. W. A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing* **31**, 826–848 (2009).
20. Farshid, J. LIBOR and swap market models and measures (*). *Finance and Stochastics* **1**, 293–330 (1997).
21. Fisher, I. *The Theory of Interest*. MacMillan Press Ltd. (1930).
22. Geman, H., El Karoui, N. & Rochet, J. Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied Probability* **32**, 443–458 (1995).
23. Giles, M. Multilevel Monte Carlo Path Simulation. *Operations Research* **56**, 607–617 (June 2008).
24. Glasserman, P. *Monte Carlo methods in financial engineering* (Springer, New York, 2004).
25. Grzelak, L. & Oosterlee, C. On Cross-Currency Models with Stochastic Volatility and Correlated Interest Rates. *Applied Mathematical Finance* **19** (June 2010).
26. Hagan, P., Kumar, D., Lesniewski, A. & Woodward, D. Managing Smile Risk. *Wilmott Magazine* **1**, 84–108 (Jan. 2002).
27. Heston, S. L. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies* **6**, 327–343 (Apr. 2015).
28. Hout, K., Bierkens, J., Ploeg, A. & Panhuis, J. A semi closed-form analytic pricing formula for call options in a hybrid Heston–Hull–White model. *Proceedings of the 58th European Study Group Mathematics with Industry* (Apr. 2007).
29. Jarrow, R. & Turnbull, S. A Unified Approach for Pricing Contingent Claims on Multiple Term Structures. *Review of Quantitative Finance and Accounting* **10**, 5–19 (Jan. 1998).
30. Jarrow, R. & Yildirim, Y. Pricing Treasury Inflation Protected Securities and Related Derivative Securities using an HJM Model. *Journal of Financial and Quantitative Analysis* **38**, 337–359 (2003).
31. Korn, R. & Kruse, S. A simple model to value inflation-linked financial products (in German). *Blätter der DGVM* **26**, 351–367 (2004).
32. Kruse, S. On the Pricing of Inflation-Indexed Caps. *European Actuarial Journal* **1**, 379–393 (2011).
33. Kruse, S. & Nögel, U. On the pricing of forward starting options in Heston’s model on stochastic volatility. *Finance and Stochastics* **9**, 233–250 (2005).

34. Lord, R., Koekkoek, R. & Van Dijk, D. *A comparison of biased simulation schemes for stochastic volatility models* Discussion Paper TI 2006-046/4 (Tinbergen Institute, Erasmus University, Rotterdam, Netherlands, 2008).
35. Mercurio, F. Pricing Inflation-Indexed Derivatives. *Quantitative Finance* **5**, 289–302.
36. Mercurio, F. & Moreni, N. Inflation with a smile. *Risk* **19**, 70–75 (Jan. 2006).
37. Mercurio F. and Moreni, N. A Multi-Factor SABR Model for Forward Inflation Rates. *SSRN Electronic Journal* (Jan. 2009).
38. Michael, S. Introduction to Stochastic Finance with Market Examples, 2nd ed. *The American Statistician* **78**, 129–130 (2024).
39. Mickel, A. & Neuenkirch, A. *The weak convergence order of two Euler-type discretization schemes for the log-Heston model* 2022.
40. OECD. *Pension Markets in Focus 2023* 56 (OECD Publishing, 2023).
41. Pelsser, A. & Haastrecht, A. Efficient, Almost Exact Simulation of the Heston Stochastic Volatility Model. *International Journal of Theoretical and Applied Finance (IJTAF)* **13**, 1–43 (Feb. 2010).
42. Pham, H. *Feynman-Kac representation of fully nonlinear PDEs and applications* 2014.
43. Singor, S., Grzelak, L., van Bragt, D. & Oosterlee, C. Pricing inflation products with stochastic volatility and stochastic interest rates. *Insurance: Mathematics and Economics* **52**, 286–299. ISSN: 0167-6687 (2013).
44. Zheng, C. Multilevel Monte Carlo simulation for the Heston stochastic volatility model. *Advances in Computational Mathematics* **49** (Nov. 2023).

