MSc thesis Applied Mathematics

A Fast Option Valuation Method in the Stochastic Multi-state Economic Model: CP2022

Vlad Popa

2025

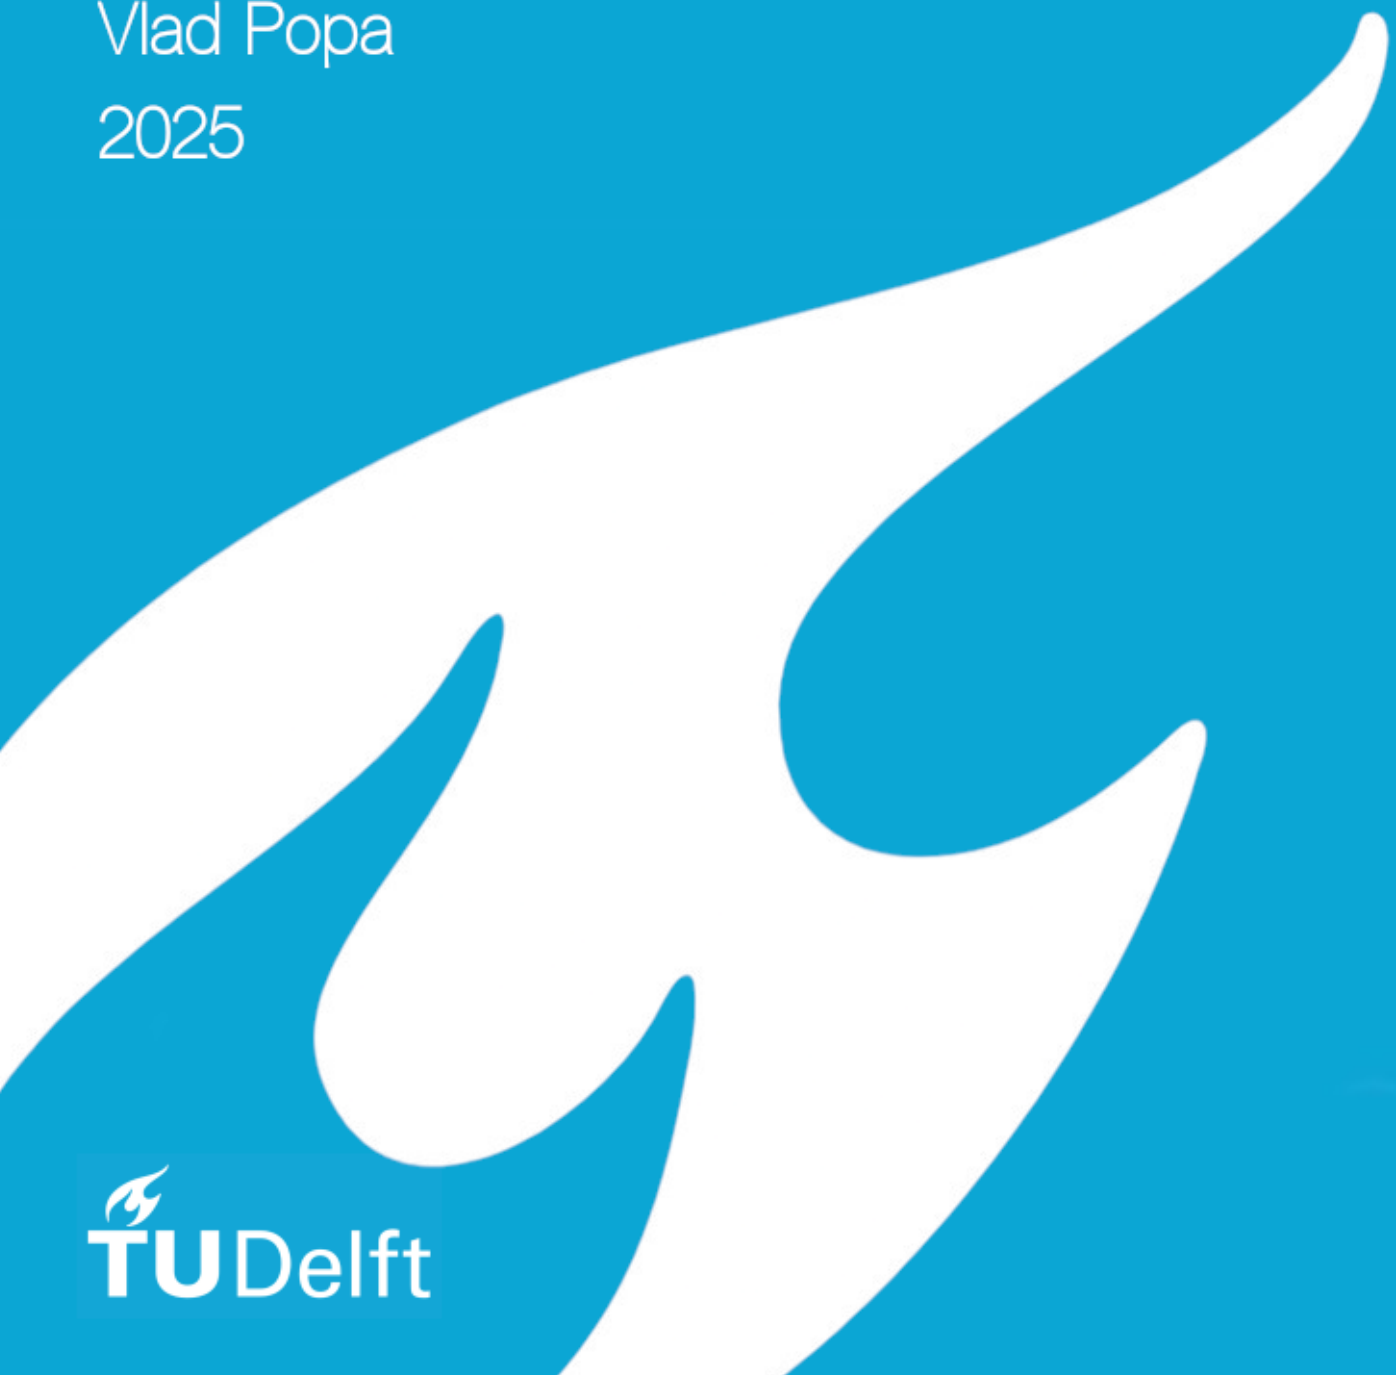

TU Delft

This page was intentionally left blank.

A FAST OPTION VALUATION METHOD IN THE STOCHASTIC MULTI-STATE ECONOMIC MODEL: CP2022

A thesis in collaboration with De Nederlandsche Bank submitted to the Delft University of Technology in fulfillment of the requirements for the degree of

Master of Applied Mathematics

by

Vlad Popa

March 17, 2025

Student number:	5113628
Supervisor:	Dr. F. Fang
Co-reader 1:	Prof. Dr. ir. C. Vuik
Co-reader 2:	Dr. J.H.M. Anderluh
Daily Supervisor:	S.M. Bye

TU Delft

DeNederlandscheBank
EUROSYSTEEM

The contents of this thesis is the responsibility of the author only and may not reflect the views of De Nederlandsche Bank.

ABSTRACT

Dutch pension funds are under the supervision by De Nederlandsche Bank (DNB) and they must adhere to the Financial Assessment Framework (FTK), which outlines the methods for calculating liabilities, buffer reserves, and risk factors. As part of the FTK, a feasibility test must be performed, which is a scenario-based analysis of pension funds' investment strategies and their pension policies based on economic scenario sets projected 100 years into future. Every quarter, DNB publishes (De Nederlandsche Bank [2022]) these economic scenarios ($\mathbb{P}$-scenarios) and equivalent risk-neutral scenarios ($\mathbb{Q}$-scenarios) are used to calculate the effects of the pension fund reform in the Netherlands. These scenarios are generated by a stochastic model, which is revised every 5 years by a commission appointed by the Dutch government, called the Commission Parameters. The model currently in use is based on Commission Parameters 2022 and is referred to as CP2022 (Commission Parameters [2022b]).

The CP2022 model is an affine and arbitrage-free model framework for correlated interest rate, inflation rate, stock index and consumer price index, each being driven by Heston's stochastic volatility process. To generate realistic economic scenarios, the model is calibrated quarterly using bond and option market data. Currently, as part of the calibration process, option prices are computed under the dynamics stated by CP2022 using a Monte Carlo-based approach. As an alternative, we resort to a Fourier method called the COS method (Fang and Oosterlee [2009a]), which relies on the availability of the characteristic function of the state factors. When interest rates are stochastic, the pricing of options can be reduced to the one-dimensional COS method by performing a cosine-series expansion of the discounted probability density function. This is achieved by incorporating the integrated process inside the characteristic function, which combines the dynamics of both the underlying and integrated state factors. CP2022 does not have a closed-form analytical characteristic function and must therefore be computed numerically, which is time-consuming. However, we accelerate the process by approximating the characteristic function using Chebyshev polynomials—a technique combined with the COS method we refer to as the Cheby-COS method. We use this scheme to efficiently price European call and put options, and derive the transitional characteristic function defined at two future time points for pricing cliquet or ratchet options under CP2022.

Compared to the Monte Carlo method, both the COS and Cheby-COS method demonstrate superior accuracy and computational efficiency. In particular, the Cheby-COS method excels in scenarios where the characteristic function must be computed numerically from large systems of Riccati ordinary differential equations that arise in affine models for option pricing.

Keywords: CP2022, multidimensional stochastic processes, affine, option pricing, COS method, Monte Carlo, vanilla options, cliquet options, ratchet options

This page was intentionally left blank.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to several individuals who have supported me throughout this journey. First and foremost, I extend my deepest appreciation to Dr. Damiaan Chen for introducing me to this project and for insightful suggestions. Working on such a fascinating and impactful project has been a privilege.

I am also deeply grateful to Sebastiaan Bye for the daily supervision, unwavering support, constructive feedback and valuable insights. The experience at the Expert Center for Financial Risk in the Pension Sector was both unique and enriching, and I am thankful for the opportunity to learn and grow in such an inspiring environment.

I would also like to express my gratitude to Dr. Fang Fang for her guidance, insightful feedback, thoughtful suggestions, and continuous belief in me. Her encouragement and emphasis on critical thinking have greatly contributed to the development of this work.

Furthermore, I am thankful to Dr. Daniel Dimitrov, Wouter Honig, Dr. Calmer Roos, and Prof. Dr. Sweder van Wijnbergen for the fruitful and stimulating discussions.

To all my colleagues at De Nederlandsche Bank, I truly appreciate the collaborative and pleasant working environment.

I would also like to thank Prof. Dr. ir. Cornelis Vuik and Dr. Jasper Anderluh for being part of the thesis committee.

Finally, my heartfelt gratitude go to my family and friends for their relentless support and encouragement throughout this journey.

Vlad Popa
Zoetermeer, March 2025

This page was intentionally left blank.

CONTENTS

This page was intentionally left blank.

LIST OF FIGURES

This page was intentionally left blank.

LIST OF TABLES

SYMBOLS & ABBREVIATIONS

Table 1: List of symbols and abbreviations used throughout the thesis.

symbols and abbreviations	definition
$\mathbb{P}$	real-world probability measure
$\mathbb{Q}$	risk-neutral measure
$\mathbb{Q}^T$	T-forward measure
$\mathcal{F}_t$	filtration at time t
$\mathbb{E}_t$	expectation conditioned on $\mathcal{F}_t$
$\left(\frac{d\mathbb{A}}{d\mathbb{B}}\Big\vert_{\mathcal{F}_t}\right)_{t\geq 0}$	Radon-Nikodym density to change from measure $\mathbb{A}$ to measure $\mathbb{B}$
W_t	Brownian Motion
$B(t)$	bank account or money-market account
$D(t,T)$	stochastic discount factor
r_t	short rate
$\tau(T,S)$	time to maturity with exact day count convention
$P(t,T), p_t(\tau)$	price of a zero-coupon bond at time t
$R(t,T)$	continuously-compounded spot rate
$y_t(\tau)$	bond yield, which for zero-coupon bonds is $R(t,T)$ implied from $P(t,T)$
$L(t,T)$	simply-compounded spot rate, also equals LIBOR with 360 day count convention
$F(t;T,S)$	simply-compounded forward rate
$f(t,T)$	instantaneous forward rate
LIBOR	London Interbank Offered Rate
EURIBOR	Euro Interbank Offered Rate
ZCB	Zero-Coupon Bond
FRA	Forward-Rate Agreement
IRS	Interest-Rate Swap
PFS	Payer Interest-Rate Swap
RFS	Receiver Interest-Rate Swap
RSW	Receiver Swaption
PSW	Payer Swaption
HICP	Harmonised Index of Consumer Prices
ZCIIS	Zero-Coupon Inflation-Indexed Swap
YYIIS	Year-on-Year Inflation-Indexed Swap
IC	Inflation-Indexed Cap
IF	Inflation-Indexed Floor
YIC	Year-on-Year Inflation-Indexed Cap
YIF	Year-on-Year Inflation-Indexed Floor
UFR	Ultimate Forward Rate
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation

This page was intentionally left blank.

1 INTRODUCTION

The Netherlands is in the middle of a transition to a new pension system defined under The Future of Pensions Act (Ministry of Social Affairs and Employment [2022]) effective as of July 1st, 2023. All Dutch pension funds must comply with the new legislation, ultimately by January 1st, 2028. This transition is necessary to overcome limitations of the current pension system, which the new system aims to address.

Under the current defined benefits system, pension funds promise specific benefits to their participants. However, if investment returns fall short, they may face challenges in fulfilling these promises. While collective reserves held by the funds are meant to back these commitments, if the reserves prove insufficient, the deficit is evenly distributed among participants of all age groups. This approach can lead to an unequal intergenerational playing field. Furthermore, the old system had Uniform Contribution Policies, such as the "Doorsneesystematiek". In this old system, all employees paid the same premium percentage and accrued the same pension benefits for each euro contributed, regardless of age. Contributions from the younger generation are invested for a longer period and as a result yield higher returns by retirement. As a consequence, the young generations effectively subsidized the older generations, with the expectation that they will be subsidized later. This system worked effectively when employees remained in the same pension plan throughout their careers and the age distribution among participants is balanced. This is no longer the case due to current demographic shifts and the increasing tendency of employees to change jobs more frequently. In combination with the historic low interest rates from 2009 till 2022, pension funds, according to buffer requirements, had to hold much regulatory capital to meet their obligations, which led to them cutting future indexations of pensions.

This, among the aforementioned reasons, mandated changes in the regulations. The new pension system is defined contribution-based, eliminates promises about future benefit amounts, and provides each participant a clear view of their share of collective assets under management. In other words, each participant's contributions are invested by pension funds and record the personal part of the collective assets under management in individual pension capital funds. Additionally, pension funds now offer life-cycle investment strategies, where investment risk gradually decreases as participants approach retirement age. Moreover, employers and employees can also allocate a portion of the collective pension capital as a buffer in advance for setbacks, under established prior agreements.

Currently, to ensure that pension funds are able to fulfil their financial obligations, they must adhere to the Financial Assessment Framework (FTK), which outlines the methods for calculating liabilities, buffer reserves, and the quantification of various types of risks. As part of the FTK, a feasibility test must be performed, which is a scenario-based analysis of pension funds' investment strategies and their pension policies based on economic scenario sets projected 100 years into future. These scenarios are generated by a stochastic model, which is calibrated based on market data. The economic scenario models are revised every 5 years by a commission appointed by the Dutch government, called the Commission Parameters, and is implemented by De Nederlandsche Bank (DNB).

Every quarter, DNB publishes (De Nederlandsche Bank [2022]) these sets of economic scenarios ($\mathbb{P}$-scenarios) and equivalent risk-neutral scenarios ($\mathbb{Q}$-scenarios). The current economic scenarios are generated by a model called the CP2022 model (Commission Parameters [2022b]). The $\mathbb{P}$-sets are used for defining a variety of risk attitudes, communication with participants on fu-

ture pension benefits, and establishing an annual tax-free contribution limit for participants. The Q-sets are used for the valuation of pension contracts and quantifying transition effects on participants with net benefits. The net benefit represents the difference between the market-consistent value of current and future accruals of pension benefits and the market-consistent valuation of future premium contributions. Changes in the net benefit are used to quantify the redistribution effects between participant groups that arise during the pension system transition (Chen and van Wijnbergen [2020]). The Q-sets are also used to determine compensation for age groups disadvantaged by the abolition of the "doorsneesystematiek".

The CP2022 model is an affine and arbitrage-free model framework for correlated interest rate, inflation rate, stock index and consumer price index, each being driven by Heston's stochastic volatility process. Affinity is an important property that allows for tractability in many ways (Duffie et al. [2000]). To be able to generate realistic economic scenarios in the future based on rapidly changing markets, the model is calibrated quarterly based on available bond and option prices. Currently, as part of the calibration process, option prices are computed under the dynamics stated by CP2022 using a Monte Carlo-based approach, which is time-consuming. There exist alternative methods in the literature that use the affinity property to derive the characteristic function as the solution to a system of coupled Riccati ordinary differential equations. Using this property, options can be priced using Fourier methods (Fang and Oosterlee [2009a]) that are faster than Monte Carlo. Unfortunately, the CP2022 model does not have a closed-form analytical characteristic function, and must therefore be solved using numerical methods. For the purpose of model calibration, the speed of the option pricing method is crucial, given that pension funds' policy changes are also based on the outcome of these sets.

The contributions of this thesis are twofold. First, we present a comprehensive overview of the model and address any missing gaps in clarity in Commission Parameters [2022b], to make it more accessible and understandable to practitioners. Second, we propose an efficient numerical scheme for pricing options, which is superior to the Monte Carlo methods employed in production. In particular, we focus on efficiently pricing European call and put options and cliquet-style options under CP2022.

The remainder of this thesis is structured as follows. Chapter 2 covers the mathematical foundations of stochastic calculus, mathematical finance, and interest-rate theory, along with a literature review, which are necessary for understanding the remainder of the thesis. In chapter 3, we give a complete framework for CP2022. Afterwards, we propose a highly efficient scheme for option pricing under the model framework and provide an error analysis in chapter 4. In chapter 5 we compare our approach with the Monte Carlo-based approach, perform convergence tests and compare the different option pricing methods during calibration. Finally, chapter 6 concludes the thesis and gives recommendations for future work.

This thesis was written as part of the Master of Science programme Applied Mathematics at Delft University of Technology.

2 PRELIMINARIES

In this chapter, we outline the preliminary theory required to understand the thesis. In particular, we address the topics of stochastic processes, option pricing, mathematical finance, and interest-rate theory.

2.1 STOCHASTIC PROCESSES

A stochastic process is a sequence of random variables that are indexed in continuous time. Examples of stochastic processes are the paths of stock prices or interest rates. Today, the exact historical paths can be observed but their future path are unknown and can only be modelled using simulations based on the process dynamics. Unless stated otherwise, the mathematical theorems with proofs can be found in the works of Shreve et al. [2004] and Øksendal [2003].

Definition 2.1.1 (Probability space)**.** *We define the sample space Ω as the set of all possible outcomes. The event space $\mathcal{F}$, which is a σ-algebra, represents the set of subsets of Ω corresponding to specific events. The probability measure $\mathbb{P}$ assigns a probability value to each event in $\mathcal{F}$. Together these quantities form the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The sequence of σ-algebras $\{\mathcal{F}_t, t \in \mathbb{R}_{\geq 0}\}$ is a filtration if*

1. *the sequence is increasing: $\mathcal{F}_0 \subset ... \subset \mathcal{F}_t \subset ...$, and*
2. *$\forall t \in \mathbb{R}_{\geq 0}, \quad \mathcal{F}_t \in \mathcal{F}$.*

A stochastic process X_t is adapted to the filtration $\mathcal{F}_t$ if X_t is $\mathcal{F}_t$-measurable. In other words, this means that the stochastic process is known up to time t. A special type of stochastic processes are martingales.

Definition 2.1.2 (Martingale)**.** *Suppose we have a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with filtration $\mathcal{F}_t$. We say that the process X_t is a martingale if*

1. *X_t is $\mathcal{F}_t$-measurable,*
2. *X_t is integrable, i.e. $\mathbb{E}|X_t| < \infty$, and*
3. *the martingale property holds $\forall t < T$: $\mathbb{E}\left[X_T | \mathcal{F}_t\right] = X_t$.*

The Wiener process, known as Brownian Motion, is the fundamental martingale stochastic process in mathematical finance.

Definition 2.1.3 (Brownian Motion)**.** *The stochastic process W_t is called a Brownian Motion if*

1. *$W(0) = 0$,*
2. *W_t is almost surely continuous,*
3. *W_t has independent increments, that is $\forall t_1 < t_2 < t_3 < t_4$, $W_{t_2} - W_{t_1}$ and $W_{t_4} - W_{t_3}$ are independent and*
4. *$\forall s < t$: $W_{t-s} \sim N(0, t-s)$.*

Stochastic differential equations (SDEs) describe the continuous-time dynamics of stochastic processes. They play a crucial part in financial mathematics, modelling random behaviours of asset prices and interest rates. We introduce a special class of systems of SDEs that play a central role in this thesis. For any vector or matrix A, we denote its transpose by A'.

Definition 2.1.4 (Affine class). *Suppose we have a system of SDEs of the form*

$$\mathrm{d}\boldsymbol{X}_t = \bar{\mu}\left(\boldsymbol{X}_t\right)\mathrm{d}t + \bar{\sigma}\left(\boldsymbol{X}_t\right)\mathrm{d}\boldsymbol{W}_t,$$

where $\boldsymbol{W}_t$ is a n-dimensional standard Brownian Motion. The processes belonging to the affine diffusion (AD) class are assumed to have drift, volatility, and interest-rate components of the affine form:

$$\begin{aligned}
&\bar{\mu}\left(\boldsymbol{X}_t\right) = a_0 + a_1\boldsymbol{X}_t, \quad (a_0, a_1) \in \mathbb{R}^n \times \mathbb{R}^{n\times n},\\
&\bar{\sigma}\left(\boldsymbol{X}_t\right)\bar{\sigma}\left(\boldsymbol{X}_t\right)' = c_0 + \sum_{i=1}^{n}\left(c_1^{(i)}\right)\boldsymbol{X}_t^{(i)}, \quad \left(c_0, c_1^{(i)}\right) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n}, \forall i \in \{1,...,n\},\\
&r\left(\boldsymbol{X}_t\right) = r_0 + r_1'\boldsymbol{X}_t, \quad (r_0, r_1) \in \mathbb{R} \times \mathbb{R}^n.
\end{aligned}$$

Additionally, the above may also be generalized to the cases that the coefficients $(a_0(t), a_1(t), r_0(t), r_1(t), c_0(t), c_1^1(t), ..., c_1^n(t))$ are time-dependent as well, as long as the coefficients are continuous. The assumption can be relaxed to piecewise continuous coefficients $(a_0(t), r_0(t), c_0(t))$, which in the literature is known as "regime switches". For technical details we refer to Filipović [2005] and Van Beek et al. [2020].

The defining characteristic of the affine class is that it provides explicit analytic formulas for pricing zero-coupon bonds, which we discuss in Section 2.4. The results were proved by Duffie et al. [2000] and are shown below.

Theorem 2.1.5 (Duffie et al. [2000]). *Suppose $\boldsymbol{X}_t \in \mathbb{R}^n$ is affine as in Definition 2.1.4 with $u \in \mathbb{R}^n$ and $\tau = T - t$. Then*

$$\xi(\boldsymbol{X}_t, t, T, u) \equiv \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r(X_s)ds}e^{iu\boldsymbol{X}_T}\middle|\mathcal{F}_t\right],$$

has a solution of the following form:

$$\xi(\boldsymbol{X}_t, t, T, u) = e^{\phi(u,t,T)+\Psi(u,t,T)'\boldsymbol{X}_t}.$$

The coefficients $\phi(u,t,T)$ and $\Psi(u,t,T)$ satisfy the complex-valued system of Riccati-type ordinary differential equations:

$$\begin{cases}
\frac{d}{d\tau}\phi(u,\tau) = -r_0 + \Psi' a_0 + \frac{1}{2}\Psi' c_0 \Psi,\\
\frac{d}{d\tau}\Psi(u,\tau)_j = -r_1^{(j)} + \Psi' a_1 + \frac{1}{2}\Psi' c_1^{(j)}\Psi, \quad \forall j \in \{1,...,n\},
\end{cases}$$

with terminal conditions $\phi(u,T,T) = 0, \Psi(u,T,T) = u$.

Brownian Motion has some important properties which makes stochastic calculus different from "deterministic" or normal calculus.

Definition 2.1.6 (Quadratic variation). *Suppose $\Pi = \{0 = t_0, t_1, ..., t_n = T\}$ is a partition of $[0,T]$ and denote $\|\Pi\| = \max_{j=0,...,n-1}(t_{j+1} - t_j)$. Then, the quadratic variation of a continuous function $f(t) : [0,T] \to \mathbb{R}$ is*

$$[f,f] = \lim_{\|\Pi\|\to 0}\sum_{j=0}^{n-1}[f(t_{j+1}) - f(t_j)]^2.$$

Denote $d[f] = d[f,f] = [df, df]$ as the quadratic variation from above. Roughly, the quadratic variation is equivalent to writing $d[f] = (df)^2$. The following holds for Brownian Motion W_t and time t:

$$d[t,t] = 0, \quad d[t,W_t] = d[W_t,t] = 0, \quad d[W_t,W_t] = dt.$$

The main concepts in stochastic calculus rely on the Itô integral, which differs from the standard Riemann integral due to the non-zero quadratic variation of Brownian motion. In this framework, both the integrands and integrators are stochastic processes.

Definition 2.1.7 (Itô integral). *Let $\{f(t) : t \geq 0\}$ be a stochastic process adapted to the filtration $\mathcal{F}_t$. The Itô integral is defined as*

$$\int_0^t f(t)dW_t = \lim_{n\to\infty}\sum_{i=0}^{n-1} f(t_i)(W(t_{i+1}) - W(t_i)).$$

The Itô integral has some important properties, such as continuity, linearity, and, under a condition, is a martingale.

Lemma 2.1.8 (Itô integral is a martingale)**.** *Let $\{f(t) : t \geq 0\}$ be a stochastic process adapted to the filtration $\mathcal{F}_t$. If*

$$\mathbb{E}\left[\int_0^T f^2(t)dt\right] < \infty,$$

then $I(T) = \int_0^T f(t)dW_t$ is a martingale.

Unlike in regular calculus, differentiating in a stochastic setting yields different results.

Lemma 2.1.9 (Itô's lemma)**.** *Let $f(t,x)$ be a function such that the partial derivative f_t, f_x, f_{xx} are continuous. Suppose there is an Itô process*

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t,$$

where $\mu(t, X_t)$ and $\sigma(t, X_t)$ are coefficients that satisfy the following Lipschitz conditions:

$$|\mu(t,x) - \mu(t,y)|^2 + |\sigma(t,x) - \sigma(t,y)|^2 \leq K_1|x-y|^2,$$
$$|\mu(t,x)|^2 + |\sigma(t,x)|^2 \leq K_2(1+|x|^2),$$

for some $K_1, K_2 \in \mathbb{R}_{>0}$. Then $f(t, X_t)$ is also an Itô process with dynamics:

$$df(t,X) = \frac{\partial f(t,X)}{\partial t}dt + \frac{\partial f(t,X)}{\partial X}dX + \frac{1}{2}\frac{\partial^2 f(t,X)}{\partial X^2}d[X,X],$$

where $d[X,X]$ is the quadratic variation of the stochastic process X_t.

The above properties can be generalized to the multivariate case.

Lemma 2.1.10 (Multivariate Itô's lemma)**.** *Let $f(t,\mathbf{x})$ be a function such that the partial derivative $f_t, f_{x_i}, f_{x_i x_j}$ are continuous in each component. Suppose there is an Itô process*

$$d\mathbf{X}_t = \boldsymbol{\mu}(t, \mathbf{X}_t)dt + \boldsymbol{\sigma}(t, \mathbf{X}_t)d\mathbf{W}_t,$$

where $\boldsymbol{\mu}(t, \mathbf{X}_t)$ and $\boldsymbol{\sigma}(t, \mathbf{X}_t)$ are coefficients that satisfy the Lipschitz conditions:

$$\|\boldsymbol{\mu}(t,\mathbf{x}) - \boldsymbol{\mu}(t,\mathbf{y})\|^2 + \|\boldsymbol{\sigma}(t,\mathbf{x}) - \boldsymbol{\sigma}(t,\mathbf{y})\|^2 \leq K_1\|\mathbf{x}-\mathbf{y}\|^2,$$
$$\|\boldsymbol{\mu}(t,\mathbf{x})\|^2 + \|\boldsymbol{\sigma}(t,\mathbf{x})\|^2 \leq K_2(1+\|\mathbf{x}\|^2),$$

with the norm for $\boldsymbol{\sigma} \in \mathbb{R}^{n\times m}$ defined as

$$\|\boldsymbol{\sigma}\|^2 = \sum_{i=1}^{n}\sum_{j=1}^{m}\sigma_{ij}^2,$$

for some $K_1, K_2 \in \mathbb{R}_{>0}$. Then $f(t, \mathbf{X}_t)$ is also an Itô process with dynamics:

$$\begin{aligned} df(t,\mathbf{X}) &= \frac{\partial f(t,\mathbf{X})}{\partial t}dt + \frac{\partial f(t,\mathbf{X})}{\partial \mathbf{X}'}d\mathbf{X} + \frac{1}{2}d\mathbf{X}'\frac{\partial^2 f(t,\mathbf{X})}{\partial \mathbf{X}\partial \mathbf{X}'}d\mathbf{X} \\ &= \left(\frac{\partial f(t,\mathbf{X})}{\partial t} + \boldsymbol{\mu}(t,\mathbf{X})'\frac{\partial f(t,\mathbf{X})}{\partial \mathbf{X}} + \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{\sigma}(t,\mathbf{X})\boldsymbol{\sigma}(t,\mathbf{X})'\frac{\partial^2 f(t,\mathbf{X})}{\partial \mathbf{X}\partial \mathbf{X}'}\right]\right)dt + \frac{\partial f(t,\mathbf{X})}{\partial \mathbf{X}'}\boldsymbol{\sigma}(t,\mathbf{X})d\mathbf{W}_t. \end{aligned}$$

Note that we use $\mathrm{Tr}[\,\cdot\,]$ to denote the trace of a square matrix. While we have seen that the Itô integral is a martingale, the converse is also true: any martingale can be expressed as an Itô integral.

Theorem 2.1.11 (Martingale representation theorem)**.** *Suppose W_t is a Brownian Motion with probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with filtration $\mathcal{F}_t$ and X_t is a martingale with respect to $\mathcal{F}$. Then, for some adapted $f(t)$ with $\mathbb{E}\left[\int_0^T f(s)^2 ds\right] < \infty$, X_t is driftless:*

$$X_t = \mathbb{E}X_0 + \int_0^T f(s)dW_s.$$

We conclude this section with a simple numerical scheme to find the solution of an Itô process. Not every Itô process admits a solution, though under some conditions, a solution exists almost surely.

Theorem 2.1.12 (Existence and uniqueness of a solution of an Itô process). *Assume an Itô process with measurable drift vector $\boldsymbol{\mu}(t, \mathbf{x})$ and diffusion matrix $\boldsymbol{\sigma}(t, \mathbf{x})$. If the Lipschitz conditions hold as in Lemma 2.1.10, and $\mathbb{E}\left[|\mathbf{X}_0|^2\right] < \infty$, then there exists an almost surely t-continuous unique solution $\mathbf{X}_t$ to the Itô process and*

$$\mathbb{E}\left[\int_0^T |\mathbf{X}_t|^2 dt\right] < \infty.$$

Analytically solving an Itô process is not always feasible. In that case, we need to resort to a numerical scheme to find the solution. The remainder of this section is based on Kloeden and Platen [1992]. Many numerical schemes exist, for an extensive overview, we refer to the before-mentioned literature.

First we introduce the notion of order of convergence of a numerical scheme in a stochastic setting.

Definition 2.1.13 (Order of convergence). *A general time discrete approximation $\tilde{\mathbf{X}}_T$ with constant step-size h converges weakly with order $\alpha > 0$ to $\mathbf{X}_T$, if for any continuous function f, $\exists C > 0$ independent of h such that:*

$$\left|\mathbb{E}[f(\mathbf{X}_T) - f(\tilde{\mathbf{X}}_T)]\right| \leq Ch^\alpha.$$

Similarly, $\tilde{\mathbf{X}}_T$ is said to converge strongly with order $\beta > 0$ if $\exists C > 0$ independent of h such that:

$$\mathbb{E}|\mathbf{X}_T - \tilde{\mathbf{X}}_T| \leq Ch^\beta.$$

The simplest numerical scheme is the Euler–Maruyama scheme, which is stated below.

Lemma 2.1.14 (Euler–Maruyama). *Partition $[0, T]$ into N equidistant intervals with $h = T/N$: $\{0 = t_0, t_1, ..., t_N = T\}$. A numerical solution to an Itô process is found recursively by setting $\tilde{\mathbf{X}}_0 = \mathbf{X}_0$ and*

$$\tilde{\mathbf{X}}_{n+1} = \tilde{\mathbf{X}}_n + \boldsymbol{\mu}(t_n, \tilde{\mathbf{X}}_n)(t_{n+1} - t_n) + \boldsymbol{\sigma}(t_n, \tilde{\mathbf{X}}_n)(\mathbf{W}_{t_{n+1}} - \mathbf{W}_{t_n}).$$

The Euler–Maruyama scheme has a weak order of convergence of $\alpha = 1$ and a strong order of convergence of $\beta = 0.5$.

2.2 OPTIONS

In this section we introduce the terminology for option pricing. For additional information we refer to Higham [2004], Hull [2017] and Shreve et al. [2004].

Consider an asset, which is a financial entity, where the value is known up till time t and future values are unknown but can be modelled with a stochastic process. In its simplest form, an option is a contract between a holder and a writer, where the option's value is derived from an underlying asset according to specific predetermined rules. The underlying asset can be stocks, interest rates, inflation rates, currencies, bonds, commodities, other options, weather variables and more. The most common option type are vanilla options, also known as European options, which we define below.

Definition 2.2.1 (European Call). *A European call option gives the holder the right, but not the obligation, to purchase from the underwriter a prescribed asset (also known as the underlying) at an agreed-upon price (the strike price K) at a prescribed time in the future (the expiration date T)*

Definition 2.2.2 (European Put). *A European put option gives the holder the right, but not the obligation, to sell to the underwriter a prescribed asset (also known as the underlying) at an agreed-upon price (the strike price K) at a prescribed time in the future (the expiration date T)*

At maturity, the holder can choose to exercise its right to either buy or sell an underlying asset S_t. It goes without saying that for call options if $S_T > K$, the holder will exercise and buy at price K. If $S_T < K$, the holder will not exercise. A similar reasoning holds for European puts, thus the payout at maturity equals

$$V^{\text{call}}(t = T, T) = \max\{S_T - K, 0\} \equiv (S_T - K)^+,$$
$$V^{\text{put}}(t = T, T) = \max\{K - S_T, 0\} \equiv (K - S_T)^+.$$

Note that European calls and puts are related with one another, since $\max\{S_T - K, 0\} - \max\{K - S_T, 0\} = S_T - K$. Evaluating this relationship at an earlier time $t < T$ yields the well-known put-call parity. European options' payoff depend only on the terminal value of the asset S_T. However, so-called exotic options can feature more complex payoffs, which can be path-dependent of the underlying. We list only a few common examples, as in reality many more exist.

- American options can be exercised at any time before the expiration date.
- Bermudan options can be exercised at a finite number of dates before the expiration date.
- Asian options' payoff at the expiration date depends on the average price of the underlying instead of the terminal value.
- Barrier options introduces a barrier level B, where depending if the underlying has crossed the barrier level or not, the holder is allowed to exercise the option.
- Basket options' payoff depends on two or more different underlying assets.
- Forward-starting options start at a specified future date and expire further in the future. A series of consecutive forward-starting options are known as cliquet or ratchet options.

The fair price $V(t, T)$ of an option depends on the assumed stochastic dynamics of the underlying asset(s). Closed-form solutions for $V(t, T)$ are available only for a limited number of models with simplified market assumptions. In most cases, numerical methods are required, with Monte Carlo simulations being among the most flexible and straightforward to implement. However, they are also very time-consuming to obtain accurate option prices. In the following section, we present the necessary tools for pricing options.

2.3 MATHEMATICAL FINANCE

In this section we introduce the theory behind martingale measures, changing measures and pricing options. Unless stated otherwise, for proofs we refer to Shreve et al. [2004] and Øksendal [2003].

Theorem 2.3.1 (Radon-Nikodym). *Let $\mathbb{P}, \mathbb{Q}$ be two probability measures on the space $(\Omega, \mathcal{F})$.*
$\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, meaning that they are equivalent, and we write $\mathbb{Q} \sim \mathbb{P}$.
$\iff$
There exists a random variable $X \geq 0$ and $X \in L^1(\mathbb{P})$ such that $\mathbb{E}_{\mathbb{Q}}(\mathbb{1}_A) = \mathbb{E}_{\mathbb{P}}(X\mathbb{1}_A)$, $\forall A \in \mathcal{F}$. Note that in this context, $L^p(\mathbb{P}) = \{X : X \text{ is } \mathcal{F}\text{-measurable and } \mathbb{E}|X|^p < \infty\}$. We say X is the density of $\mathbb{Q}$ with respect to $\mathbb{P}$ and is denoted by $\frac{d\mathbb{Q}}{d\mathbb{P}}$. Also, X is $\mathbb{P}$-a.s. unique.

Definition 2.3.2 (Equivalent martingale measure). *A probability measure $\mathbb{Q}$ is a martingale measure for a process X_t if the discounted price process $\tilde{X}_t = X_t e^{-\int_t^T r(s)ds}$ is a martingale under $\mathbb{Q}$, where $r(s)$ is a risk-free interest-rate process. A probability measure $\mathbb{Q}$ is an equivalent martingale measure if it is a martingale measure and it is equivalent to the "real world" measure $\mathbb{P}$.*

The following theorems will be essential for understanding how to price options.

Theorem 2.3.3 (Fundamental theorem of asset pricing 1). *A financial market is free of arbitrage if and only if there exists an equivalent martingale measure $\mathbb{Q}$.*

For a technical proof we refer to Delbaen and Schachermayer [1998]. In the context of financial markets, arbitrage simply means that there is an opportunity to make a *riskless* profit with no initial capital. This excludes the trivial scenario of returns from a money savings account. The formulation above is a foundational result upon which most, if not all, concepts in mathematical finance are built (hence its name). This allows us to price options via the following methods:

Theorem 2.3.4 (Martingale pricing)**.** *Suppose the financial market is free of arbitrage. Let $r(t)$ be an $\mathcal{F}_t$-measurable risk-free interest-rate process and let S_t be a stochastic asset price process. Suppose an option has a payoff $V(T, S_T)$ on the underlying S_t. Then, by the Fundamental Theorem of Asset Pricing 1, there exists an equivalent martingale measure $\mathbb{Q}$ such that*

$$V(t, S_t) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r(s)ds} V(T, S_T) \middle| \mathcal{F}_t \right].$$

An alternative method exists to price options. Suppose there are two different portfolios of financial products which produce the same cash flows in the future. Then they must have the same price today. If the prices were different, an arbitrage opportunity would arise, allowing traders to buy the cheaper portfolio and sell the more expensive one, securing a risk-free profit. We call this principle the law of 1 price.

Theorem 2.3.5 (Law of 1 price)**.** *Suppose a financial market is free of arbitrage. Suppose there are two portfolios Y_t and Z_t, that have an identical payoff at future time T. Then both portfolios must have the same price, i.e. $Y_0 = Z_0$.*

We state under what conditions a financial market model is complete.

Definition 2.3.6 (Completeness)**.** *A financial market is complete if every option can be replicated.*

Theorem 2.3.7 (Fundamental theorem of asset pricing 2)**.** *There exists a unique equivalent martingale measure $\mathbb{Q}$ if and only if the financial market is arbitrage-free and complete.*

For a technical proof we refer to Delbaen and Schachermayer [1994]. Sometimes it is easier to price options under a different measure. For that, we introduce the notion of numéraire.

Definition 2.3.8 (Numéraire)**.** *A numéraire is the unit of account in which other assets are denominated. Note that any tradeable asset can be taken as a numéraire, with the only condition being that it needs to be observable and tradeable. Otherwise, option prices cannot be computed using neither the martingale approach nor any replicating portfolio utilizing the law of 1 price relative to the numéraire.*

We remark that it is not possible to take non-observable variables as numéraires like volatilities or the instantaneous spot rate. The numéraire is usually set as a country's currency but can be changed to another currency when needed, such as in foreign exchange markets. Adjusting the numéraire can simplify calculations when pricing options.