A Appendix

A.1 Proof of Theorem 1

Proof. Recall that $h_\ell = \frac{1}{M}h_{\ell-1}$ for all $\ell = 1, \dots, L$ for some $M \in \mathbb{N} \setminus \{1\}$. Without loss of generality, assume $h_0 = 1$. If this is not the case simply scale the constants c_1, c_2 and c_3 accordingly.

Using the notation $\lceil x \rceil$ to denote rounding up to the nearest integer, start with choosing L to be

$$L = \lceil \alpha^{-1} \log_M(\sqrt{2}c_1\epsilon^{-1}) \rceil < \alpha^{-1} \log_M(\sqrt{2}c_1\epsilon^{-1}) + 1 \quad (\text{A.1})$$

so that

$$M^{-\alpha} \frac{\epsilon}{\sqrt{2}} < c_1 M^{-\alpha L} \leq \frac{\epsilon}{\sqrt{2}} \quad (\text{A.2})$$

and hence due to assumptions 1 and 2,

$$\left(\mathbb{E}[\hat{P}_\ell - P] \right)^2 \leq \frac{1}{2} \epsilon^2 \quad (\text{A.3})$$

This $\frac{1}{2}\epsilon^2$ upper bound on the square of the bias error together with the $\frac{1}{2}\epsilon^2$ upper bound on the variance of the estimator (proved later) produces an ϵ^2 upper bound on the mean squared error estimator.

Using the left side of equation (A.2) and the standard result for the geometric series, the following inequality is obtained which will be used later

$$\sum_{\ell=0}^L h_\ell^{-\gamma} = h_L^{-\gamma} \sum_{\ell=0}^L M^{-\ell} < \frac{M}{M-1} h_L^{-\gamma} \leq \frac{M^2}{M-1} (\sqrt{2}c_1)^{\gamma/\alpha} \epsilon^{-\gamma/\alpha}. \quad (\text{A.4})$$

Now the different values of β are considered.

(a) If $\beta = \gamma$, set $N_\ell = \lceil 2\epsilon^{-2}(L+1)c_2h_\ell \rceil$ so that

$$\mathbb{V}(Y) = \sum_{\ell=0}^L \mathbb{V}(Y_\ell) \leq \sum_{\ell=0}^L c_2 N_\ell^{-1} h_\ell \leq \frac{1}{2} \epsilon^2, \quad (\text{A.5})$$

which is the required upper bound on the variance of the estimator. To bound the computational complexity C , begin with an upper bound on L given by

$$L \leq \frac{\log \epsilon^{-1}}{\alpha \log M} + \frac{\log(\sqrt{2}c_1 T^\alpha)}{\alpha \log M} + 1 \quad (\text{A.6})$$

given that $1 < \log \epsilon^{-1}$ for $\epsilon < e^{-1}$, it follows that

$$L + 1 \leq c_5 \log \epsilon^{-1}, \quad (\text{A.7})$$

where

$$c_5 = \frac{1}{\alpha \log M} + \max\left\{0, \frac{\log(\sqrt{2}c_1 T^\alpha)}{\alpha \log M} + 2\right\}. \quad (\text{A.8})$$

Hence, the computational complexity is bounded by

$$C \leq c_3 \sum_{\ell=0}^L N_\ell h_\ell^{-\gamma} \leq c_3 \left(2\epsilon^{-2} c_2 (L+1)^2 + \sum_{\ell=0}^L h_\ell^{-\gamma} \right). \quad (\text{A.9})$$

Using the upper bound for $L+1$ and inequality (A.4) it follows that

$$C \leq c_4 \epsilon^{-2} (\log \epsilon)^2, \quad (\text{A.10})$$

where

$$c_4 = 2c_2 c_3 c_5^2 + c_3 \frac{M^2}{M-1} (\sqrt{2}c_1)^{\gamma/\alpha} \quad (\text{A.11})$$

(b) For $\beta > \gamma$, setting

$$N_\ell = \lceil 2\epsilon^{-2} T^{\beta-\gamma} c_2 (1 - M^{-(\beta-\gamma)/2})^{-1} h_\ell^{(\beta+\gamma)/2} \rceil \quad (\text{A.12})$$

using the standard result for geometric series,

$$\sum_{\ell=0}^L h_\ell^{(\beta-\gamma)/2} = T^{(\beta-\gamma)/2} \sum_{\ell=0}^L (M^{-(\beta-\gamma)/2})^\ell < T^{(\beta-\gamma)/2} (1 - M^{-(\beta-\gamma)/2})^{-1}, \quad (\text{A.13})$$

and hence a $\frac{1}{2}\epsilon^2$ upper bound on the variance of the estimator is obtained. Using the N_ℓ upper bound, the computational complexity is bounded by

$$C \leq c_3 \left(2\epsilon^{-2} c_2 T^{(\beta-\gamma)/2} (1 - M^{-(\beta-\gamma)/2})^{-1} \sum_{\ell=0}^L h_\ell^{(\beta-\gamma)/2} + \sum_{\ell=0}^L h_\ell^{-\gamma} \right). \quad (\text{A.14})$$

Using inequalities (A.13) and (A.4) gives $C \leq c_4 \epsilon^{-2}$, where

$$c_4 = 2c_2 c_3 T^{\beta-\gamma} (1 - M^{-(\beta-\gamma)/2})^{-2} + c_3 \frac{M^2}{M-1} (\sqrt{2}c_1)^{\gamma/\alpha} \quad (\text{A.15})$$

(c) for $\beta < \gamma$, setting

$$N_\ell = \lceil 2c_2 \epsilon^{-2} h_L^{-(\gamma-\beta)/2} (1 - M^{-(\gamma-\beta)/2}) h_\ell^{(\beta+\gamma)/2} \rceil. \quad (\text{A.16})$$

Then

$$\sum_{\ell=0}^L \mathbb{V}(Y) < \frac{1}{2}\epsilon^2 h_L^{(\gamma-\beta)/2} (1 - M^{-(\gamma-\beta)/2}) \sum_{\ell=0}^L h_\ell^{-(\gamma-\beta)/2}. \quad (\text{A.17})$$

Since

$$\sum_{\ell=0}^L h_\ell^{-(\gamma-\beta)/2} = h_L^{-(\gamma-\beta)/2} \sum_{\ell=0}^L (M^{-(\gamma-\beta)/2})^\ell < h_L^{-(\gamma-\beta)/2} (1 - M^{-(\gamma-\beta)/2})^{-1}, \quad (\text{A.18})$$

again an upper bound $\frac{1}{2}\epsilon^2$ is obtained for the variance of the estimator.

Using the upper bound for N_ℓ , the computational complexity is bounded by

$$C \leq c_3 \left(2c_2 \epsilon^{-2} h_L^{-(\gamma-\beta)/2} (1 - M^{-(\gamma-\beta)/2})^{-1} \right) \sum_{\ell=0}^L h_\ell^{-(\gamma-\beta)/2} + \sum_{\ell=0}^L h_\ell^{-\gamma}. \quad (\text{A.19})$$

Using inequality (A.18) gives

$$h_L^{-(\gamma-\beta)} < (\sqrt{2}c_1)^{(\gamma-\beta)/\alpha} M^{\gamma-\beta} \epsilon^{-(\gamma-\beta)/\alpha}. \quad (\text{A.20})$$

Combining the inequalities with inequality (A.4) results in

$$C \leq c_4 \epsilon^{-2-(\gamma-\beta)/\alpha}, \quad (\text{A.21})$$

where

$$c_4 = 2c_2 c_3 (\sqrt{2}c_1)^{(\gamma-\beta)/\alpha} M^{\gamma-\beta} (1 - M^{-(\gamma-\beta)/2})^{-2} + c_3 \frac{M^2}{M-1} (\sqrt{2}c_1)^{\gamma/\alpha} \quad (\text{A.22})$$