In principle, we can take any positively priced asset as a numéraire and denominate all other assets in terms of the chosen numéraire. Associated with each numéraire is a risk-neutral measure. When making this association, we take only non-dividend-paying assets as numéraires. In particular, we regard $\mathbb{Q}$ as the risk-neutral measure associated with the money market account.

Changing measures is equivalent to changing the numéraire.

Theorem 2.3.9 (Change of measure)**.** *Let $\mathbb{A} \sim \mathbb{B}$. Assume that the associated numéraire with $\mathbb{A}$ is A and the associated numéraire with $\mathbb{B}$ is B. This means that the price of any traded asset X relative to a numéraire is a martingale under its respective measure:*

$$\mathbb{E}^{\mathbb{A}}\left[\frac{X(T)}{A(T)} \middle| \mathcal{F}_t \right] = \frac{X(t)}{A(t)}, \qquad \mathbb{E}^{B}\left[\frac{X(T)}{B(T)} \middle| \mathcal{F}_t \right] = \frac{X(t)}{B(t)}.$$

Then for $0 \leq t_0 \leq t \leq T$, the Radon-Nikodym derivative that defines the measure $\mathbb{A}$ is the density process

$$(Z_t)_{t \geq 0} = \left(\frac{d\mathbb{A}}{d\mathbb{B}}\bigg|_{\mathcal{F}_t} \right)_{t \geq 0} = \frac{A(t)B(t_0)}{A(t_0)B(t)}$$

and if $X \in L^1(\mathbb{A})$ and $X \in \mathcal{F}_t$, then

$$\mathbb{E}^{\mathbb{A}}[X|\mathcal{F}_t] = \frac{\mathbb{E}^{\mathbb{B}}\left[X\left(\frac{d\mathbb{A}}{d\mathbb{B}}\Big|_{\mathcal{F}_T}\right)_{T\geq 0}\Big|\mathcal{F}_t\right]}{\mathbb{E}^{\mathbb{B}}\left[\left(\frac{d\mathbb{A}}{d\mathbb{B}}\Big|_{\mathcal{F}_T}\right)_{T\geq 0}\Big|\mathcal{F}_t\right]} = \frac{\mathbb{E}^{\mathbb{B}}[XZ_T|\mathcal{F}_t]}{\mathbb{E}^{\mathbb{B}}[Z_T|\mathcal{F}_t]} = \frac{1}{Z_t}\mathbb{E}^{\mathbb{B}}[XZ_T|\mathcal{F}_t].$$

This theorem allows us to switch measures when calculating conditional expectations. Whenever we wish to change the drift of a stochastic differential equation, that is whenever we want to derive the stochastic dynamics under a new equivalent measure, Girsanov's theorem is fundamental.

Theorem 2.3.10 (Girsanov's theorem)**.** *Suppose we have a stochastic process with Lipschitz coefficients μ and σ under measure $\mathbb{A}$:*

$$dX = \mu^{\mathbb{A}}(X,t)dt + \sigma(X,t)dW_t^{\mathbb{A}}.$$

Suppose under the new measure $\mathbb{B}$ we wish to have drift $\mu^{\mathbb{B}}$. Define the adapted process $\Theta_t = (\mu^{\mathbb{B}}(X,t) - \mu^{\mathbb{A}}(X,t))/\sigma(X,t)$. Define the Radon-Nikodym density as

$$Z_t = \left(\frac{d\mathbb{B}}{d\mathbb{A}}\Big|_{\mathcal{F}_t}\right)_{t\geq 0} = \exp\left(-\int_0^t \Theta_s dW_s - \frac{1}{2}\int_0^t \Theta_s^2 ds\right),$$

and assume

$$\mathbb{E}^{\mathbb{B}}\left[\int_0^t \Theta_s^2 Z_s^2 ds\right] < \infty.$$

Then $\mathbb{B} \sim \mathbb{A}$ and

$$dW^{\mathbb{B}} = dW^{\mathbb{A}} - \Theta_t dt,$$

where $W^{\mathbb{B}}$ is a Brownian Motion under the new measure $\mathbb{B}$ such that

$$dX = \mu^{\mathbb{B}}(X,t)dt + \sigma(X,t)dW_t^{\mathbb{B}}.$$

We conclude this section by noting that, instead of calculating option prices through conditional expectations or changing measures, it is also possible to reformulate the option pricing problem as the solution of a partial differential equation (PDE).

Theorem 2.3.11 (Feynman-Kac)**.** *Suppose we have the partial differential equation with coefficients $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ that satisfy the Lipschitz conditions and has the form*

$$\frac{\partial f}{\partial t} + \boldsymbol{\mu}(\mathbf{X},t)'\frac{\partial f}{\partial \mathbf{X}} + \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{\sigma}(t,\mathbf{X})\boldsymbol{\sigma}(t,\mathbf{X})'\frac{\partial^2 f(t,\mathbf{X})}{\partial \mathbf{X}\partial \mathbf{X}'}\right] - V(\mathbf{X},t)f(\mathbf{X},t) = 0,$$

subject to the terminal condition: $f(\mathbf{X}_T, T) = g(\mathbf{X}_T)$.

Then, the solution to the PDE above can be written as:

$$f(\mathbf{X},t) = E(e^{-\int_t^T V(\mathbf{X}_\tau,\tau)d\tau}g(\mathbf{X}_T)|\mathbf{X}_t = \mathbf{x}),$$

with $d\mathbf{X}_t = \boldsymbol{\mu}(\mathbf{X},t)dt + \boldsymbol{\sigma}(\mathbf{X},t)d\mathbf{W}_t$.

In summary, we have outlined three key pillars of option pricing. On one hand, a replicating portfolio is constructed to align the cash flows with the payouts of the option in question. On the other hand, the underlying asset price is modelled using a stochastic process, and options are priced either by solving a conditional expectation or by solving a PDE.

2.4 INTEREST-RATE THEORY

In this section, we present the various concepts of interest rates and introduce the fundamentals of interest-rate theory, based on Brigo and Mercurio [2006]. Afterwards, we discuss well-known interest-rate options: swaps, caps, floors, swaptions and inflation-indexed options.

Various types of interest rates exist. In particular, we distinguish between interbank rates and government rates. Government rates are derived from bonds issued by the government, while interbank rates refer to the rates at which deposits are lent between banks or the rates applied in swap transactions between banks. Zero-coupon rates, derived from bonds in both government and interbank markets, create distinct zero-coupon curves. The mathematical modelling of these rates is similar in both cases. The most well-known interbank rate was the London InterBank Offered Rate (LIBOR), which fixed daily in London. However, other interbank rates exist, such as the Euro Interbank Offered Rate (EURIBOR), which is fixed in Brussels.

In 2012, incidents came to light where major banks were found to have colluded to artificially influence LIBOR rates to benefit their trading positions and enhance their financial positions (Hou and Skeie [2014]). This scandal, known as the "LIBOR scandal", prompted a major regulatory and legal overhaul, which eventually led to the discontinuation of the LIBOR rates in 2023 (Bank of England [2023], FCA [2024]).

Bank Account and the Short Rate

When depositing money into a bank account, individuals expect their funds to grow over time at a certain rate. Investors anticipate a reward for the risk of lending to a party that could potentially default, however small that risk may be. This compensation reflects the time value of money, as receiving a specific amount in the future differs from receiving the same amount today. While these concepts are well-known, expressing them mathematically requires introducing several definitions to establish a coherent framework.

We consider the bank account, also known as the money-market account. In this account, continuous profit accrues at the prevailing risk-free rate in the market at any given moment in time.

Definition 2.4.1 (Bank account)**.** *We define $B(t)$ to be the value of a bank account at time $t \geq 0$. We assume $B(0) = 1$ and that the bank account evolves according to the following differential equation:*

$$dB(t) = r_t B(t)dt, \quad B(0) = 1,$$

where r_t is a positive function of time. As a consequence the solution to the equation above is

$$B(t) = \exp\left(\int_0^t r_s ds\right).$$

In here, r_t is the instantaneous rate at which the bank account accrues over a short time dt. This instantaneous rate is usually referred to as instantaneous spot rate, or briefly as short rate. Note that the short rate cannot be observed directly; it is used to define all other observable types of interest rates that are quoted in the market from different bonds. Calculating the value of cash flows at future time T at present time t can be done via the stochastic discount factor. After all, if the bank account grows from $B(0)$ to $B(T)$, we can also discount the above to get back to $B(0)$.

Definition 2.4.2 (Stochastic discount factor)**.** *The (stochastic discount) factor $D(t,T)$ between two time instants t and T is the amount at time t that is "equivalent" to one unit of currency payable at time T, and is given by*

$$D(t,T) = \frac{B(t)}{B(T)} = \exp\left(-\int_t^T r_s ds\right).$$

Making r_t stochastic is significant because it directly affects the underlying asset in question, the bank account numéraire $B(t)$. In numerous pricing scenarios, particularly when the Black-Scholes formula is used in equity or foreign-exchange (FX) markets, r_t is often assumed to be a deterministic function of time. This assumption results in both the bank account and the discount factors at any future time becoming deterministic functions of time. This assumption is based on the belief that interest-rate movements have a relatively small impact on the pricing of equity or FX options compared to movements in the underlying assets.

However, when working with interest-rate products, the primary factor influencing product prices are interest rates themselves. The most well-known interest-rate product is a zero-coupon bond.

Definition 2.4.3 (Nominal zero-coupon bond)**.** *A T-maturity (nominal) zero-coupon bond is a contract that guarantees its holder the payment of one unit of currency at time T, with no intermediate payments. The contract value at time $t < T$ is denoted by $P(t,T)$. Therefore $P(T,T) = 1$.*

If r_t is deterministic, then $D(t,T)$ is deterministic as well and $D(t,T) = P(t,T)$ by Theorem 2.3.5. In case rates are stochastic, then by Theorem 2.3.4, the zero-coupon-bond price equals

$$P(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r(s)ds}\middle|\mathcal{F}_t\right]. \tag{2.1}$$

Determining the present value of a future payment is thus equivalent to knowing the zero-coupon bond price for that particular future time. These prices are fundamental to interest-rate theory. They often serve as fundamental reference points from which all types of interest rates can be derived and expressed in terms of zero-coupon bond prices. It is important to note that while interest rates are commonly quoted in the markets, zero-coupon bonds exist as theoretical instruments and are not directly observable in the market.

We refer to the nominal economy when inflation is not accounted for, whereas in the real economy, inflation is adjusted for. For instance, if $I(t)$ represents the consumer price index, the price of a real zero-coupon bond is given by:

$$P^R(t,T) = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r(s)\,ds}\,\frac{I(T)}{I(t)}\middle|\mathcal{F}_t\right]. \tag{2.2}$$

In the absence of any specification, we assume a nominal economy.

Transitioning between zero-coupon bond prices and interest rates requires knowledge of the compounding type and the day-count convention, which gives rise to multiple types of interest rates.

Definition 2.4.4 (Continuously-compounded spot interest rate)**.** *The continuously-compounded spot interest rate at time t with maturity T is denoted by $R(t,T)$ and is the constant rate at which an investment of $P(t,T)$ units of currency at time t accrues continuously to yield a unit amount of currency at maturity T:*

$$R(t,T) \equiv -\frac{\ln P(t,T)}{\tau(t,T)}$$

Note that by default we define the time to maturity as $\tau(t,T) = T - t$ in years.

An alternative to continuous compounding is simple compounding, which applies when accruing occurs proportionally to the time of the investment.

Definition 2.4.5 (Simply-compounded spot interest rate)**.** *The simply compounded spot interest rate prevailing at time t for the maturity T is denoted by $L(t,T)$ and is the constant rate at which an investment has to be made to produce an amount of one unit of currency at maturity, starting from $P(t,T)$ units of currency at time t, when accruing occurs proportionally to the investment time:*

$$L(t,T) \equiv \frac{1 - P(t,T)}{\tau(t,T)P(t,T)}.$$

EURIBOR rates are simply-compounded rates and are linked to zero-coupon bond prices with a 360 day count convention. We can express bond prices in terms of simply-compounded rates as:

$$P(t,T) = \frac{1}{1 + L(t,T)\tau(t,T)}.$$

There are multiple compounding methods like annual compounding, but we will not delve further into the specifics, other than mention that in a similar fashion we can derive the price of zero-coupon bonds in terms of annual compounded rates.

All preceding definitions of spot interest rates are equivalent in the limit of infinitesimal time intervals. That is, the short rate can be derived as a limit of the various rates defined earlier using L'Hôpital's rule:

$$\begin{aligned} r(t) &= \lim_{T \to t^+} R(t,T) \\ &= \lim_{T \to t^+} L(t,T). \end{aligned}$$

Forward Rates

Forward rates are defined by three distinct time points: the time at consideration t, the expiry $T \geq t$, and maturity $S \geq T$. Forward rates are interest rates that can be secured today for future investments.

A forward rate can be understood through a forward-rate agreement (FRA). This contract ensures an interest-rate payment between T and S. At maturity S, a fixed payment based on a pre-set rate K is exchanged for a floating payment based on the spot rate $L(T,S)$ reset at T and maturing at S. This agreement allows locking in a simply compounded interest rate between T and S at a specific value K. In particular, at time S, one receives $\tau(T,S)KN$ currency units and pays $\tau(T,S)L(T,S)N$ currency units, with N representing the contract's nominal value. Therefore, the contract's value at time S is

$$\mathbf{FRA}(t = S, T, S, \tau(T,S), N, K) = N\tau(T,S)(K - L(T,S)),$$

assuming all rates use the same day-count convention. In terms of the zero-coupon bond prices, this expression can be rewritten as

$$\mathbf{FRA}(t = S, T, S, \tau(T,S), N, K) = N\left[\tau(T,S)K - \frac{1}{P(T,S)} + 1\right],$$

and the contract's value at present time t is

$$\mathbf{FRA}(t, T, S, \tau(T,S), N, K) = N[P(t,S)\tau(T,S)K - P(t,T) + P(t,S)].$$

There is one value of K such that the price is zero at time t. This resulting rate defines the simply-compounded forward interest rate.

Definition 2.4.6 (Simply-compounded forward interest rate)**.** *The simply-compounded forward interest rate prevailing at time t for the expiry $T > t$ and maturity $S > T$ is denoted by $F(t;T,S)$ and equals*

$$F(t;T,S) \equiv \frac{1}{\tau(T,S)}\left(\frac{P(t,T)}{P(t,S)} - 1\right).$$

Rewriting the FRA in terms of the simply-compounded forward interest rate gives:

$$\mathbf{FRA}(t, T, S, \tau(T,S), N, K) = NP(t,S)\tau(T,S)(K - F(t;T,S)). \tag{2.3}$$

If we use a zero-coupon bond as a numéraire, the corresponding measure is the T-forward measure and for any traded asset $V(t)$ we have

$$\frac{V(t)}{P(t,T)} = \mathbb{E}^{\mathbb{Q}^T}\left[\frac{V(T)}{P(T,T)}\middle|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}^T}\left[V(T)|\mathcal{F}_t\right].$$

Note that by Definition 2.4.6,

$$P(t,S)F(t;T,S) = \frac{1}{\tau(T,S)}(P(t,T) - P(t,S)),$$

and the left-hand side is a traded asset, because it is a multiple of the difference of zero-coupon bonds. Thus using $P(t,S)$ as numéraire, we have

$$F(t;T,S) = \frac{P(t,S)F(t;T,S)}{P(t,S)} = \mathbb{E}^{\mathbb{Q}^T}\left[\frac{P(S,S)F(S;T,S)}{P(S,S)}\middle|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{Q}^T}\left[F(S;T,S)|\mathcal{F}_t\right].$$

In other words, under the T-forward measure, $F(t,T,S)$ is a martingale. This property is what gave this measure its name. Remark that $F(T;T,S) = L(T,S)$ and because $F(t,T,S)$ is a martingale,

$$\mathbb{E}^{\mathbb{Q}^T}[L(T,S)|\mathcal{F}_t] = F(t;T,S).$$

As the maturity of the forward rate approaches its expiry, it converges to the instantaneous forward rate, which is

$$\begin{aligned}\lim_{S\to T^+} F(t;T,S) &= -\lim_{S\to T^+}\frac{1}{P(t,S)}\frac{P(t,S)-P(t,T)}{S-T}\\ &= -\frac{1}{P(t,T)}\frac{\partial P(t,T)}{\partial T}\\ &= -\frac{\partial \ln P(t,T)}{\partial T},\end{aligned}$$

where $\tau(T,S) = S - T$.

Definition 2.4.7 (Instantaneous forward interest rate)**.** *The instantaneous forward interest rate prevailing at time t for the maturity $T > t$ is denoted by $f(t,T)$ and is defined as*

$$f(t,T) \equiv \lim_{S\to T^+} F(t;T,S) = -\frac{\partial \ln P(t,T)}{\partial T},$$

so that we also have

$$P(t,T) = \exp\left(-\int_t^T f(t,u)du\right).$$

An assumption in Definition 2.4.7 is that the zero-coupon-price function $T \mapsto P(t,T)$ is smooth. Instantaneous forward rates play an important role in determining whether arbitrage is present or not in interest-rate models (Heath et al. [1992a]).

Before we continue, we need a lemma from calculus. This lemma will also be used in the following chapter.

Lemma 2.4.8 (Leibniz integral rule)**.** *Suppose $f(x,t), f_x(x,t), f_t(x,t), a(x), a_x(x), b(x), b_x(x)$ are continuous. Then*

$$\frac{d}{dx}\left(\int_{a(x)}^{b(x)} f(x,t)dt\right) = \frac{db(x)}{dx}f(x,b(x)) - \frac{da(x)}{dx}f(x,a(x)) + \int_{a(x)}^{b(x)}\frac{\partial f(x,t)}{\partial x}dt.$$

Proof. See Protter and Morrey [2012]. □

Finally, we show the relationship between the short rate and the instantaneous forward interest rate:

$$\begin{aligned}
\mathbb{E}^{\mathbb{Q}^T}[r_T \mid \mathcal{F}_t] &= \left(\frac{d\mathbb{Q}^T}{d\mathbb{Q}}\bigg|_{\mathcal{F}_t}\right)^{-1} \mathbb{E}^{\mathbb{Q}}\left[\frac{d\mathbb{Q}^T}{d\mathbb{Q}}\bigg|_{\mathcal{F}_T} r_T \,\middle|\, \mathcal{F}_t\right] \\
&= \frac{P(t_0,T)B(t)}{P(t,T)B(t_0)}\mathbb{E}^{\mathbb{Q}}\left[\frac{P(T,T)B(t_0)}{P(t_0,T)B(T)} r_T \,\middle|\, \mathcal{F}_t\right] \\
&= \frac{1}{P(t,T)}\mathbb{E}^{\mathbb{Q}}\left[\frac{B(t)}{B(T)} r_T \,\middle|\, \mathcal{F}_t\right] \\
&= \frac{1}{P(t,T)}\mathbb{E}^{\mathbb{Q}}\left[r_T e^{-\int_t^T r_s ds} \,\middle|\, \mathcal{F}_t\right] \\
&\overset{\text{(i)}}{=} -\frac{1}{P(t,T)}\mathbb{E}^{\mathbb{Q}}\left[\frac{\partial}{\partial T} e^{-\int_t^T r_s ds} \,\middle|\, \mathcal{F}_t\right] \\
&\overset{\text{(ii)}}{=} -\frac{1}{P(t,T)}\frac{\partial P(t,T)}{\partial T} \\
&= f(t,T),
\end{aligned}$$

where in (i) we use the Leibniz integration rule and exchanging the order of the partial derivative and expectation in (ii) is justified by Lebesgue's dominated convergence theorem as the stochastic discount factor is bounded by 1. For the remainder of this section, we will discuss well-known interest-rate options, such as swaps, caps, floors, swaptions, and inflation-indexed interest-rate options.

Interest-Rate Swaps

We introduce the notion of an Interest-Rate Swap (IRS), which is a generalization of the forward-rate agreement. In this contract, exchange payments are made between two differently indexed legs, starting from a future leg. We refer to a leg as a pre-specified time instance when a payment is made. At every leg T_i in a pre-specified set of dates $T_{\alpha+1}, \ldots, T_\beta$, the fixed leg pays out the amount $N\tau_i K$ corresponding to a fixed interest rate K, a nominal value N and a year fraction τ_i between T_{i-1} and T_i. The floating leg pays the amount $N\tau_i L(T_{i-1}, T_i)$ corresponding to the interest rate $L(T_{i-1}, T_i)$ resetting at the previous instant T_{i-1} for the maturity given by the current payment instant T_i for a given date T_α. Note that the floating-leg rate resets at dates $T_\alpha, T_{\alpha+1}, \ldots, T_{\beta-1}$ and pays at dates $T_{\alpha+1}, \ldots, T_\beta$. We set $\mathcal{T} = \{T_\alpha, \ldots, T_\beta\}$ and $\tau = \{\tau_{\alpha+1}, \ldots, \tau_\beta\}$.