□

A.2 Calibrations performed in Literature

As part of any model the pricing methods described, contain unknown parameters which must be calibrated using historical data and/or quoted market data. Calibration is an optimization procedure that estimates the model parameters such that the market prices are replicated by the model as closely as possible. This generally consists of a least squares estimation like $\min_{\Omega} \|C - \hat{C}\|$ where C is the market price, \hat{C} the model price, $\|\cdot\|$ some norm, and Ω the set of parameters. In this section some of the performed calibrations are discussed.

For the LIBOR market model in Section 3.3.3, calibration is performed on caplet prices, which are stripped from the cap prices by taking differences between consecutive maturities. The time 0 value of a T_i -forward CPI where $T_i = i$ years, can be obtained from the market quote $S(T_i)$ of the corresponding zero-coupon inflation-indexed swap by applying the relation $\mathcal{I}_i(0) = I(0)(1 + S(T_i))^i$. Cap prices for different strikes and maturities are provided in Table 1. The discount factors, zero coupon swap rates and implied forward CPI's are quoted in Table 1.

The model parameters to be calibrated are for maturity T_M :

- Volatility parameters: $\sigma_V, \alpha, \theta, V(0)$,
- Forward CPI's volatility coefficients: $\sigma_i^I, i = 1, \dots, M$,
- Correlations between consecutive forward CPI's: $\rho_{i-1,i}^I, i = 2, \dots, M$,
- Correlations between forward CPI's and the volatility $\rho_i^{I,V}, i = 1, \dots, M$

To reduce the degrees of freedom the correlations between forward CPI's are parameterized by a decreasing growth function.

$$\rho_{i,i-1}^I = 1 - (1 - \rho_0) e^{-\lambda T_{i-2}}, \quad i = 2, \dots, M \quad (\text{A.23})$$

Table 5: Inflation-indexed Cap prices (in bps) for different strikes and maturities for November 3, 2004 in the USD market. Source: ICAP. Taken from [36].

T_i	K				
	-1.5%	-1.0%	-0.5%	0%	0.5%
1y	416.7	368.2	319.8	271.8	224.3
2y	822.4	727.6	633.3	540.1	448.7
3y	1212.7	1074.0	936.5	800.7	667.8
4y	1588.4	1408.0	1229.2	1052.8	880.5
5y	1952.9	1732.8	1514.6	1299.4	1089.1
6y	2288.2	2030.7	1775.6	1524.1	1278.3
7y	2612.4	2319.7	2029.8	1744.2	1465.2
8y	2911.5	2585.9	2263.6	1946.2	1636.5
9y	3197.2	2840.7	2488.0	2140.9	1802.5
10y	3467.7	3082.3	2701.1	2326.3	1961.1

Table 6: USD discount factors, Zero Coupon swap rates, and implied forward CPI's. November 3, 2004. Source: ICAP. Taken from [36].

T_i	$P(0, T_i)$	ZC rates	$\mathcal{I}_i(0)$
1y	0.97701	2.111%	194.94
2y	0.94982	2.188%	199.35
3y	0.91835	2.240%	204.03
4y	0.88433	2.278%	208.91
5y	0.84862	2.293%	213.82
6y	0.81179	2.300%	218.82
7y	0.77460	2.310%	224.00
8y	0.73785	2.320%	229.36
9y	0.70218	2.325%	234.78
10y	0.66773	2.335%	240.48

where $\lambda > 0$. As a result a total of $2M + 5$ parameters are calibrated by minimizing the sum of squared differences. For short maturity and close to zero strike caplets the modelled prices coincide very well with market prices from 2004, see Figure A.1. However, increasing the maturity to beyond 3 years reduces the fitting quality significantly which is highly undesirable.