We assume that fixed-rate payments and floating-rate payments occur at the same dates and with the same day count convention. When the fixed leg is paid and the floating leg is received, the IRS is called a Payer IRS (PFS), while conversely the IRS is a Receiver IRS (RFS).

The discounted payoff at a time $t < T_\alpha$ of the PFS is

$$\sum_{i=\alpha+1}^{\beta} D(t, T_i) N\tau_i \left(L(T_{i-1}, T_i) - K\right),$$

while for the RFS we have

$$\sum_{i=\alpha+1}^{\beta} D(t, T_i) N\tau_i \left(K - L(T_{i-1}, T_i)\right).$$

Thus, the RFS can be seen as a portfolio of FRAs and using Equation 2.3 yields:

$$\begin{aligned}\mathbf{RFS}(t,\mathcal{T},\tau,N,K) &= \sum_{i=\alpha+1}^{\beta} \mathbf{FRA}\left(t,T_{i-1},T_i,\tau_i,N,K\right)\\ &= N\sum_{i=\alpha+1}^{\beta}\tau_i P\left(t,T_i\right)\left(K-F\left(t;T_{i-1},T_i\right)\right)\\ &= -NP\left(t,T_\alpha\right)+NP\left(t,T_\beta\right)+N\sum_{i=\alpha+1}^{\beta}\tau_i KP\left(t,T_i\right).\end{aligned}$$

The strike K that makes contract value zero at time t is the forward swap rate.

Definition 2.4.9 (Forward swap rate)**.** *The forward swap rate $S_{\alpha,\beta}(t)$ at time t for the sets of times $\mathcal{T}$ and year fractions τ is the rate in the fixed leg of the FRS that makes the FRS a fair contract at the present time. It is the fixed rate K for which $\mathbf{RFS}(t,\mathcal{T},\tau,N,K)=0$:*

$$S_{\alpha,\beta}(t)=\frac{P\left(t,T_\alpha\right)-P\left(t,T_\beta\right)}{\sum_{i=\alpha+1}^{\beta}\tau_i P\left(t,T_i\right)}\equiv\frac{P\left(t,T_\alpha\right)-P\left(t,T_\beta\right)}{P_{\alpha+1,\beta}(t)},$$

where we have defined $P_{\alpha+1,\beta}(t)=\sum_{i=\alpha+1}^{\beta}\tau_i P\left(t,T_i\right)$.

Caps and Floors

A cap is a contract that can be viewed as a payer IRS where each exchange payment is executed only if it is positive. The cap price at time t equals:

$$\mathbf{Cap}\left(t,\mathcal{T},\tau,N,K\right)=\sum_{i=\alpha+1}^{\beta}\mathbb{E}^{\mathbb{Q}}\left[D(t,T_i)N\tau_i\left(L\left(T_{i-1},T_i\right)-K\right)^{+}\middle|\,\mathcal{F}_t\right],$$

Analogously, a floor is equivalent to a receiver IRS where each exchange payment is executed only if it is positive. The floor price at time t equals:

$$\mathbf{Floor}\left(t,\mathcal{T},\tau,N,K\right)=\sum_{i=\alpha+1}^{\beta}\mathbb{E}^{\mathbb{Q}}\left[D(t,T_i)N\tau_i\left(K-L(T_{i-1},T_i)\right)^{+}\middle|\,\mathcal{F}_t\right].$$

Using the identity $(a-b)^{+}-(b-a)^{+}=a-b,\forall a,b\in\mathbb{R}$, caps and floors can be related to swaps:

$$\mathbf{Floor}\left(t,\mathcal{T},\tau,N,K\right)-\mathbf{Cap}\left(t,\mathcal{T},\tau,N,K\right)=\mathbf{RFS}(t,\mathcal{T},\tau,N,K)$$

A market practice is to price these caps and floors under the Black model (Black [1976]), which is done as

$$\mathbf{Cap}^{\text{Black}}\left(0,\mathcal{T},\tau,N,K,\sigma_{\alpha,\beta}\right)=N\sum_{i=\alpha+1}^{\beta}P\left(0,T_i\right)\tau_i\,\mathrm{Bl}\left(K,F\left(0,T_{i-1},T_i\right),v_i,1\right)$$

where Φ is the standard Gaussian cumulative distribution function and

$$\begin{aligned}\mathrm{Bl}(K,F,v,\omega) &= F\omega\Phi\left(\omega d_1(K,F,v)\right)-K\omega\Phi\left(\omega d_2(K,F,v)\right),\\ d_1(K,F,v) &= \frac{\ln(F/K)+v^2/2}{v},\\ d_2(K,F,v) &= \frac{\ln(F/K)-v^2/2}{v},\\ v_i &= \sigma_{\alpha,\beta}\sqrt{T_{i-1}}.\end{aligned}$$

The volatility parameter $\sigma_{\alpha,\beta}$ represents the instantaneous volatility of the forward rates between T_α and T_β and are quoted on the market. Analogously, the corresponding floor is priced as:

$$\textbf{Floor}^{\text{Black}}\left(0,\mathcal{T},\tau,N,K,\sigma_{\alpha,\beta}\right)=N\sum_{i=\alpha+1}^{\beta}P\left(0,T_i\right)\tau_i\,\text{Bl}\left(K,F\left(0,T_{i-1},T_i\right),v_i,-1\right).$$

Swaptions

Swap options, more commonly known as swaptions, are options on an IRS. There are two main types of swaptions, a payer version and a receiver version.

A European payer swaption is an option giving the right (but not the obligation) to enter a payer IRS at a given future time, which we denote by the swaption maturity. A European receiver swaption gives the holder the right to enter a receiver IRS. Therefore, a receiver option will be exercised only if the receiver IRS is positive, and similarly to the payer case. Hence, the option price of a receiver and payer swaption is

$$\begin{aligned}\textbf{RSW}\left(t,\mathcal{T},\tau,N,K\right)&=\mathbb{E}^{\mathbb{Q}}\left[D(t,T_\alpha)\left(-NP\left(T_\alpha,T_\alpha\right)+NP\left(T_\alpha,T_\beta\right)+NK\sum_{i=\alpha+1}^{\beta}\tau_iP\left(T_\alpha,T_i\right)\right)^{+}\middle|\,\mathcal{F}_t\right]\\&=N\mathbb{E}^{\mathbb{Q}}\left[D(t,T_\alpha)P_{\alpha+1,\beta}(T_\alpha)\,\left(K-S_{\alpha,\beta}(T_\alpha)\right)^{+}\middle|\,\mathcal{F}_t\right],\end{aligned}$$

$$\begin{aligned}\textbf{PSW}\left(t,\mathcal{T},\tau,N,K\right)&=\mathbb{E}^{\mathbb{Q}}\left[D(t,T_\alpha)\left(NP\left(T_\alpha,T_\alpha\right)-NP\left(T_\alpha,T_\beta\right)-NK\sum_{i=\alpha+1}^{\beta}\tau_iP\left(T_\alpha,T_i\right)\right)^{+}\middle|\,\mathcal{F}_t\right]\\&=N\mathbb{E}^{\mathbb{Q}}\left[D(t,T_\alpha)P_{\alpha+1,\beta}(T_\alpha)\,\left(S_{\alpha,\beta}(T_\alpha)-K\right)^{+}\middle|\,\mathcal{F}_t\right].\end{aligned}$$

Again, it is market practice to value swaptions under the Black model (Black [1976]):

$$\textbf{Swaption}^{\text{Black}}\left(0,\mathcal{T},\tau,N,K,\sigma_{\alpha,\beta}\right)=N\text{Bl}\left(K,S_{\alpha,\beta}(0),\sigma_{\alpha,\beta}\sqrt{T_\alpha},w\right)\sum_{i=\alpha+1}^{\beta}\tau_iP\left(0,T_i\right),$$

where $w=-1$ for a RSW and $w=1$ for a PSW. Analogous to how caps and floors are related to swaps, using the same identity, a similar relationship exists between receiver swaptions, payer swaptions and swaps:

$$\textbf{RSW}\left(t,\mathcal{T},\tau,N,K\right)-\textbf{PSW}\left(t,\mathcal{T},\tau,N,K\right)=\textbf{RFS}(t,\mathcal{T},\tau,N,K).$$

Inflation-indexed swaps

Given a series of payment dates $\{T_1,\ldots,T_\alpha\}$, an inflation-indexed swap is a financial contract in which two parties exchange cash flows tied to inflation. On each specified date T_i, Party A makes a payment to Party B based on the inflation rate observed over a predefined period, which forms the floating leg of the contract. Meanwhile, Party B pays Party A a fixed rate, which is often referred to as the strike or the fixed leg, regardless of the actual inflation observed during the period. The floating leg is typically linked to a widely recognized inflation index, such as the Harmonised Index of Consumer Prices (HICP), while the fixed leg represents a constant payment throughout the life of the swap. The two primary types of inflation-indexed swaps traded in the market are the zero-coupon inflation-indexed swap (ZCIIS) and the year-on-year inflation-indexed swap (YYIIS).

In a ZCIIS, Party B pays Party A the fixed amount $N\left[(1+K)^{\alpha}-1\right]$ at the maturity date T_{α}, where K is the fixed rate, also known as the break-even inflation rate. In return, Party A pays Party B the floating amount $N\left[I(T_{\alpha})/I(t)-1\right]$.

In a YYIIS, at each yearly payment date T_i, Party B pays Party A the fixed amount NK. In return, at each yearly payment date T_i, Party A pays Party B the floating amount $N\left[I(T_i)/I(T_{i-1})-1\right]$.

The pricing of these swaps follows a similar approach to pricing interest-rate swaps:

$$\mathbf{ZCIIS}(t,T_{\alpha},N,K)=N\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_{\alpha}}r_s ds}\left[\frac{I(T_{\alpha})}{I(t)}-(1+K)^{\alpha}\right]\middle|\mathcal{F}_t\right],$$

and

$$\mathbf{YYIIS}(t,T_{\alpha},N,K)=N\sum_{i=1}^{\alpha}\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_i}r_s ds}\left[\frac{I(T_i)}{I(T_{i-1})}-(1+K)\right]\middle|\mathcal{F}_t\right].$$

Inflation-indexed caps and floors

An inflation-indexed cap is a call option on the inflation rate for a given strike. A zero-coupon inflation-indexed cap gives the holder the right, but not the obligation, to enter into a ZCIIS. In an analogous way, an inflation-indexed floor is a put option on the inflation rate. Thus the payout for a zero-coupon inflation-indexed cap and floor at maturity is

$$N\left(w\left(\frac{I(T_{\alpha})}{I(T_0)}-(1+K)^{\alpha}\right)\right)^{+},$$

where $w=1$ for a cap and $w=-1$ for a floor. For a year-on-year inflation-indexed cap and floor, the payout is similar but based on YYIIS:

$$N\left(w\left(\frac{I(T_i)}{I(T_{i-1})}-(1+K)\right)\right)^{+}.$$

Analogous to pricing inflation-indexed swaps, the prices of zero-coupon inflation-indexed caps and floors and year-on-year inflation-indexed caps and floors is respectively:

$$\begin{aligned}
\mathbf{IC}(t,T_{\alpha},N,K)&=N\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_{\alpha}}r_s ds}\left(\frac{I(T_{\alpha})}{I(t)}-(1+K)^{\alpha}\right)^{+}\middle|\mathcal{F}_t\right],\\
\mathbf{IF}(t,T_{\alpha},N,K)&=N\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_{\alpha}}r_s ds}\left((1+K)^{\alpha}-\frac{I(T_{\alpha})}{I(t)}\right)^{+}\middle|\mathcal{F}_t\right],\\
\mathbf{YIC}(t,T_{\alpha},N,K)&=N\sum_{i=1}^{\alpha}\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_i}r_s ds}\left(\frac{I(T_i)}{I(T_{i-1})}-(1+K)\right)^{+}\middle|\mathcal{F}_t\right],\\
\mathbf{YIF}(t,T_{\alpha},N,K)&=N\sum_{i=1}^{\alpha}\mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^{T_i}r_s ds}\left((1+K)-\frac{I(T_i)}{I(T_{i-1})}\right)^{+}\middle|\mathcal{F}_t\right].
\end{aligned}$$

2.5 EXISTING INTEREST-RATE MODELS

There are many different types of models for interest rates. We give an overview of the most common types of models.

One method is to model the short rate using a stochastic differential equation, from which the other interest rates can be derived based on this process. We make a distinction between so-called one-factor models and multi-factor models. In one-factor models, there is only one source

of uncertainty that drives the short rate process. Some well-known examples are the Vasicek model (Vasicek [1977]) and Cox–Ingersoll–Ross (CIR) model (Cox et al. [1985]). The advantage of one-factor models is that they are very tractable. However, one-factor models are limited to only one source of uncertainty and hence replicate a limited number of curves of the term structure of interest rates. Multi-factor models address this issue by incorporating multiple sources of uncertainty, providing a more accurate representation of the term structure. However, they come with the drawback of being significantly more computationally expensive to calibrate and sacrifices tractability.

We say that modelling the short rate via a stochastic differential equation creates an endogenous term structure of interest rates. That is, the model dynamics characterizes the evolution of the complete term structure. As a result they are prone to model misspecification errors and they cannot fit the observed term structure of interest rates exactly. On the other hand, models that take the initial term structure as given and then describe the dynamics of the whole term structure are said to create exogenous term structures of interest rates. A well-known example of the latter is the Hull-White model (Hull and White [1990]). All of the above models can be further extended by incorporating stochastic jumps.

Many additional approaches exist in practice. Instead of modelling short rates, another approach is to model forward rates, which is the Heath–Jarrow–Morton framework (Heath et al. [1992b]). A different approach is to use parametric or non-parametric methods to obtain the yield curves using market data. Well-known parametric models use curves that belong to the Nelson-Siegel family (Nelson and Siegel [1987]). For an overview of more methods like stripping, splines, regressions via kernels, linear programming, time series, artificial neural networks and more, we refer to James and Webber [2000].

Throughout this thesis, we use the terms ”term structure of interest rates” and ”yield curves” interchangeably, as they refer to the same concept.

2.6 LITERATURE REVIEW

2.6.1 Past Commission Parameters Models

As part of the Financial Assessment Framework (FTK), pension funds have to perform feasibility tests based on economic scenario sets projected 100 years into the future. Regulatory supervisors assess the risk profile of pension funds by stress-testing their investment strategies and pension policies. This includes factors like participant contributions, pension indexation (adjustments for inflation or wage growth), and the term structure used to discount future cash flows. These sets were initially produced from the Koijen-Nijman-Werker (KNW) model (Koijen et al. [2009]) and was adopted by Commission Parameters [2014]. The KNW model is based on the works of Brennan and Xia [2002], and is a Gaussian Ornstein–Uhlenbeck two-factor model with two state variables without economic interpretation, but drive the dynamics of time-varying bond risk premia, interest rates and inflation rates in a nominal and real economy. With subsequent modifications to the original KNW model by Draper [2014] using Dutch market data. The model was further extended by Muns [2015], who added more restrictions during calibration to prevent convergence to undesirable local optima as a result of optimizing in a high-dimensional setting, and was later adopted by Commission Parameters [2019]. The model was designed to be simple, yet gave relatively realistic scenarios at the time. However, there were some limitations. The two factors driving the economies lacked economic interpretation. Moreover, calibration was done using a modified Kalman filter on bond yield data, equity returns and inflation data. Omitting option data prevented capturing the volatility and correlations of options in the equity and interest-rate markets. Additionally, in certain cases, the model was generating negative interest rates in the long term, which is not realistic. Since the model assumed constant volatilities, all risk factors followed a Gaussian distribution. Consequently, well-known stylized facts in empirical research, such as heavy tails and skewness, were not captured by the model.

2.6.2 Option Pricing

The literature on option pricing is vast and has evolved considerably since the Black-Scholes framework provided an analytical solution for vanilla options. Analytical solutions are available only for relatively simple underlying dynamics, where the probability density function of the asset is known. For most other dynamics, numerical methods must be employed.

A binomial tree method was introduced by Cox et al. [1979], which approximates option prices on a discrete-time lattice, and was later extended to n dimensions (Boyle et al. [1989]). Tree-based methods are particularly popular for pricing options with early exercise features, as optimal exercise strategies can be calculated through backward induction. A drawback of tree-based methods is that the number of nodes grows exponentially with the number of underlying assets, leading to very slow convergence for pricing basket options.

The most flexible and intuitive method for pricing options is via Monte Carlo simulations, where many paths of the underlying assets are simulated to approximate the expectation of discounted payoffs, based on the Central Limit Theorem. A major advantage of Monte Carlo methods is their simplicity in implementation. Additionally, they exhibit a convergence rate of $n^{-1/2}$, which is independent of the number of dimensions (or underlying assets) involved. This characteristic allows Monte Carlo simulations to avoid the curse of dimensionality that typically affects tree-based methods, making them more scalable and efficient for high-dimensional problems. However, while the convergence rate of Monte Carlo methods is relatively slow for a single underlying asset, various variance-reduction techniques can improve the efficiency. For a comprehensive overview of (Quasi-) Monte Carlo methods and these techniques, we refer to Glasserman [2004].

In cases where the probability density function of the underlying asset is unknown, but its characteristic function can be analytically derived as outlined by Duffie et al. [2000], option prices can be computed through Fourier inversion, which can be efficiently implemented using Fast Fourier Transform algorithms. This approach, initially proposed by Carr and Madan [1999], has become popular for pricing vanilla options within stochastic volatility and jump-diffusion models. While the method can be extended to exotic options, its implementation is not always straightforward and becomes less efficient in higher dimensions. Since then, different Fourier-based approaches and modifications were made, see for example Lee et al. [2004], Lord and Kahl [2006] and Fang and Oosterlee [2009a].

Up till now we have discussed various methods for approximating option prices via the martingale-based approach. However, since the Feynman-Kac theorem provides a connection between PDEs and stochastic processes, stating that the solution to parabolic PDEs can be represented as the expectation of a function of a stochastic process. Thus, option prices can also be computed by numerically solving PDEs, using methods such as finite differences or spectral methods, see Foulon et al. [2010] and Ezzine et al. [2024] and the references therein. The advantage of PDE-based option pricing methods is that they allow for the calculation of option prices across the entire discretization domain simultaneously. However, these methods need to be redefined for each new pricing problem, and explicit schemes may suffer from numerical instability for certain discretization values of the domain. Additionally, the computational complexity grows significantly for higher-dimensional PDEs or fine grids.

The large international collaboration project "BENCHOP" has benchmarked a wide range of numerical methods on various option pricing problems (von Sydow et al. [2015], von Sydow et al. [2019]). The results revealed that Fourier-based methods generally outperform other numerical approaches in terms of efficiency, with the "COS method" emerging as the most competitive. Given that the model examined in this thesis has a characteristic function that can be determined, the COS method will be the primary focus.

The Fourier-Cosine Series Expansion (COS) method for pricing options was introduced by Fang and Oosterlee [2009a], which, under appropriate conditions, converges exponentially (in contrast to the Carr and Madan approach). This method has been extended to two dimensions by Ruijter and Oosterlee [2012] and to n dimensions by Junike and Stier [2024]. It has a variety

of applications in the financial industry, such as pricing Bermudan and barrier options (Fang and Oosterlee [2011]), American options (Fang and Oosterlee [2009b]), Asian options (Zhang and Oosterlee [2013]), basket options (Junike and Stier [2024]), credit default swaps (Fang et al. [2009]), solving stochastic control problems (Ruijter et al. [2013]), and calculations of risk measures (Ortiz-Gracia and Oosterlee [2014]), counterparty credit risk (De Graaf et al. [2014]) and wrong way risk (Feng and Oosterlee [2016]). Additionally, the method has also been used in a data-driven approach (Leitao et al. [2018]) and a neural network-based approach (Liu et al. [2019]).

In option pricing, the COS method is based on performing a Fourier-cosine series expansion of the unknown probability density function of the underlying in the payoff function and relating its Fourier coefficients to the known characteristic function. In principle, the density function can also be approximated using different series expansion, such as wavelets (Ortiz-Gracia and Oosterlee [2013]). Additionally, other transforms other than Fourier-based transforms exist in the literature for option pricing, such as Hilbert transforms (Feng and Linetsky [2008]) or Gauss transforms (Broadie and Yamamoto [2003]).

Since in multi-state stochastic volatility jump-diffusion models with stochastic interest rates the characteristic function can often be determined, Fourier-based numerical methods, alongside with Monte Carlo methods, are commonly used for efficient option pricing.

For pricing cliquet options, in Broadie and Kaya [2006], Detlefsen and Härdle [2006], Bianchetti et al. [2015] and Korn et al. [2017], a Monte Carlo-based method is used. Pricing cliquet options with tree methods is discussed in Gaudenzi and Zanette [2011] and in Windcliff et al. [2006], a PDE-approach is used. Kruse and Nögel [2005] considers Heston's model with deterministic interest rates and prices cliquet options using the approach of Carr and Madan [1999]. van Haastrecht and Pelsser [2011] price vanilla and cliquet options under the Heston and Schöbel-Zhu stochastic volatility model with multi-factor Gaussian interest rates with a full correlation structure by deriving the forward characteristic function under the $T-$forward measure at two future times and price these options using the Carr and Madan approach. Switching to the $T-$forward measure eliminates the stochastic discount factor in the expectation, simplifying the characteristic function by focusing solely on the dynamics of the underlying. A similar approach is used by van Haastrecht et al. [2009] in the Schöbel–Zhu–Hull–White model.

2.7 A BRIEF SUMMARY OF THE DUTCH PENSION SYSTEM

The Dutch pension system has long been recognized as one of the best in the world, achieving the top ranking in 2024 (Mercer [2024]). This distinction highlights the system's effectiveness and sustainability, which we briefly address. As of 2024, Dutch pension funds manage total assets amounting to approximately 1.8 trillion euros, which represents roughly 180% of the Dutch Gross Domestic Product (GDP) in 2023 (De Nederlandsche Bank [2024]). This substantial figure underscores the importance of the pension system to the Dutch economy and its significant role in global pension management.

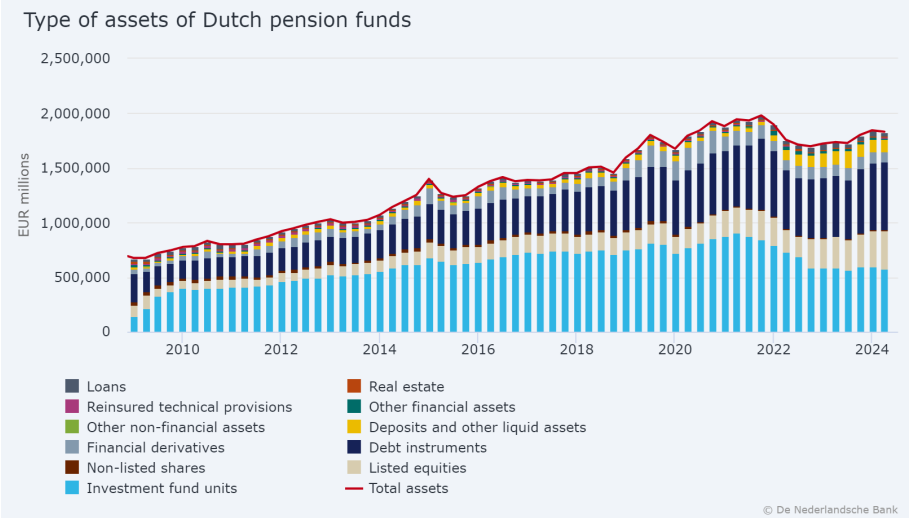

Figure 2.1: Total assets of Dutch pension funds, taken from De Nederlandsche Bank [2024].

The Dutch pension system can be summarized into 3 pillars, consisting of a public pension, occupational pension and a private saving pension.

The first pillar of the Dutch pension system is the General Old-Age Pensions Act (AOW), established in 1957. It provides a basic pension to all residents of the Netherlands upon reaching the qualifying age, regardless of employment history. Initially set at 65 years, the entitlement age has increased to 67 by 2024, with future adjustments linked to life expectancy. Residents accumulate 2% of their AOW entitlement annually over 50 years of residency and is linked to the development of the minimum wage. The AOW functions as a minimum-level provision, typically supplemented by income from additional pension pillars. The AOW is financed by tax revenues and by the workforce in the form of contributions rather than individual contributions.

The second pillar of the Dutch pension system comprises occupational pensions, which are accrued by the majority of employees during their employment. Contributions are made jointly by employers and employees and are allocated to a pension fund affiliated with the employer. These funds may serve entire industries or individual companies. As an employment benefit, participation in industry-wide pension funds are mostly mandatory. The pension amount is determined by the level of salary and the length of employment. The joint contributions by employees and employers are collectively invested by pension funds, which operate as non-profit.

The third pillar includes individual retirement saving accounts or insurance products that allow for voluntary pension contributions, with tax benefits up to a certain limit. This pillar is particularly important for individuals who cannot participate in the second pillar, such as self-employed individuals.

What sets the Dutch pension system apart from other European countries is the relatively equal distribution of funds between Pillar 1 and Pillar 2, whereas in most other European countries,

Pillar 1 is significantly larger than Pillar 2 (Eurostat [2024]). The AOW provides stability during economic downturns, while capital-funded pensions are less susceptible to demographic changes, such as population ageing and shifts in employees' career paths. The latter was, however, still a problem in the old system due to the "doorsneesystematiek". In this system, all employees paid the same premium percentage and accrued equal pension benefits per euro contributed, regardless of age. Younger generations, who invested for a longer period, generated higher returns by retirement, effectively subsidizing the older generations.

This system is under growing pressure due to increased life expectancy, which requires larger and longer-lasting pensions, while the number of young people contributing has relatively declined. In combination with the low interest-rates era of 2009 till 2022, pension funds, according to buffer requirements, had to hold much regulatory capital to meet their obligations, which led to cutting indexations of pensions.

This has resulted in a shift from a defined benefit system to one that is focused on defined contributions. Under the new system, a larger portion of the pension amount will depend on the individual contributions made during the accrual phase and the returns generated from investments managed by pension funds. These contributions will later be converted into periodic benefits paid out starting from the retirement age.

Pension funds now offer two contribution schemes: a solidarity-based scheme and a flexible scheme. The solidarity scheme provides a uniform life-cycle investment strategy, where investment risk gradually decreases as participants approach retirement age. In contrast, the flexible scheme allows participants to choose from various life-cycle investment strategies with different risk profiles, placing greater exposure to investment risks and interest-rate fluctuations on the employees.

Pension funds are also required to inform its participants about how market fluctuations could impact their future benefits. This is achieved by using an economic scenario-generating model to project the current economic scenarios and assess potential changes in future welfare. The current model is the CP2022 model, which we describe in detail in the next chapter.

3 CP2022: A COMPLETE FRAMEWORK

The $\mathbb{P}$- and $\mathbb{Q}$-economic scenario sets published by DNB are used by pension funds for many applications. The $\mathbb{Q}$-sets are used for valuing pension contracts, specifically to determine the present value of future cash flows that will be paid to participants starting from their pension age onwards. The $\mathbb{P}$-sets are used for a variety of purposes, such as defining risk attitudes, quantifying transition effects and pension fund communication with participants. In the previous model by Commission Parameters [2019], economic scenarios were generated by the Koijen-Nijman-Werker (KNW) model (Koijen et al. [2009]). The Dutch pension transition and the Wet toekomst pensioenen (Wtp) mandated an extension of the model which can generate $\mathbb{Q}$-scenarios for valuing pension contracts, in addition to $\mathbb{P}$-scenarios. This led to the development of a new model: CP2022 (Commission Parameters [2022b]). The model structure was chosen to fall within the affine class. This is a class of models that ensures that the so-called spot rates, also called zero-coupon rates with the corresponding discount factors, can be determined at any time in the future and for all possible future maturities in terms of the state variables.

The CP2022 model differs from the KNW model in a few aspects. Contrary to the KNW model, the volatility which drives all risk factors, is added as an additional stochastic state variable. In this way, a better match can be achieved with non-symmetric distributions observed in the market for future values of interest rates, inflation and stock returns. For pension funds' applications, it is crucial to estimate as accurately as possible the present value of uncertain future cash flows like pension premiums and benefits. The dynamics of the risk factors in CP2022 were designed so that the generated economic scenarios produce distributions of interest rates, inflation rates and stock returns with positive, positive and negative skewness respectively, which is supported by empirical evidence (European Central Bank [2024], World Bank [2023], Ari et al. [2023], Aparicio and Estrada [2001]). This stands in direct contrast to the KNW model, which generated only symmetric distributions.

The scenarios delivered by Commission Parameters [2019] formed a $\mathbb{P}$-set. Now, a $\mathbb{Q}$-set must be generated simultaneously and both sets must be consistent. In other words, the $\mathbb{P}$ and $\mathbb{Q}$ probability measures that generate these scenario sets must be equivalent. This equivalence can be guaranteed by assigning a market price to every random process that drives the dynamics of economic variables that compensates for the risk incurred.

The dynamics of risk premiums influence the compensation for interest rate and inflation risk, which affects the pricing of nominal and real bonds. Therefore, they affect the nominal and real interest-rate term structure, which maps the yields of bonds with different maturities. In Commission Parameters [2019], the risk premiums were assumed to be constant over time. As a result, the previous model could only create a limited number of yield curves. Although affine models are highly tractable, they cannot replicate every aspect of the curves and an exact fit of both current and all historical interest-rate term structures is not possible within the class of affine models. Given the applications of the scenario sets, it is desirable that the associated discount factors for risk-free deterministic cash flows correspond exactly to the discount curves as prescribed by DNB for pension funds. Therefore, Commission Parameters [2022a] introduces time-dependent future risk premiums so that the proposed affine model is able to generate yield curves that fit the observed market curves at the time of calibration. To simplify the calibration of historical market data, an approach with constant risk premiums was used for historical bond prices. The Commission Parameters prioritizes fitting to the yield curves at the time of calibration, as these curves are most important for the scenario set applications for pension funds. A consequence of this choice is that re-calibration will require time-dependent risk premiums for

interest and inflation to be re-estimated. This will become clear in Section 3.5, where we discuss how the model is calibrated.

Commission Parameters [2019] used in the calibration process monthly historical values of nominal zero-coupon bonds, a price index and a stock index. In the current model, options are added during calibration to ensure that the generated $\mathbb{Q}$-set is aligned with the market. Additionally, historical values of the VSTOXX, the volatility index associated with the EURO STOXX 50 index, is added. In this way, the models parameters can be estimated that determine the development of the new state variable, the stochastic volatility.

In this chapter, we discuss in greater detail the CP2022 model and address any missing gaps in clarity in Commission Parameters [2022b]. In particular, we examine the dynamics of CP2022 under both the $\mathbb{P}$ and $\mathbb{Q}$ measures with the resulting term structure of interest rates, discuss how the model is calibrated to align with current market conditions to generate consistent economic scenarios, and illustrate the simulations of these economic scenarios.

3.1 CP2022 DYNAMICS

The economic model CP2022 is based on a stochastic process $\mathbf{X}_t$:

$$\mathbf{X}_t = \begin{bmatrix} \mathbf{X}_t^s \\ \mathbf{X}_t^o \end{bmatrix}, \quad \mathbf{X}_t^s = \begin{bmatrix} v_t \\ r_t \\ \pi_t \end{bmatrix}, \quad \mathbf{X}_t^o = \begin{bmatrix} \ln(S_t) \\ \ln(\Pi_t) \end{bmatrix},$$

which consists of an unobservable state vector process $\mathbf{X}^s$, which contains a short rate process r_t, an instantaneous expected European inflation rate process π_t and a stochastic variance process v_t. Furthermore, v_t is the stochastic variance from the Heston model, which is the square of the stochastic volatility. The two additional observable variables in $\mathbf{X}_t^o$ are the logarithm of a stock price index S_t and the European consumer price index (CPI) Π_t.

The dynamics of $\mathbf{X}_t^s$ is described by

$$\begin{aligned} d\mathbf{X}_t^s &= \begin{bmatrix} K_{vv} & 0 & 0 \\ K_{vr} & K_{rr} & K_{\pi r} \\ K_{v\pi} & K_{r\pi} & K_{\pi\pi} \end{bmatrix} \left(\begin{bmatrix} \mathbb{E}v_\infty \\ \mathbb{E}r_\infty \\ \mathbb{E}\pi_\infty \end{bmatrix} - \mathbf{X}_t^s \right) dt + \begin{bmatrix} \omega & 0 & 0 & 0 & 0 \\ \sigma_{vr} & \sigma_{r1} & \sigma_{r2} & 0 & 0 \\ \sigma_{v\pi} & \sigma_{\pi 1} & \sigma_{\pi 2} & 0 & 0 \end{bmatrix} \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t \Gamma_1 \end{bmatrix}^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}} \\ &\equiv K \left(\mathbb{E}\mathbf{X}_\infty^s - \mathbf{X}_t^s \right) dt + \Sigma^{r\pi} \left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma \right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}}, \end{aligned} \tag{3.1}$$

with

$$\begin{aligned} \mathbb{E}\mathbf{X}_\infty^s &= \begin{bmatrix} \mathbb{E}v_\infty \\ \mathbb{E}r_\infty \\ \mathbb{E}\pi_\infty \end{bmatrix}, \quad K = \begin{bmatrix} K_{vv} & 0 & 0 \\ K_{vr} & K_{rr} & K_{\pi r} \\ K_{v\pi} & K_{r\pi} & K_{\pi\pi} \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 1 & 0_{1\times 4} \\ 0_{4\times 1} & \Gamma_1 \end{bmatrix}, \quad \Gamma_0 = \begin{bmatrix} 0 & 0_{1\times 4} \\ 0_{4\times 1} & I_4 \end{bmatrix}, \\ \Gamma_1 &= \begin{bmatrix} \gamma_{11} & 0 & 0 & 0 \\ 0 & \gamma_{22} & 0 & 0 \\ 0 & 0 & \gamma_{33} & 0 \\ 0 & 0 & 0 & \gamma_{44} \end{bmatrix}, \quad \Sigma^{r\pi} = \begin{bmatrix} \omega & 0 & 0 & 0 & 0 & 0 \\ \sigma_{vr} & \sigma_{r1} & \sigma_{r2} & 0 & 0 & 0 \\ \sigma_{v\pi} & \sigma_{\pi 1} & \sigma_{\pi 2} & 0 & 0 & 0 \end{bmatrix}. \end{aligned} \tag{3.2}$$

We denote $(\mathbf{X}_t^s)_1$ by the first element in $\mathbf{X}_t^s$, which is v_t in this case. Additionally, $K \in \mathbb{R}^{3\times 3}$ is the mean reversion matrix and controls the speed of convergence to the steady state vector $\mathbb{E}\mathbf{X}_\infty^s \in \mathbb{R}^3$. Moreover, $\Sigma^{r\pi} \in \mathbb{R}^{3\times 5}$ is the state covariance matrix and $W_t^{\mathbb{P}}$ is a 5 -dimensional standard Brownian Motion with independent components in $\mathbb{P}$. In the system of SDEs, the correlation

structure is determined by K and $\Sigma^{r\pi}$. Therefore, it is not necessary to work directly with correlated Brownian motions. Correlated Brownian motions can be written into independent Brownian motions using Cholesky factorization. For example, if the Brownian Motions have correlation structure

$$d\mathbf{W}_t d\mathbf{W}_t' = \begin{bmatrix} 1 & \rho_{12} & \rho_{13} \\ \rho_{12} & 1 & \rho_{23} \\ \rho_{13} & \rho_{23} & 1 \end{bmatrix} dt,$$

then $W_t = L\tilde{W}_t$, where $\tilde{W}_t$ has independent components and L follows from Cholesky decomposition:

$$L = \begin{bmatrix} 1 & 0 & 0 \\ \rho_{12} & \sqrt{1-\rho_{12}^2} & 0 \\ \rho_{13} & \frac{\rho_{23}-\rho_{12}\rho_{23}}{\sqrt{1-\rho_{12}^2}} & \sqrt{\frac{\rho_{13}^4+\rho_{12}^2+\rho_{23}^2-2\rho_{23}^2\rho_{12}\rho_{13}-1}{\rho_{12}^2-1}} \end{bmatrix}.$$

The logarithm of stock prices and European consumer prices in $\mathbf{X}_t^o$ evolve according to

$$\begin{aligned} d\mathbf{X}_t^o &= \begin{bmatrix} r_t + \eta_S - \frac{1}{2}\sigma_S' \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t\Gamma_1 \end{bmatrix} \sigma_S \\ \pi_t + \eta_\Pi - \frac{1}{2}\sigma_\Pi' \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t\Gamma_1 \end{bmatrix} \sigma_\Pi \end{bmatrix} dt + \begin{bmatrix} \sigma_S' \\ \sigma_\Pi' \end{bmatrix} \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t\Gamma_1 \end{bmatrix}^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}} \\ &= \begin{bmatrix} r_t + \eta_S \\ \pi_t + \eta_\Pi \end{bmatrix} dt - \frac{1}{2}\mathcal{D}\left(\begin{bmatrix} \sigma_S' \\ \sigma_\Pi' \end{bmatrix} \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t\Gamma_1 \end{bmatrix} \begin{bmatrix} \sigma_S' \\ \sigma_\Pi' \end{bmatrix}' \right) dt + \begin{bmatrix} \sigma_S' \\ \sigma_\Pi' \end{bmatrix} \begin{bmatrix} v_t & 0_{1\times 4} \\ 0_{4\times 1} & I_4 + v_t\Gamma_1 \end{bmatrix}^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}} \\ &\equiv \quad (\mu^o + K^o\mathbf{X}_t^s)\, dt + \Sigma^{S\Pi} \left(\Gamma_0 + (\mathbf{X}_t^s)_1\, \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}}, \end{aligned} \tag{3.3}$$

with

$$\begin{aligned} \mu^o &= \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} - \frac{1}{2}\mathcal{D}\left(\Sigma^{S\Pi}\Gamma_0\Sigma^{S\Pi\prime}\right), \quad \Sigma^{S\Pi} = \begin{bmatrix} \sigma_S' \\ \sigma_\Pi' \end{bmatrix}, \\ K^o &= \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} - \frac{1}{2}\mathcal{D}\left(\Sigma^{S\Pi}\Gamma\Sigma^{S\Pi\prime}\right) \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, \end{aligned} \tag{3.4}$$

and $\mathcal{D}(A)$ is the vector that equals the diagonal of matrix A. Additionally, $\eta_S \in \mathbb{R}$ denotes the equity risk premium, $\eta_\Pi \in \mathbb{R}$ is the risk premium on inflation and $\Sigma^{S\Pi} \in \mathbb{R}^{2\times 5}$ is the economic indices covariance matrix. Combining Equation 3.1 and 3.3 gives the full process under $\mathbb{P}$ in matrix notation:

$$\begin{bmatrix} d\mathbf{X}_t^s \\ d\mathbf{X}_t^o \end{bmatrix} = \begin{bmatrix} K(\mathbb{E}\mathbf{X}_\infty^s - \mathbf{X}_t^s) \\ \mu^o + K^o\mathbf{X}_t^s \end{bmatrix} dt + \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix} \left(\Gamma_0 + (\mathbf{X}_t^s)_1\, \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}} \tag{3.5}$$

Some restrictions are imposed on the model parameters.

- K_{vv} is the speed of the mean reversion at which v_t reverts to the long-run mean of the volatility $\mathbb{E}v_\infty$, hence $K_{vv} > 0$.
- w is by definition the volatility of volatility, so $w > 0$.
- To guarantee that $\forall v_0 > 0$ and $\forall t > 0$: $\mathbb{P}(v_t > 0) = 1$, the Feller condition is imposed: $K_{vv}\mathbb{E}v_\infty - \frac{1}{2}w^2 \geq 0$.
- To ensure that $\mathbf{X}_t^s$ converges to $\mathbb{E}\mathbf{X}_\infty^s$, it is required that $\mathbf{X}_t^s$ is stationary, which only holds if K has real positive eigenvalues (Dai and Singleton [2000]).

- For a (unique) solution $\mathbf{X}_t$ to exist, another restriction has to be enforced. Imposing that Γ_1 has real positive eigenvalues, $\left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)$ will be positive definite and, by the chosen forms of the other matrices, also symmetric. As a result, $\left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)^{\frac{1}{2}}$ can be determined uniquely so that it is also symmetric positive definite (Horn and Johnson [2012]). Would $\left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)$ not be symmetric positive definite, its square root computation could lead to multiple solutions.