In the case of a different volatility process for each forward CPI as in Section 3.3.4 sequences of caplets/floorlets with maturities of 1 year to 15 years were again derived from cap/floor quotes. Bid/ask spreads for caps/floors on the HICP excluding tobacco index from September 2008 are quoted in Table 3. The large bid/ask spreads indicate a lack of liquidity for the inspected derivatives. For maturities where caps were not available the cap prices for known maturities were inverted to find the corresponding flat implied volatilities. The flat volatilities are then interpolated to find cap prices for the missing maturities. This is a common technique in the interest rate market because calibrating to YoY caplets over YoY caps decreases computation time significantly since the iterative calculation is avoided.

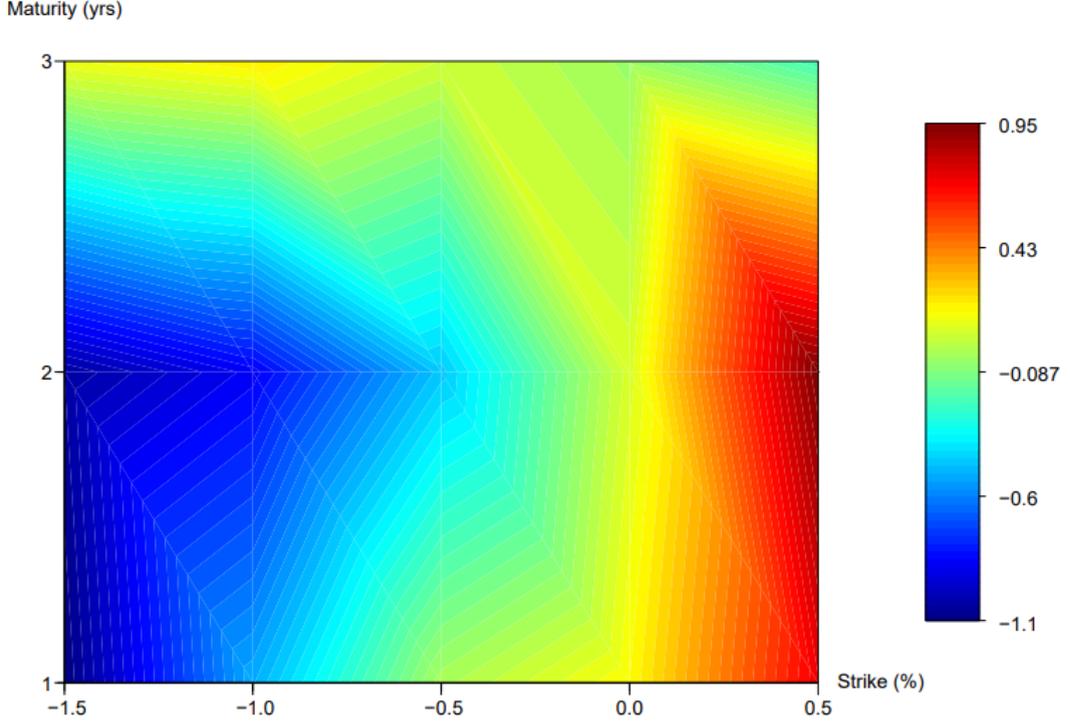


Figure A.1: Percentage errors after calibration of the LIBOR market model with Heston dynamics. Caplet maturities are 1,2 and 3 years and strikes range from -1.5% to 0.5% . Taken from [36].

From here two calibration tests are performed with the zero-drift approximation. Firstly the correlations between inflation rates and between the inflation rates and the interest rates, are assumed to be zero. A visually accurate fit is obtained, however, it is noted that the assumption made is unrealistic. For the second test the calibrations are assumed to follow a complicated structure involving idiosyncratic coupling coefficients. How this structure works is stated in the appendix without references making it unclear what the method entails. The resulting correlation assumptions are

$$\rho_{i,j}^W = e^{-\lambda|i-j|},$$

$$\rho_{i,j}^{F,W} = \frac{\sum_{k=1}^1 5c_{i,k}\rho_{k,j}^W}{\sqrt{1 + \sum_{k,k'=1}^1 5c_{i,k}c_{i,k'}\rho_{k,k}^W}},$$

with

$$c_{i,j} = ce^{-\lambda_c|i-j-1|},$$

where $\lambda, \lambda_c, c > 0$ are freely chosen.

The calibration results are shown in Figure A.2 where the model implied volatilities are plotted together with the market data. While the volatilities are well modelled, the lack of error analysis, reliably defined correlation structure, and liquid market data makes the results hard to verify.

Table 7: HICP-exT cap (C) and floor (F) bid/ask quotes (in bps) for different maturities T in years and strikes K . Taken from [37].

option type	F	F	F	F	C	C	C	C
T / K	-1%	0%	1%	2%	2.5%	3%	4%	5%
3	3/6	7/14	26/40	98/118	109/128	67/84	26/37	11/18
5	5/17	14/34	51/73	165/189	218/238	143/161	62/79	29/45
7	14/22	39/49	73/104	220/256	321/348	212/241	97/119	49/67
10	23/38	45/78	102/152	293/351	467/510	315/362	152/189	82/114
15	40/70	72/130	150/233	405/499	676/746	465/543	237/299	136/191
20	56/100	96/181	191/306	499/629	834/930	578/684	302/389	179/253
30	82/154	136/268	256/432	650/841	1047/1186	730/885	390/517	238/348

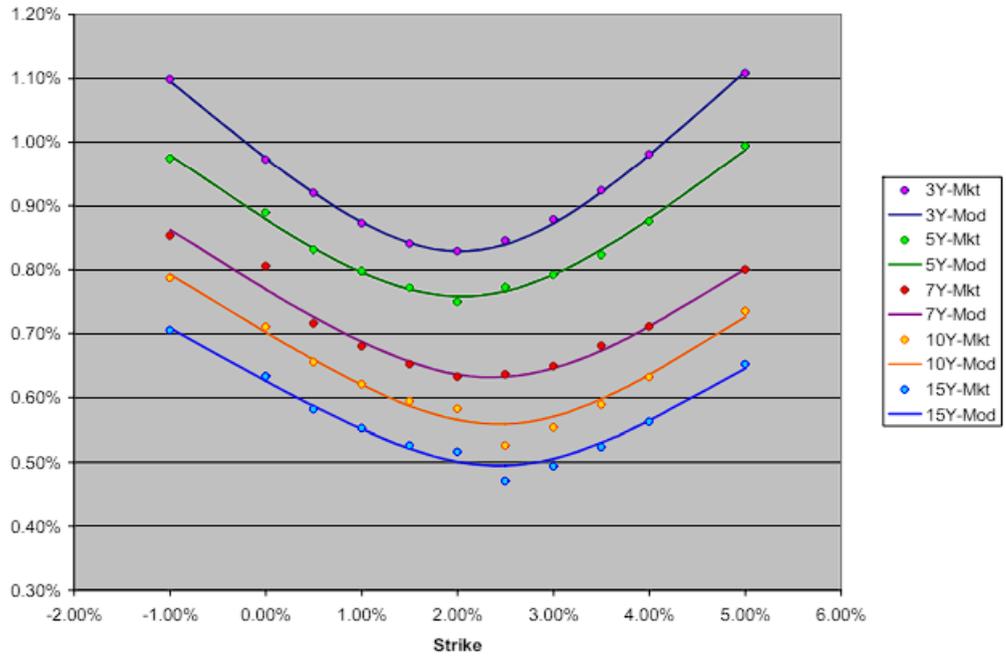


Figure A.2: Market and model implied volatilities for caplets/floorlets maturing in 3, 5, 7, 10, 15 years, correlated case. Taken from [37]

Lastly the Heston and Hull-White dynamics in Section 3.3.2 are considered. A calibration procedure is proposed by Singor et al. [43] that is more involved than just a least squares estimation. In their case with the market data split in pairs (\bar{T}, \bar{K}) where \bar{T} is the option maturity and \bar{K} the strike level, and using the Euclidean norm, the function to be minimized is

$$\min_{\Omega} \left[\left(\sum_{i=1}^m \sum_{k=1}^n |C(\bar{T}_i, \bar{K}) - \hat{C}(\bar{T}_j, \bar{K}_j)|^2 \right)^{\frac{1}{2}} \right]. \quad (\text{A.24})$$

The authors note that the difference between market and model implied volatilities could also be minimized. However this requires an additional numerical inversion with every iteration step which is undesirable.