It is clear that in $\mathbb{P}$, the dynamics in Equation 3.5 belongs to the affine class (Definition 2.1.4). We define $\Sigma = \begin{bmatrix}\Sigma^{r\pi} & \Sigma^{SII}\end{bmatrix}'$ and by taking

$$a_0 = \begin{bmatrix} K\mathbb{E}\mathbf{X}_\infty^s \\ \mu^0 \end{bmatrix}, \quad a_1 = \begin{bmatrix} -K \\ K^0 \end{bmatrix}, \quad c_0 = \Sigma\Gamma_0\Sigma', \quad c_1^{(1)} = \Sigma\Gamma\Sigma',$$
$$c_1^{(2)} = c_1^{(3)} = c_1^{(4)} = c_1^{(5)} = 0_{5\times 5}, \quad r_0 = 0, \quad r_1' = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

shows the dynamics $\mathbf{X}_t$ belong to the class of affine models.

3.2 TRANSFORMATION FROM $\mathbb{P}$ TO $\mathbb{Q}$ UNDER CONSTANT MARKET PRICES OF RISK

Up till now the dynamics of $\mathbf{X}_t$ was modelled under the real-world measure $\mathbb{P}$. However, since $\mathbb{P}$ depends on investors' judgement of the market, it is impossible to calculate unique option prices. Under the assumption of no arbitrage, we know from the Fundamental Theorem of Asset Pricing 1 (Delbaen and Schachermayer [1998]) that there exists an equivalent martingale measure $\mathbb{Q} \sim \mathbb{P}$. In this section we illustrate how to find $\mathbb{Q}$.

The dynamics described by the CP2022 model includes an unobservable state process which are not directly tradeable. As a result, the economy market model is incomplete. Hence, for every source of uncertainty, a market price of risk Λ_t must be introduced, reflecting investors' risk tolerance. By definition, the market price of risk (also referred to as the Sharpe ratio in asset pricing terminology) is the ratio of the risk premium (the excess return over the risk-free rate) and volatility:

$$\Sigma\mathbf{\Lambda}_t = \eta. \tag{3.6}$$

Because the market model is incomplete, the market prices of risk Λ_t are not unique. Hence the transformation from $\mathbb{P}$ to $\mathbb{Q}$ is not unique either. As a result, the market prices of risk have to be determined by calibrating the generated yields by the model to the empirical yield curve observed in the market (Dai and Singleton [2003]).

To find the risk-neutral measure transformation, we have to choose its associated numéraire, which in this case is the money market account. We illustrate how to find the transformation from $\mathbb{P}$ to $\mathbb{Q}$ for a given Λ_t in the general case, to understand the transformation used in the CP2022 model.

Suppose $d\mathbf{X}_t = \boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X}_t)dt + \boldsymbol{\sigma}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X}_t)d\mathbf{W}_t$ is a stochastic process that drives the economy and the price of a zero-coupon bond is $P(t,T) = P(t,T,\mathbf{X}_t)$. Then applying Itô's lemma (Lemma 2.1.10) gives

$$\frac{dP(t,T)}{P(t,T)} = \mu^{\mathbb{P}}dt + \sigma^{\mathbb{P}}dW_t^{\mathbb{P}},$$

with

$$\begin{aligned}\mu^{\mathbb{P}} &= \frac{1}{P(t,T)}\left(\frac{\partial P(t,T,\mathbf{X})}{\partial t} + \boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X})'\frac{\partial P(t,T,\mathbf{X})}{\partial \mathbf{X}} + \frac{1}{2}\mathrm{Tr}\left[\boldsymbol{\sigma}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X})\boldsymbol{\sigma}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X})'\frac{\partial^2 P(t,T,\mathbf{X})}{\partial \mathbf{X}\partial \mathbf{X}'}\right]\right) \\ \sigma^{\mathbb{P}} &= \frac{1}{P(t,T)}\frac{\partial P(t,T,\mathbf{X})}{\partial \mathbf{X}'}\boldsymbol{\sigma}_{\mathbf{X}}^{\mathbb{P}}(t,\mathbf{X})\end{aligned} \tag{3.7}$$

Then, by definition of the market prices of risks, we have

$$\eta \equiv \mu^{\mathbb{P}} - r_t = \sigma^{\mathbb{P}} \Lambda_t. \tag{3.8}$$

Plugging the obtained parameters from Equation 3.7 in Equation 3.8 gives the PDE:

$$\frac{\partial P}{\partial t} + (\boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{P}} - \sigma_{\mathbf{X}}^{\mathbb{P}} \boldsymbol{\Lambda}_t)' \frac{\partial P}{\partial X} + \frac{1}{2} \mathrm{Tr}\left[\sigma_{\mathbf{X}}^{\mathbb{P}}(t, \mathbf{X}) \sigma_{\mathbf{X}}^{\mathbb{P}}(t, \mathbf{X})' \frac{\partial^2 P(t, T, \mathbf{X})}{\partial \mathbf{X} \partial \mathbf{X}'} \right] = r(t) P(t, T).$$

Additionally, a direct application of Feynman-Kac (Theorem 2.3.11) to $P(t, T)$ in eq. (2.1) gives the PDE:

$$\frac{\partial P}{\partial t} + \boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{Q}\prime} \frac{\partial P}{\partial X} + \frac{1}{2} \mathrm{Tr}\left[\sigma_{\mathbf{X}}^{\mathbb{Q}}(t, \mathbf{X}) \sigma_{\mathbf{X}}^{\mathbb{Q}}(t, \mathbf{X})' \frac{\partial^2 P(t, T, \mathbf{X})}{\partial \mathbf{X} \partial \mathbf{X}'} \right] = r(t) P(t, T).$$

By comparing the results, we obtain the required transformations:

$$\begin{aligned} \boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{Q}} &= \boldsymbol{\mu}_{\mathbf{X}}^{\mathbb{P}} - \sigma_{\mathbf{X}}^{\mathbb{P}} \boldsymbol{\Lambda}_t, \\ \sigma_{\mathbf{X}}^{\mathbb{Q}} &= \sigma_{\mathbf{X}}^{\mathbb{P}}. \end{aligned} \tag{3.9}$$

Ultimately, this shows that the diffusion remains the same, but in the risk-neutral setting, the drift is adjusted so that risk premiums are absent.

In CP2022, the market prices of risk Λ_t with $\Lambda_1 \in \mathbb{R}^{5\times 3}$ and $\lambda_0 \in \mathbb{R}^5$ are assumed to take the form

$$\Lambda_t = \lambda_0 + \Lambda_1 \mathbf{X}_t^s. \tag{3.10}$$

They are affine in $\mathbf{X}_t^s$, to ensure that nominal zero-coupon bond prices are exponentially affine in the state variables under both measures.

The model assumes the equity risk premium η_S and the risk premium on inflation η_Π does not change over time. This translates to the following constraints:

$$\Sigma^{S\Pi} \lambda_0 = \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} \tag{3.11}$$

$$\Sigma^{S\Pi} \Lambda_1 = 0_{2\times 3} \tag{3.12}$$

The first constraint follows from the definition of market prices of risks (Equation 3.6) and the second constraint follows from the assumption that the risk premiums are constant for any $\mathbf{X}_t$.

In CP2022, the kernel defining the transformation that preserves the affine structure and ensures $\mathbb{Q} \sim \mathbb{P}$ is given by

$$d\mathbf{W}_t^{\mathbb{P}} = d\mathbf{W}_t^{\mathbb{Q}} - \left((\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}} \right)^{-1} (\lambda_0 + \Lambda_1 \mathbf{X}_t^s)\, dt. \tag{3.13}$$

Then, the dynamics of $\mathbf{X}_t^o$ in $\mathbb{Q}$ becomes:

$$\begin{aligned} d\mathbf{X}_t^o &= (\mu^o + K^o \mathbf{X}_t^s)\, dt + \Sigma^{S\Pi} (\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}}\, d\mathbf{W}_t^{\mathbb{P}} \\ &= (\mu^o + K^o \mathbf{X}_t^s)\, dt + \Sigma^{S\Pi} (\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}} \left(d\mathbf{W}_t^{\mathbb{Q}} - \left((\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}} \right)^{-1} (\lambda_0 + \Lambda_1 \mathbf{X}_t^s)\, dt \right) \\ &= \left(\mu^o - \Sigma^{S\Pi} \Lambda_t + K^o \mathbf{X}_t^s \right) dt + \Sigma^{S\Pi} (\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}}\, d\mathbf{W}_t^{\mathbb{Q}} \\ &\overset{(3.11),(3.12)}{=} \left(\mu^o - \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} + K^o \mathbf{X}_t^s \right) dt + \Sigma^{S\Pi} (\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma)^{\frac{1}{2}}\, d\mathbf{W}_t^{\mathbb{Q}}. \end{aligned}$$

As we have seen in the general risk-neutral setting, risk premiums must be absent in the drift. This is indeed the case here, as the risk premiums in the drift get subtracted from μ^0 in Equation 3.4. This shows that the transformation as proposed is consistent with the definition of market prices of risk and the dynamics are risk-neutral.

Similarly, the dynamics of $\mathbf{X}_t^s$ in $\mathbb{Q}$ is

$$\begin{aligned} d\mathbf{X}_t^s &= K\left(\mathbb{E}\mathbf{X}_\infty^s - \mathbf{X}_t^s\right) dt + \Sigma^{r\pi}\left(\Gamma_0 + \left(\mathbf{X}_t^s\right)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{P}} \\ &= K\left(\mathbb{E}\mathbf{X}_\infty^s - \mathbf{X}_t^s\right) dt + \Sigma^{r\pi}\left(\Gamma_0 + \left(\mathbf{X}_t^s\right)_1 \Gamma\right)^{\frac{1}{2}} \left(d\mathbf{W}_t^{\mathbb{Q}} - \left(\left(\Gamma_0 + \left(\mathbf{X}_t^s\right)_1 \Gamma\right)^{\frac{1}{2}}\right)^{-1} \left(\lambda_0 + \Lambda_1 \mathbf{X}_t^s\right) dt\right) \\ &= \left(K\mathbb{E}\mathbf{X}_\infty^s - \Sigma^{r\pi}\lambda_0 - \left(K + \Sigma^{r\pi}\Lambda_1\right)\mathbf{X}_t^s\right) dt + \Sigma^{r\pi}\left(\Gamma_0 + \left(\mathbf{X}_t^s\right)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{Q}} \end{aligned}$$

To ensure the dynamics in $\mathbb{Q}$ can be written in the same form as in $\mathbb{P}$, we define a risk-neutral matrix M of K as:

$$M = K + \Sigma^{r\pi}\Lambda_1. \tag{3.14}$$

We denote the steady state vector $\mathbb{E}\mathbf{X}_\infty^s$ under $\mathbb{Q}$ by $\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s$ and denote the drift as $M(\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s - \mathbf{X}_\infty^s)$. Then, we require
$M(\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s - \mathbf{X}_\infty^s) = (K\mathbb{E}\mathbf{X}_\infty^s - \Sigma^{r\pi}\lambda_0 - (K + \Sigma^{r\pi}\Lambda_1)\mathbf{X}_t^s)$, which yields the condition

$$\Sigma^{r\pi}\lambda_0 = -M\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s + K\mathbb{E}\mathbf{X}_\infty^s \tag{3.15}$$

that must hold to keep $\mathbb{P}$ and $\mathbb{Q}$ consistent under the chosen kernel. To ensure that the dynamics are equivalent in $\mathbb{P}$ and $\mathbb{Q}$, the following constraints are imposed

- K and M have the same structure. Since $K_{rv} = K_{\pi v} = 0$, we get $M_{rv} = M_{\pi v} = 0$ by enforcing $(\Lambda_1)_{1,2} = (\Lambda_1)_{1,3} = 0$. This follows from Equation 3.14 and the structure of $\Sigma^{r\pi}$ as in Equation 3.2.

- M has real and positive eigenvalues, like K.

- the Feller condition must also hold in $\mathbb{Q}$, i.e. we get $\mathbb{Q}(v_t > 0) = 1$ if $M_{vv}\mathbb{E}^{\mathbb{Q}}v_\infty - \frac{1}{2}\omega^2 \geq 0$.

The dynamics in $\mathbb{Q}$ can be described as:

$$\begin{bmatrix} d\mathbf{X}_t^s \\ d\mathbf{X}_t^o \end{bmatrix} = \begin{bmatrix} M\left(\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s - \mathbf{X}_t^s\right) \\ \mu^o - \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} + K^o\mathbf{X}_t^s \end{bmatrix} dt + \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix} \left(\Gamma_0 + \left(\mathbf{X}_t^s\right)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{Q}}. \tag{3.16}$$

The model dynamics under both $\mathbb{P}$ and $\mathbb{Q}$ are complete, assuming constant market prices of risk.

In the next section, we discuss the term structure of interest rates, which plays a major part for calibrating the model.

3.3 TERM STRUCTURE OF INTEREST RATES UNDER CP2022

The key component that pension funds use for valuing liabilities and assets is the nominal and real interest-rate term structure. Both curves are the basis for determining the discount factor to calculate the present value for future cash flows. The real interest rates are adjusted for inflation,

compared to the nominal rates. As a result, the primary focus during calibration is to ensure that the yield of nominal and real zero-coupon bonds align with those quoted on the market.

An important property of affine term-structure models is that they can price zero-coupon bonds analytically (Duffie et al. [2000]), where yields can be inferred as in Definition 2.4.4. Additionally, under some technicality conditions, the characteristic function of $\mathbf{X}_t$ can be readily obtained, if $\mathbf{X}_t$ is of the affine class, by solving a system of Riccati differential equations. Via the same reasoning as in Duffie et al. [2000], the results are easily extended when the drift and diffusion in $\mathbf{X}_t$ are time-dependent.

For the purpose of clarity, we write $\phi(t,T,u,v) = \phi_{uv}(t,T)$, $\Psi(t,T,u,v) = \Psi_{uv}(t,T)$, $\mathbf{X}_t^i = (\mathbf{X}_t)_i$ and $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathcal{F}_t]$.

Theorem 3.3.1 (Characteristic function of an affine process)**.** *If an n-dimensional process $\mathbf{X}_t$ satisfies*

$$d\mathbf{X}_t = (\zeta(t) - L\mathbf{X}_t)\,dt + \Sigma\left(G_0 + \sum_{i=1}^{n} G_i \mathbf{X}_t^i\right)^{\frac{1}{2}} d\mathbf{W}_t, \tag{3.17}$$

with $\mathbf{W}_t$ a standard n-dimensional Brownian Motion (under any measure) and with L, Σ, G_0, $G_i \in \mathbb{R}^{n\times n}$ and deterministic function $\zeta(t) : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$. Then $\forall t \leq T$ and $u, v \in \mathbb{C}^n$:

$$\mathbb{E}_t\left[e^{u'\mathbf{X}_T + v'\int_t^T \mathbf{X}_s ds}\right] = e^{\phi_{uv}(t,T) + \Psi_{uv}(t,T)'\mathbf{X}_t} \tag{3.18}$$

with deterministic functions $\phi_{uv}(t,T) : \mathbb{C}^n \times \mathbb{C}^n \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{C}$ and $\Psi_{uv}(t,T) : \mathbb{C}^n \times \mathbb{C}^n \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{C}^n$ that solve

$$\frac{\partial}{\partial t}\Psi_{uv}(t,T)_i = -\frac{1}{2}\Psi_{uv}(t,T)'\Sigma G_i \Sigma' \Psi_{uv}(t,T) + \left(L'\Psi_{uv}(t,T)\right)_i - v_i, \quad \Psi_{uv}(T,T) = u \tag{3.19}$$

$$\phi_{uv}(t,T) = \int_t^T \left(\Psi_{uv}(s,T)'\zeta(s) + \frac{1}{2}\Psi_{uv}(s,T)'\Sigma G_0 \Sigma' \Psi_{uv}(s,T)\right) ds, \tag{3.20}$$

and the additional technical conditions hold:

1. $\mathbb{E}\left|e^{u'\mathbf{X}_T + v'\int_0^T \mathbf{X}_s ds}\right| < \infty$,
2. $\mathbb{E}\left[\left(\int_0^T e^{\phi_{uv}(t,T) + \Psi_{uv}(t,T)'\mathbf{X}_t} \Psi_{uv}(t,T)'\Sigma\left(G_0 + \sum_{i=1}^n G_i \mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T) e^{\phi_{uv}(t,T)' + \mathbf{X}_t'\Psi_{uv}(t,T)} dt\right)^{\frac{1}{2}}\right] < \infty$.

Proof. Note that Equation 3.17 is affine and thus a similar approach can be used as in Duffie et al. [2000]. We define

$$Z_t = \mathbb{E}_t\left[e^{u'\mathbf{X}_T + v'\int_0^T \mathbf{X}_s ds}\right] = e^{\phi_{uv}(t,T) + \Psi_{uv}(t,T)'\mathbf{X}_t + v'\int_0^t \mathbf{X}_s ds}. \tag{3.21}$$

It is enough to prove that Z_t is a martingale, since then

$$\mathbb{E}_t[Z_T] = \mathbb{E}_t\left[e^{\phi_{uv}(T,T) + \Psi_{uv}(t,T)'\mathbf{X}_T + v'\int_0^T \mathbf{X}_s ds}\right] \overset{(*)}{=} \mathbb{E}_t\left[e^{u'\mathbf{X}_T + v'\int_0^T \mathbf{X}_s ds}\right] = Z_t,$$

where in $(*)$ we impose the terminal conditions of the ODEs, which equal $\phi_{uv}(T,T) = 0$ and $\Psi_{uv}(T,T) = u$.

We apply Multivariate Itô's lemma to Z_t:

$$\begin{aligned}
dZ_t &= \frac{\partial Z}{\partial t}dt + \frac{\partial Z}{\partial \mathbf{X}'}d\mathbf{X} + \frac{1}{2}d\mathbf{X}'\frac{\partial^2 Z}{\partial \mathbf{X}\partial \mathbf{X}'}d\mathbf{X} \\
&= Z_t\left(\frac{\partial}{\partial t}\phi_{uv}(t,T) + \frac{\partial}{\partial t}\Psi_{uv}(t,T)'\mathbf{X}_t + v'\mathbf{X}_t\right)dt \\
&+ Z_t\Psi_{uv}(t,T)'\left((\zeta(t) - L\mathbf{X}_t)\,dt + \Sigma\left(G_0 + \sum_{i=1}^n G_i\mathbf{X}_t^i\right)^{\frac{1}{2}}d\mathbf{W}_t\right) \\
&+ \frac{1}{2}Z_t\Psi_{uv}(t,T)'\Sigma\left(G_0 + \sum_{i=1}^n G_i\mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T)dt \\
&= Z_t\left(\frac{\partial}{\partial t}\phi_{uv}(t,T) + \frac{\partial}{\partial t}\Psi_{uv}(t,T)'\mathbf{X}_t + v'\mathbf{X}_t + \Psi_{uv}(t,T)'\left(\zeta(t) - L\mathbf{X}_t\right)\right. \\
&\left.+ \frac{1}{2}\Psi_{uv}(t,T)'\Sigma\left(G_0 + \sum_{i=1}^n G_i\mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T)\right)dt + Z_t\Psi_{uv}(t,T)'\Sigma\left(G_0 + \sum_{i=1}^n G_i\mathbf{X}_t^i\right)^{\frac{1}{2}}d\mathbf{W}_t.
\end{aligned}$$

Note that Z_t is a martingale, by the Martingale Representation Theorem (see Theorem 2.1.11), if and only if dZ_t is driftless. This is true if dZ_t has no dt terms:

$$\frac{\partial}{\partial t}\phi_{uv}(t,T) + \frac{\partial}{\partial t}\Psi_{uv}(t,T)'\mathbf{X}_t + v'\mathbf{X}_t + \Psi_{uv}(t,T)'\left(\zeta(t) - L\mathbf{X}_t\right) + \frac{1}{2}\Psi_{uv}(t,T)'\Sigma\left(G_0 + \sum_{i=1}^n G_i\mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T) = 0$$

$$\iff$$

$$\begin{aligned}
\frac{\partial}{\partial t}\phi_{uv}(t,T) + \frac{\partial}{\partial t}\Psi_{uv}(t,T)'\mathbf{X}_t =& -\Psi_{uv}(t,T)'\zeta(t) - \frac{1}{2}\Psi_{uv}(t,T)'\Sigma G_0\Sigma'\Psi_{uv}(t,T) \\
&-v'\mathbf{X}_t + \Psi_{uv}(t,T)'L\mathbf{X}_t - \frac{1}{2}\Psi_{uv}(t,T)'\Sigma\left(\sum_{i=1}^n G_i\mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T).
\end{aligned}$$