Equation (A.24) is solved iteratively by sampling random starting points and using the local minimization Levenberg-Marquardt least-squares algorithm. This is repeated until the best local minimum is found. It is worth mentioning that the speed of calibration is not analyzed in any of the above examples despite the fast valuation of options being a topic of interest.

First the correlation parameters between the inflation and nominal/real interest rates $\rho_{I,n}$, $\rho_{I,r}$ and $\rho_{n,r}$ are estimated using historical data. Considering the past rates are known exactly this seems only natural, however this method is not applied in the previous calibration examples where some or all correlation parameters are parameterized or assumed to be zero. The correlations between the interest rates and the variance process $\rho_{r,V}$, and $\rho_{n,V}$ are derived with a conditional sampling method, while the correlation between the inflation and the variance process $\rho_{I,V}$ is the only correlation parameter estimated in the calibration process. Conditional sampling in the literature is stated as restricting samples to a specific condition which makes the method for obtaining these correlations quite vague.

The authors utilize another unique step is calibrating the mean-reversion and volatility parameters in the one-factor Hull-White interest rate model to options on the interest rate. Conditional on the estimated interest rate parameters, the inflation model is calibrated to inflation indexed caplets/floorlets. Since interest rate option data are much more liquid than inflation option data, this should be standard practice when dealing with interest rate models.

Calibration is essential for accurately pricing inflation-indexed derivatives. This process, as demonstrated through the various models, highlights significant challenges such as fitting long-term maturities and handling complex correlation structures. Different procedures were used which often lacked error analysis and were not compared to other methods. This emphasizes the need for refinement of the calibration procedure to increase model robustness and reliability for when companies and investors are interested in using them.

A.3 MLMC for the Heston Hull-White model

In the case of Hull-White dynamics for the interest rates the integrated interest is no longer constant. That is, the integrated interest rate terms found in the inflation index dynamics and payoff functions are not trivial. Using the trapezoidal rule in both cases is the straightforward approach. Because of the path dependent structure, the discount factor can be approximated with the composite trapezoidal rule because the nominal interest rate is known at every time step,

$$\int_0^T r_n(u)du \approx \frac{T}{n} \left(\frac{r_n(0)}{2} + \sum_{k=1}^{N_{steps}-1} \left(\hat{r} \left(k \frac{T}{n} \right) \right) + \frac{\hat{r}_n(T)}{2} \right). \quad (\text{A.25})$$

For the inflation index, approximating the integrated interest term with the trape-

zoidal rule results in:

$$\begin{aligned}
\hat{I}(t + \Delta t) = \hat{I}(t) \exp & \left(\int_t^{t+\Delta t} (\hat{r}_N(s) - \hat{r}_R(s)) ds - \frac{1}{2} \int_t^{t+\Delta t} \hat{V}(s) ds \right. \\
& + \frac{\rho}{\sigma_V} \left(\hat{V}(t + \Delta t) - \hat{V}(t) - \alpha \bar{V} \Delta t + \alpha \int_t^{t+\Delta t} \hat{V}(s) ds \right) \\
& \left. + \sqrt{1 - \rho^2} \sqrt{\int_t^{t+\Delta t} \hat{V}(s) ds} N \right).
\end{aligned}
\tag{A.26}$$

With these approximations the use of MLMC is similar to the Heston case with constant interest rates. The only additional requirement is when simulating the interest rates exactly, the Brownian path between the coarse level and the fine level must be established in the same manner as for the inflation index.