Note that after the equivalence, the right-hand side on the first line is independent of $\mathbf{X}_t$ and the remaining part depends on $\mathbf{X}_t$. This means that we get the ODEs:

$$\begin{aligned}
\frac{\partial}{\partial t}\phi_{uv}(t,T) &= -\Psi_{uv}(t,T)'\zeta(t) - \frac{1}{2}\Psi_{uv}(t,T)'\Sigma G_0\Sigma'\Psi_{uv}(t,T) \\
\frac{\partial}{\partial t}\Psi_{uv}(t,T)'\mathbf{X}_t &= -v'\mathbf{X}_t + \Psi_{uv}(t,T)'L\mathbf{X}_t - \frac{1}{2}\Psi_{uv}(t,T)'\Sigma\left(\sum_{i=1}^n G_i\mathbf{X}_t^i\right)\Sigma'\Psi_{uv}(t,T).
\end{aligned}$$

Enforcing $\Psi_{uv}(T,T) = u$ and $\phi_{uv}(T,T) = 0$ such that Equation 3.21 holds and matching components for each $\mathbf{X}_t^i$ in the second equation gives the desired result. Because $e^{v'\int_0^t \mathbf{X}_s ds}$ is $\mathcal{F}_t$-measurable, from the definition of Z_t in Equation 3.21 we get

$$e^{v'\int_0^t \mathbf{X}_s ds}\mathbb{E}_t\left[e^{u'\mathbf{X}_T + v'\int_t^T \mathbf{X}_s ds}\right] = e^{\phi_{uv}(t,T) + \Psi_{uv}(t,T)'\mathbf{X}_t + v'\int_0^t \mathbf{X}_s ds}.$$

This gives directly Equation 3.18. Finally, remark that the two technical conditions are required for Z_t to be a proper martingale. Indeed, the first condition coincides with the integrability condition for martingales in Definition 2.1.2 and the second condition (combined with Jensen's inequality) is equivalent to the condition in Theorem 2.1.11, which ensures that the Itô integral (the diffusion of dZ_t) is a martingale, see Theorem 2.1.8. □

The yield of a nominal zero-coupon bond at time t with time with maturity T and time to maturity $\tau = T - t$ is

$$\begin{aligned} y_t(\tau) &= -\tau^{-1} \ln P(t,T) = -\tau^{-1} \ln \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r_u du} \right] = -\tau^{-1} \ln \mathbb{E}_t^{\mathbb{Q}} \left[e^{\begin{bmatrix} 0 & -1 & 0 \end{bmatrix} \int_t^{t+\tau} \mathbf{x}_u^s du} \right] \\ &= -\tau^{-1} \left(\phi(t, t+\tau) + \Psi(t, t+\tau)' \mathbf{X}_t^s \right). \end{aligned} \tag{3.22}$$

The deterministic functions ϕ and Ψ solve the Riccati differential equations in Theorem 3.3.1 with the parameters that characterize the dynamics of $\mathbf{X}_t^s$ in $\mathbb{Q}$ as in Equation 3.16:

$$\begin{aligned} &L = M, \quad \zeta(t) = M \mathbb{E}^{\mathbb{Q}} \mathbf{X}_\infty^s, \quad \Sigma = \Sigma^{r\pi}, \quad u' = \begin{bmatrix} 0 & 0 & 0 \end{bmatrix}, \quad v' = \begin{bmatrix} 0 & -1 & 0 \end{bmatrix} \\ &G_0 = \Gamma_0, \quad G_1 = \Gamma, \quad G_2 = G_3 = G_4 = G_5 = 0_{5\times 5}. \end{aligned} \tag{3.23}$$

Real yields satisfy

$$\begin{aligned} y_t^R(\tau) &= -\tau^{-1} \ln \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^T r_u du + (\ln \Pi_T - \ln \Pi_t)} \right] = -\tau^{-1} \ln \mathbb{E}_t^{\mathbb{Q}} \left[e^{\int_t^{t+\tau} \begin{bmatrix} 0 & -1 & 0 \end{bmatrix} \mathbf{x}_u^s du + \begin{bmatrix} 0 & 1 \end{bmatrix} (\mathbf{x}_{t+\tau}^o - \mathbf{x}_t^o)} \right] \\ &= -\tau^{-1} \left(\phi_R(t, t+\tau) + \Psi_R^s(t, t+\tau)' \mathbf{X}_t^s + \left(\Psi_R^o(t, t+\tau)' - \begin{bmatrix} 0 & 1 \end{bmatrix} \right) \mathbf{X}_t^o \right), \end{aligned} \tag{3.24}$$

for ϕ_R, Ψ_R^s and Ψ_R^o that solve the Riccati differential equations in Theorem 3.3.1 with parameters

$$\begin{aligned} &L = \begin{bmatrix} M & 0_{3\times 2} \\ -K^0 & 0_{2\times 2} \end{bmatrix}, \quad \zeta(t) = \begin{bmatrix} M \mathbb{E}^{\mathbb{Q}} \mathbf{X}_\infty^s \\ \mu^o - \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix}, \quad u' = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \\ &v' = \begin{bmatrix} 0 & -1 & 0 & 0 & 0 \end{bmatrix}, \quad G_0 = \Gamma_0, \quad G_1 = \Gamma, \quad G_2 = G_3 = G_4 = G_5 = 0_{5\times 5}. \end{aligned} \tag{3.25}$$

The analytic solution for the real yields can be simplified further, since for the chosen model parameters in Equation 3.19, we have

$$\Psi_R^o(t, t+\tau)' = [0 \quad 1]. \tag{3.26}$$

As a result,

$$y_t^R(\tau) = -\tau^{-1} \left(\phi_R(t, t+\tau) + \Psi_R^s(t, t+\tau)' \mathbf{X}_t^s \right). \tag{3.27}$$

In the next section we explore how the model fits the term structure of interest rates.

3.4 CALIBRATION OF THE TERM STRUCTURE OF INTEREST RATES WITH TIME-VARYING MARKET PRICES OF RISK

Since affine models can only replicate a limited number of types of term structures of interest rates, adjustment functions are introduced which bridge the gap between bond prices generated by constant and time-varying market prices of risk. Up till now market prices of risk were assumed to be constant for simplicity. Though in practice, this is not realistic since the market prices of risk are influenced by many factors. Economic conjuncture, market volatility, investor sentiment, and liquidity fluctuate over time, affecting the compensation investors demand for bearing risk. Hence, CP2022 extends the KNW model by introducing time-dependent risk premiums which arise from assuming time-dependent market prices of risk. For fitting the term structure of interest rates at the time of calibration t_0, the market prices of risk are assumed to be constant for $t < t_0$ and time-varying for $t \geq t_0$:

$$\tilde{\lambda}_0(t) = \lambda_0 + \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix}^{-1} \begin{bmatrix} 0 & f(t) & 0 & 0 & f^R(t) \end{bmatrix} \mathbb{1}_{\{t \geq t_0\}}, \tag{3.28}$$

where f and f^R are the adjustment functions. Note that because Commission Parameters prioritizes fitting to the yield curves at the time of calibration, the market prices of risk are assumed to be constant when determining historic bond prices. The deterministic adjustment functions adjust the short rate r_t and the European consumer price index $\ln \Pi_t$, which are the risk factors that determine the nominal and real term structures.

The adjustment functions f and f^R are assumed to be piecewise constant functions over a region Δ: $f(t) = f_i$ for $t \in [i\Delta, (i+1)\Delta)$. The functions will be determined using the nominal and real bond prices generated by constant and time-varying market prices of risk. Choosing Δ involves a trade-off between the resulting smoothness of the adjustment functions and the time complexity of calibration, as these values need to be recalculated at each step of the optimization process. Hence monthly time steps $\Delta = 1/12$ are used, and the corresponding nominal and real yields observed in the market are determined through linear interpolation to obtain intermediate values between observations. For calculating the adjustment functions, nominal zero-coupon rates with yearly maturities ranging from 1 to 50 years are used and real zero-coupon rates with maturities {1-10, 12, 15, 20, 25, 30, 40, 50} are used. Extrapolation for maturities beyond these is performed by assuming that the yearly forward rates equal a constant forward rate, referred to as the Ultimate Forward Rate, which we discuss in Equation 3.5.2.

As a result of making the market prices of risk time-varying, the dynamics in $\mathbb{Q}$ change compared to the case with constant prices of risk ($f = f^R = 0$). Using $\tilde{\lambda}_0(t)$ as defined in Equation 3.28 into the defined transformation from $\mathbb{P}$ to $\mathbb{Q}$ in Equation 3.13 gives newly generated $\mathbb{Q}$-dynamics:

$$\begin{aligned} d\mathbf{X}_t^s &= M\left(\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s - \mathbf{X}_t^s\right) dt + \Sigma^{r\pi} \left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{Q}} - [0 \quad 1 \quad 0]' f(t) \mathbb{1}_{\{t \geq t_0\}} dt, \\ d\mathbf{X}_t^o &= \left(\mu^o - \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} + K^o \mathbf{X}_t^s\right) dt + \Sigma^{S\Pi} \left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{Q}} - [0 \quad 1]' f^R(t) \mathbb{1}_{\{t \geq t_0\}} dt. \end{aligned}$$

The combined $\mathbb{Q}$-dynamics are:

$$\begin{bmatrix} d\mathbf{X}_t^s \\ d\mathbf{X}_t^o \end{bmatrix} = \begin{bmatrix} M\left(\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s - \mathbf{X}_t^s\right) - \begin{bmatrix} 0 \\ f(t) \\ 0 \end{bmatrix} \\ \mu^o + K^o \mathbf{X}_t^s - \begin{bmatrix} \eta_S + 0 \\ \eta_\Pi + f^R(t) \end{bmatrix} \end{bmatrix} dt + \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix} \left(\Gamma_0 + (\mathbf{X}_t^s)_1 \Gamma\right)^{\frac{1}{2}} d\mathbf{W}_t^{\mathbb{Q}}. \tag{3.29}$$

Denote $p_{t_0}(\tau)$ and $\tilde{p}_{t_0}(\tau)$ by nominal bond prices generated by the market prices of risk λ_0 and $\tilde{\lambda}_0(t)$ respectively, and $p_{t_0}^R(\tau)$ and $\tilde{p}_{t_0}^R(\tau)$ for the corresponding real bond prices.

Comparing the dynamics generated by $\tilde{\lambda}_0(t)$ in Equation 3.29 with those generated by λ_0 in Equation 3.16 affects only one model parameter when determining the nominal and real yields: $\tilde{\zeta}(t) = \zeta(t) - [0 \quad f(t) \quad 0]' \mathbb{1}_{\{t \geq t_0\}}$ for the nominal case and $\tilde{\zeta}(t) = \zeta(t) - [0 \quad f(t) \quad 0 \quad 0 \quad f^R(t)]' \mathbb{1}_{\{t \geq t_0\}}$ for the real case. Denote $\widetilde{\phi}_{uv}(t,T)$ and $\widetilde{\Psi}_{uv}(t,T)$ by the solutions to the Riccati differential equations with $\tilde{\zeta}(t)$, and note that $\widetilde{\Psi}_{uv}(t,T) = \Psi_{uv}(t,T) = \Psi(\tau)$. For constructing $f(t)$ and $f^R(t)$, the Leibniz integration rule (see Theorem 2.4.8) is used.

Using Equation 3.22 and 3.20 results in

$$\ln \frac{\tilde{p}_{t_0}(\tau)}{p_{t_0}(\tau)} = \tilde{\phi}\left(t_0, t_0+\tau\right) - \phi\left(t_0, t_0+\tau\right) = -\int_{t_0}^{t_0+\tau} \Psi\left(s, t_0+\tau\right)' \begin{bmatrix} 0 & f(s) & 0 \end{bmatrix}' ds, \tag{3.30}$$

$$\frac{\partial}{\partial \tau} \ln \frac{\tilde{p}_{t_0}(\tau)}{p_{t_0}(\tau)} = -\frac{\partial}{\partial \tau} \int_{t_0}^{t_0+\tau} \Psi\left(t_0+\tau-s\right)_2 f(s) ds = -\int_{t_0}^{t_0+\tau} \dot{\Psi}\left(t_0+\tau-s\right)_2 f(s) ds, \tag{3.31}$$

where the Leibniz integration rule is used with $\Psi(0)_2 = 0$, which follows from the chosen u in Equation 3.23 and the terminal condition in Equation 3.19. For determining f_i, Equation 3.31 is used, which relates the derivative of the log differences between the observed market prices $\tilde{p}_{t_0}$ and model prices p_{t_0} to the original Ψ terms. Approximating the derivative by finite differences on interval $[t_0 + i\Delta, t_0 + (i+1)\Delta]$ where $f(t) = f_i$ gives

$$\begin{aligned} E_i \equiv \Delta^{-1} \ln \frac{\tilde{p}_{t_0}((i+1)\Delta)/\tilde{p}_{t_0}(i\Delta)}{p_{t_0}((i+1)\Delta)/p_{t_0}(i\Delta)} &= -\int_{t_0}^{t_0+(i+1)\Delta} \dot{\Psi}\left(t_0+(i+1)\Delta-s\right)_2 f(s) ds \\ &= -\sum_{j=0}^{i} \int_{t_0+j\Delta}^{t_0+(j+1)\Delta} \dot{\Psi}\left(t_0+(i+1)\Delta-s\right)_2 f(s) ds \\ &= -\sum_{j=0}^{i} f_j \int_{t_0+j\Delta}^{t_0+(j+1)\Delta} \dot{\Psi}\left(t_0+(i+1)\Delta-s\right)_2 ds \\ &= \sum_{j=0}^{i} f_j (\Psi((i-j)\Delta) - \Psi((i+1-j)\Delta))_2. \end{aligned}$$

Now f_i can be obtained recursively as:

$$\begin{aligned} f_0 &= -\frac{E_0}{\Psi(\Delta)_2} \\ f_i &= \frac{-E_i}{\Psi(\Delta)_2} + \sum_{j=0}^{i-1} f_j \frac{\Psi((i-j)\Delta)_2 - \Psi((i+1-j)\Delta)_2}{\Psi(\Delta)_2}, \quad \forall i > 0. \end{aligned}$$

The real adjustment function f^R can be constructed in a similar fashion as in the nominal case, where we use the Leibniz integration rule and note that $\Psi^R_5 = 1$ (Equation 3.26):

$$\ln \frac{\tilde{p}^R_{t_0}(\tau)}{p^R_{t_0}(\tau)} = \tilde{\phi}^R\left(t_0, t_0+\tau\right) - \phi^R\left(t_0, t_0+\tau\right) = -\int_{t_0}^{t_0+\tau} \Psi^R\left(s, t_0+\tau\right)' \begin{bmatrix} 0 & f(s) & 0 & 0 & f^R(s) \end{bmatrix}' ds$$

$$\frac{\partial}{\partial \tau} \ln \frac{\tilde{p}^R_{t_0}(\tau)}{p^R_{t_0}(\tau)} = -\int_{t_0}^{t_0+\tau} \dot{\Psi}^R\left(t_0+\tau-s\right)_2 f(s) ds - f^R\left(t_0+\tau\right),$$

and analogous to the nominal case

$$E^R_i \equiv \Delta^{-1} \ln \frac{\tilde{p}^R_{t_0}((i+1)\Delta)/\tilde{p}^R_{t_0}(i\Delta)}{p^R_{t_0}((i+1)\Delta)/p^R_{t_0}(i\Delta)} = -f^R_i + \sum_{j=0}^{i} f_j \left(\Psi^R((i-j)\Delta) - \Psi^R((i+1-j)\Delta)\right)_2.$$

Again, f^R_i can be obtained recursively:

$$\begin{aligned} f^R_0 &= -\frac{E^R_0}{1+\Psi^R(\Delta)_2} \\ f^R_i &= \frac{-E^R_i}{1+\Psi^R(\Delta)_2} + \sum_{j=0}^{i-1} f^R_j \frac{\Psi^R((i-j)\Delta)_2 - \Psi^R((i+1-j)\Delta)_2}{1+\Psi^R(\Delta)_2}, \quad \forall i > 0. \end{aligned}$$

3.5 CALIBRATION OF CP2022

There are two distinct assumptions that can be made. The first assumptions states that all yields can be observed with measurement errors. It is reasonable to assume that errors are present in all observed yields. Such errors can arise from data entry mistakes and the interpolation methods used to generate zero-coupon yields. When every yield is subject to measurement error, it is not possible to invert the yields in Equation 3.22 to determine the state vector.

Alternatively, in an n-factor model, one may use observations with at most n different maturities to obtain the state factors by inverting Equation 3.22 and 3.27. When more distinct maturities are used than state factors, then Equation 3.22 and 3.27 cannot hold for all maturities. When adding more maturities, measurement errors need to be introduced. A drawback of this approach is that results may be sensitive to the chosen maturities with no measurement errors.

The Commission Parameters opted for the second approach, based on Joslin et al. [2013]. The choice of which maturity to choose with no measurement errors was based on a combination of two elements. First, the maturity with the most liquid bonds in the market are used, ensuring that measurement errors converge fastest to zero. Additionally, the maturity was chosen with the goal of spanning as much of the term structure as possible.

Assume the price indices S_t, Π_t, the squared volatility v_t and the yields of one nominal and one real bond with maturities τ_N^* and τ_R^* respectively, can be observed with no measurement error. Then, the state $(\mathbf{X}_t^s, \mathbf{X}_t^o) = (v_t, r_t, \pi_t, \ln S_t, \ln \Pi_t)$ can be observed by inverting Equation 3.22 and 3.27 to get

$$\mathbf{X}_t^{s,\text{obs}} = \begin{bmatrix} 1 & 0 & 0 \\ \Psi(t, t+\tau_N^*)' \\ \Psi_R(t, t+\tau_R^*)' \end{bmatrix}^{-1} \begin{bmatrix} v_t^{\text{obs}} \\ -\tau_N^* y_t(\tau_N^*)^{\text{obs}} - \phi(t, t+\tau_N^*) \\ -\tau_R^* y_t(\tau_R^*)^{\text{obs}} - \phi(t, t+\tau_R^*) \end{bmatrix}, \tag{3.32}$$

where for the observed squared volatility in the market, EURO STOXX 50 is used as the VSTOXX index (Stoxx [2024]).

For other maturities, the nominal and real yields $Y_t^{\text{obs}}(\tau)$ and $Y_t^{R,\text{obs}}(\tau)$ respectively are assumed to have independent and identically distributed (iid) measurement errors, where the measurements are collected into vectors $\mathbf{Y}_t^{\text{obs}}(\tau)$ and $\mathbf{Y}_t^{R,\text{obs}}(\tau)$. The corresponding maturity vectors τ and τ_R have length n_y and n_y^R respectively, and stacking the nominal and real yields gives the measurement equation assuming Gaussian errors:

$$\begin{bmatrix} \mathbf{Y}_t^{\text{obs}}(\tau_1) \\ \vdots \\ \mathbf{Y}_t^{\text{obs}}\left(\tau_{\tau_y}\right) \\ \mathbf{Y}_t^{\text{R,obs}}\left(\tau_1^R\right) \\ \vdots \\ \mathbf{Y}_t^{\text{R,obs}}\left(\tau_{n_y^R}^R\right) \end{bmatrix} = \begin{bmatrix} y_t(\tau_1) \\ \vdots \\ y_t\left(\tau_{n_y}\right) \\ y_t^R\left(\tau_1^R\right) \\ \vdots \\ y_t^R\left(\tau_{n_y^R}^R\right) \end{bmatrix} + \epsilon_t^y, \quad \epsilon_t^y \overset{iid}{\sim} N\left(0_{(n_y+n_y^R)\times 1}, \Sigma^y\right), \tag{3.33}$$

where $\Sigma^y \in \mathbb{R}^{(n_y+n_y^R)\times(n_y+n_y^R)}$ is a diagonal matrix with values h_i^2 for $1 \leq i \leq n_y + n_y^R$, where h_i is the standard deviation of the measurement equation error. Another assumptions made was that there is no cross-sectional correlation between the measurement errors. A work that allows cross-sectional correlation between the measurement errors in affine models is de Jong [2000].

Besides measurement errors, model misspecification errors from the assumed dynamics in Equation 3.5 are approximated as:

$$\begin{bmatrix} \mathbf{X}^{s,\text{obs}}_{t+dt} - \mathbf{X}^{s,\text{obs}}_{t} \\ \mathbf{X}^{o,\text{obs}}_{t+dt} - \mathbf{X}^{o,\text{obs}}_{t} \end{bmatrix} = \begin{bmatrix} K\left(\mathbb{E}\mathbf{X}^{s}_{\infty} - \mathbf{X}^{s,\text{obs}}_{t}\right) \\ \mu^{o} + K^{o}\mathbf{X}^{s,\text{obs}}_{t} \end{bmatrix} dt + \epsilon^{\text{so}}_{t} \tag{3.34}$$

with

$$\epsilon^{\text{so}}_{t} \overset{iid}{\sim} N\left(0_{5\times 1}, \Sigma^{\text{so}}_{t}\right) \quad \left(\Sigma^{\text{so}}_{t}\right)^{\frac{1}{2}} = \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix} \left(\Gamma_0 + \left(\mathbf{X}^{s,\text{obs}}_{t}\right)_1 \Gamma\right)^{\frac{1}{2}} \sqrt{dt}.$$

Before formulating the log-likelihood function that needs to be maximized, consider the following.

Suppose $\mathbf{X} = (X_1, ..., X_d)$ is iid and follows a multivariate Normal distribution $\mathbf{X} \sim \mathcal{N}_d(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} \in \mathbb{R}^d$ with $\mu_i = \mathbb{E}[X_i]$ is the mean vector and $\boldsymbol{\Sigma} \in \mathbb{R}^{d\times d}$ with $\boldsymbol{\Sigma}_{i,j} = \mathbb{E}[(X_i - \mu_i)((X_j - \mu_j))]$ is the covariance matrix. The density function equals

$$f_{\mathbf{X}}(x_1, ..., x_d) = (2\pi)^{-d/2}(\det(\boldsymbol{\Sigma}))^{-1/2} \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{x} - \boldsymbol{\mu})\right).$$

Then, the log-likelihood is

$$\ln L(\boldsymbol{x}) = -\frac{1}{2}[d\ln(2\pi) + \ln(\det\boldsymbol{\Sigma}) + (\boldsymbol{x} - \boldsymbol{\mu})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{x} - \boldsymbol{\mu})].$$

Maximum likelihood estimation is based on finding $\bar{}$ and $\boldsymbol{\Sigma}$ that maximize the log-likelihood. If we assume $\mathbf{X} \sim \mathcal{N}_d(\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X)$ and $\mathbf{Y} \sim \mathcal{N}_d(\boldsymbol{\mu}_Y, \boldsymbol{\Sigma}_Y)$ and we assume $\mathbf{X}$ and $\mathbf{Y}$ are independent, then $f_{\mathbf{X},\mathbf{Y}} = f_{\mathbf{X}} f_{\mathbf{Y}}$ and thus

$$\begin{aligned} \ln L(\boldsymbol{x}, \boldsymbol{y}) = -\frac{1}{2}[&(\boldsymbol{x} - \boldsymbol{\mu}_X)'\boldsymbol{\Sigma}_X^{-1}(\boldsymbol{x} - \boldsymbol{\mu}_X) + (\boldsymbol{y} - \boldsymbol{\mu}_Y)'\boldsymbol{\Sigma}_Y^{-1}(\boldsymbol{y} - \boldsymbol{\mu}_Y) \\ &+ \ln(\det(\boldsymbol{\Sigma}_X)) + \ln(\det(\boldsymbol{\Sigma}_Y))] - d\ln(2\pi). \end{aligned}$$

Because the last term is a constant, it can be omitted during optimization. The errors ϵ^{so} are assumed to be independent from ϵ^{y}. This means that the following log-likelihood needs to be maximized under $\mathbb{P}$:

$$\ln L^{\mathbb{P}} = -\frac{1}{2}\sum_{t=1}^{n}\left(\epsilon^{y'}_{t}\left(\Sigma^{y}\right)^{-1}\epsilon^{y}_{t} + \epsilon^{\text{so}'}_{t}\left(\Sigma^{\text{so}}_{t}\right)^{-1}\epsilon^{\text{so}}_{t} + \ln\left(\det\left(\Sigma^{y}\right)\right) + \ln\left(\det\left(\Sigma^{\text{so}}_{t}\right)\right)\right), \tag{3.35}$$

over the parameter space $\Theta = \{\mathbb{E}\mathbf{X}^{s}_{\infty}, \mathbb{E}^{\mathbb{Q}}\mathbf{X}^{s}_{\infty}, K, M, \Gamma_1, \Sigma^{r\pi}, \Sigma^{S\Pi}, \eta_S, \eta_\Pi\}$. The market prices of risk λ_0 and Λ_1 are determined using the constraints in Equation 3.11, 3.12, 3.14 and 3.15.

It is assumed that nominal and real zero-coupon bond observations with maturities $\{\tau_1, ..., \tau_5\} = \{1, 5, 10, 20, 30\}$ and $\{\tau^R_1, ..., \tau^R_5\} = \{1, 5, 15, 20, 30\}$ respectively have measurement errors and the bonds corresponding with maturities $\tau^*_N = 15$ and $\tau^*_R = 10$ have no measurement errors.

3.5.1 Optimization Algorithm

To maximize the log-likelihood, an optimization algorithm is required. A widely used class is gradient descent algorithms. Consider a differentiable one-dimensional function $L(x)$. For any differentiable function, taking a step of size ϵ in the direction of the slope will increase the function value for sufficiently small ϵ and smooth L. That is, $L(x + \epsilon\ \text{sign}(L'(x)) \geq L(x)$. This idea can be extended to multiple dimensions, where all model parameters can be updated as:

$$\boldsymbol{\Theta}_{n+1} = \boldsymbol{\Theta}_n + \epsilon\nabla_{\boldsymbol{\Theta}}L(\Theta),$$

where $\nabla_{\boldsymbol{\Theta}}L(\Theta)$ is the gradient with respect to all model parameters. Using a fixed learning rate throughout the entire optimization process is undesirable. A large value for ϵ could cause divergence, while a very small ϵ leads to slow convergence. Therefore, advanced algorithms typically use adaptive step sizes. Additionally, gradient descent is to the initialization of model parameters and can result in convergence to an undesirable local maximum.

To improve the accuracy of gradient descent, one can incorporate information about the curvature of L, known as the Hessian. By performing a second-order Taylor expansion of L around $\boldsymbol{\Theta}_n$ yields

$$L(\boldsymbol{\Theta}_{n+1}) = L(\boldsymbol{\Theta}_n) + (\boldsymbol{\Theta}_{n+1} - \boldsymbol{\Theta}_n)'\nabla_{\boldsymbol{\Theta}} L(\Theta_n) + \frac{1}{2}(\boldsymbol{\Theta}_{n+1} - \boldsymbol{\Theta}_n)' H(L(\Theta_n))(\boldsymbol{\Theta}_{n+1} - \boldsymbol{\Theta}_n) + R_2(\boldsymbol{\Theta}_{n+1} - \boldsymbol{\Theta}_n), \quad (3.36)$$

where R_2 is the remainder term. According to Newton's method, $L(\Theta_{n+1})$ in Equation 3.36 can be maximized by taking:

$$\Theta_{n+1} = \Theta_n + H^{-1}(L(\Theta_n))(\nabla_\Theta L(\Theta_n)). \quad (3.37)$$

However, calculating the exact Hessian matrix and its inverse is often too computationally expensive. Therefore, modern optimization algorithms, such as the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm, employs a Quasi-Newton approximation. An approximation of the Hessian is computed at each iteration using a secant method (Dennis and Schnabel [1983]) and an update according to Equation 3.37 is made for some step size ϵ. In BFGS, the optimal step size is determined via a line search that satisfies the Wolfe conditions, which ensures that the gradient converges to zero (Wolfe [1969]). The default optimization algorithm for calibrating CP2022 is the BFGS algorithm.

3.5.2 Constraints

Since optimization is performed over a large parameter space, there are many local maxima. To avoid converging to a suboptimal local maximum, five well-justified macroeconomic constraints are imposed at the time of calibration. The Ultimate Forward Rate (UFR) is defined as the limit for $\tau \to \infty$ of the zero-coupon rates. The constraints are:

$$\mathbb{E}^{\mathbb{P}} y_{60}(10) = -\frac{1}{10}\left(\phi(60,70) + \Psi(60,70)\mathbb{E}^{\mathbb{P}}\mathbf{X}^s_\infty\right) = \ln(1+0.020), \quad (3.38)$$

$$\mathbb{E}^{\mathbb{P}} y^R_{60}(10) = -\frac{1}{10}\left(\phi_R(60,70) + \Psi_R(60,70)\mathbb{E}^{\mathbb{P}}\mathbf{X}^s_\infty\right) = \ln(1+0.000), \quad (3.39)$$

$$\text{UFR} = \lim_{\tau\to\infty} y_t(\tau) = \lim_{\tau\to\infty} -\tau^{-1}(\phi(t,t+\tau) + \Psi(t,t+\tau)\mathbf{X}^s_t) = f_{t_0}(30,50), \quad (3.40)$$

$$\lim_{t\to\infty} \mathbb{E}^{\mathbb{P}}\left[\ln S_{t+1} - \ln S_t\right] = \mathbb{E}^{\mathbb{P}} r_\infty + \eta_S - \frac{1}{2}\sigma_S(\Gamma_0 + \mathbb{E}^{\mathbb{P}} v_\infty \Gamma)\sigma_S = \ln(1+0.054), \quad (3.41)$$

$$\lim_{t\to\infty} \mathbb{E}^{\mathbb{P}}\left[\ln \Pi_{t+1} - \ln \Pi_t\right] = \mathbb{E}^{\mathbb{P}} \pi_\infty + \eta_\Pi - \frac{1}{2}\sigma_\Pi(\Gamma_0 + \mathbb{E}^{\mathbb{P}} v_\infty \Gamma)\sigma_\Pi = \ln(1+0.020). \quad (3.42)$$

We discuss what the constraints imply. The authors of the KNW model imposed $\mathbb{P}(y_{60}(10) < 0) < 0.025$. In other words, the probability that the zero-coupon rate 60 years after calibration for a bond with maturity 10 years later, falls below zero is less than 2.5%. The restriction was imposed to limit the amount of economic scenarios with extremely high and low interest rates. Limiting these amounts is equivalent to fixing the expectation of the rates, assuming that the state $\mathbf{X}^s_t$ has converged to its long-term expectation under $\mathbb{P}$. The Commission Parameters decided, based on the long-term interest rate estimations done by Brand et al. [2018], to fix the expected nominal and real long-term interest rate at 2.0% and 0.0% respectively.

The third constraint fixes the Ultimate Forward Rate to the forward rate between maturities 30 and 50 year at the time of calibration, t_0. To understand the reasoning, we describe the UFR method first.

UFR

The basis of market-consistent valuation is to align the valuation of future cash flows with the market prices of risk-free bonds, which also generate future cash flows. For very long maturities,

few or no risk-free bonds are traded. Thus, to value cash flows far into the future, an extrapolation method is used. This extrapolation method is based on the level to which the interest rate converges, known as the Ultimate Forward Rate (UFR), and the path that approaches the UFR. The UFR method consists of 3 parts:

- A starting point for the extrapolation method, which is called the First Smoothing Point. This is the time from which no market data is used anymore to construct the term structure of interest rates.
- The extrapolation method used for the term structure from the starting point to converge to the UFR.
- The UFR level, to which the method converges.

The Commission Parameters decided:

- to choose the first smoothing point at 50 years,
- to use extrapolation to the term structure of interest rates beyond 50 years with a constant forward rate,
- to fix the UFR level to the 30 year forward rate with duration 20 years at the time of calibration.

During the 2020 Review of Solvency II, EIOPA has made an analysis of liquidity in the market for interest-rate swap contracts (EIOPA [2020]). The Commission Parameters used the findings of EIOPA, which concluded that the market is sufficiently liquid for interest-rate swap contracts with maturities up to 50 years, and set the first smoothing point at this maturity.

As seen earlier in Section 2.4, for long maturities, the forward rates and the spot rate converge to the same limit. Thus the UFR method should also converge to this same limit, and the forward rates with long maturities can be used as an approximation to this limit. Although it is more accurate to use forward rates with an expiry of 49 years and a maturity of 50 years, these rates are much less accurate to determine since the corresponding interest-rate products with these maturities are extremely illiquid. Since these products with maturities of 30 and 50 years are much more liquid, the Commission Parameters decided to use the latter approach instead. Within DNB's methodology of nominal forward rates:

$$f_{t_0}(T,S) = \ln\left(\left(\frac{P(t_0,t_0+T)}{P(t_0,t_0+S)}\right)^{\frac{1}{S-T}}\right),$$

the UFR is set to

$$UFR = f_{t_0}(30,50) = \ln\left(\left(\frac{P(t_0,t_0+30)}{P(t_0,t_0+50)}\right)^{\frac{1}{20}}\right) = \frac{50y_{t_0}(50) - 30y_{t_0}(30)}{20}.$$

We return to the third constraint and analyse its implications in CP2022. First, we state the following corollary.

Corollary 3.5.0.1. *Let f be a function which is continuous on $[a,b]$ and differentiable on (a,b). Then $\exists c \in (a,b)$ such that*

$$\int_a^b f(x)dx = f(c)(b-a).$$

Proof. By the Mean Value theorem, $\exists c \in (a,b)$ such that $F'(c) = f(c) = \frac{F(b)-F(a)}{b-a}$, where F is the antiderivative of f. Additionally, by the Fundamental Theorem of Calculus, $\int_a^b f(x)dx = F(b) - F(a)$. Combining results finishes the proof. □

Remark that in Equation 3.18, we may write $\Psi(t, t+\tau) = \Psi(\tau)$ and if $\zeta(t)$ is constant, then we also have $\phi(t, t+\tau) = \phi(\tau)$. This means that $\forall t \leq t_0$:

$$\begin{aligned}
UFR &= \lim_{\tau\to\infty} -\tau^{-1}(\phi(t,t+\tau) + \Psi(t,t+\tau)\mathbf{X}_t^s) \\
&= \lim_{\tau\to\infty} -\tau^{-1}(\phi(\tau) + \Psi(\tau)\mathbf{X}_t^s) \\
&\overset{\text{(i)}}{=} \lim_{\tau\to\infty} -\tau^{-1}\phi(\tau) \\
&\overset{3.20}{=} \lim_{\tau\to\infty} -\tau^{-1}\int_{t_0}^{t_0+\tau}\left(\Psi(t_0+\tau-s)'\zeta(s) + \frac{1}{2}\Psi(t_0+\tau-s)'\Sigma G_0\Sigma'\Psi(t_0+\tau-s)\right)ds \\
&\overset{3.5.0.1}{=} -\lim_{\tau\to\infty}\tau^{-1}\tau\left(\Psi(t_0+\tau-c)'\zeta(c) + \frac{1}{2}\Psi(t_0+\tau-c)'\Sigma G_0\Sigma'\Psi(t_0+\tau-c)\right) \\
&\overset{\text{(ii)}}{=} -\Psi_\infty'\left(M\mathbb{E}^{\mathbb{Q}}\mathbf{X}_\infty^s + \frac{1}{2}\Sigma^{r\pi}\Gamma_0(\Sigma^{r\pi})'\Psi_\infty\right),
\end{aligned}$$

where in step (i) $\lim_{\tau\to\infty}\Psi(\tau)\mathbf{X}_t^s/\tau$ vanishes to zero since $\mathbf{X}_t^s$ is stationary and in step (ii) we use parameters as in Equation 3.23, and $\Psi_\infty \in \mathbb{R}^3$ is determined from Equation 3.19:

$$-\frac{1}{2}\Psi_\infty'\Sigma^{r\pi}\Gamma(\Sigma^{r\pi})'\Psi_\infty\mathbb{1}_{\{i=1\}} + (M'\Psi_\infty)_i + \mathbb{1}_{\{i=2\}} = 0 \quad \forall i \in \{1,2,3\},$$

The final two constraints represent the expected long-term annual geometric returns for the stock index S_t and consumer price index Π_t. These converge, since $\mathbf{X}_t$ is stationary, given K has positive eigenvalues. Furthermore, the constraints follow directly from Equation 3.3. The Commission Parameters set the expected long-term European consumer price index to 2.0%, in line with DNB's objective. The Commission Parameters set the expected long-term annual return on the stock index to 5.4%, based on Dimson et al. [2022].

Eventually, the state variables v_t, r_t and π_t converge to an equilibrium distribution. This means that their values will continue to fluctuate in the long term. Although the means converge to a constant expected value, the variances of these processes converge to strictly positive constants. The rate at which the state variables converge to their equilibrium distribution depends on the parameter K resulting from calibration. The expected values in equilibrium themselves are difficult to estimate using the available market data, since financial instruments with very long maturities are illiquid. Therefore, they were imposed implicitly as in Equation 3.41 and 3.42.

3.5.3 Joint Optimization in $\mathbb{P}$ and $\mathbb{Q}$

To find a set of model parameters under $\mathbb{P}$ and $\mathbb{Q}$ to fit the market, an optimization routine has been set up that consists of two parts. In the first part, a likelihood criterium under $\mathbb{P}$ with constraints is set up (see Section 3.5) to fit historical nominal and real bond prices using the state variables v_t, r_t and π_t. The first variable is assumed to be observable from the VSTOXX index and the last two can be determined as in Equation 3.22 and 3.27 respectively from the nominal bond prices 15 years maturity and real bond prices with 10 year maturity respectively, which were the ones assumed to have no measurement errors.

In the second part, constraints under $\mathbb{Q}$ are set up to ensure that the mean squared errors between option prices observed in the market and option prices under the CP2022 model at the time of calibration are below a certain threshold:

$$e_S(\Theta)^2 \leq (1.50\%)^2, \quad e_r(\Theta)^2 \leq (0.15\%)^2, \quad e_\Pi(\Theta)^2 \leq (0.50\%)^2, \tag{3.43}$$

with

$$e_i(\Theta)^2 = \frac{1}{n_i}\sum_{k=1}^{n_i}\left(\frac{p_k^{i,\text{observed}} - p_k^{i,\text{model}}(\Theta)}{\mathcal{U}_k^i}\right)^2, \tag{3.44}$$

where i denotes the underlying of an option, which can be stock index S_t, interest rate r_t or inflation rate Π_t. Moreover, $p_k^{i,\text{observed}}$ is the observed option price in the market of the k-th

option with underlying i, $p_k^{i,\text{model}}(\Theta)$ is the corresponding option price under the CP2022 model for a fixed model parameter set Θ and υ_k^i is the vega of the k-th option, which measures the sensitivity of an option's price to changes in the volatility of the underlying. In here, vega are the weights added to the mean squared errors during optimization, which are determined from Black's model (Black [1976]), where the implied volatility inserted are quoted in the market and the discount factors can be determined by the term structure of interest rates (see Appendix A.2). We call this loss function "the implied volatility mean squared error", because option pricing errors in terms of prices become option pricing errors in terms of the implied volatilities via a linear approximation. This can be seen by performing a first order Taylor expansion of p^{model} with respect to its implied volatility:

$$p^{\text{model}} = p^{\text{obs}} + \left.\frac{\partial p^{\text{obs}}}{\partial \sigma}\right|_{\sigma=\sigma^{\text{obs}}} (\sigma^{\text{model}} - \sigma^{\text{obs}}) + \mathcal{O}(\sigma^{\text{model}} - \sigma^{\text{obs}})^2.$$

Since the volatility of option prices close to maturity and around at-the-money strike prices have a high volatility, minimizing the implied volatility mean squared error will prevent creating long-term scenarios of 100 year in the future, which are highly biased to options close to maturity. Option prices under the CP2022 model are determined under $\mathbb{Q}$ with time-varying market prices of risk using a Monte Carlo-based approach, which will be discussed in Section 3.8.

The joint optimization in $\mathbb{P}$ and $\mathbb{Q}$ is a constrained optimization problem, which is transformed to an unconstrained optimization problem via a penalty method:

$$\max_{\Theta} \ln L^{\mathbb{P}}(\Theta) - p\left(\max(0, e_S(\Theta) - 0.015)^2 + \max(0, e_r(\Theta) - 0.0015)^2 + \max(0, e_{\Pi}(\Theta) - 0.005)^2\right), \tag{3.45}$$

where p is a high penalty coefficient such as 10^{12}.

Within the context of pensions, assets and liabilities are valued via the scenario sets under $\mathbb{Q}$ generated by CP2022. A participant's future pension cash flows can be seen as a complex option product, where the value depends on factors such as the development of wages, inflation, interest rates, the pension fund's returns and more. For example, the present value is determined using expected pension payments, and uncertain events in the future, such as not indexing the pension payments with inflation or the possibility of unexpected setbacks in the investments made by pension funds, will negatively impact the pension payments. A recent example is the regional bank runs in the USA in 2023, which also affected Dutch pension funds. Young people are much more sensitive to these future events than people close to their pension age. Hence the Commission Parameters has chosen to create $\mathbb{Q}$-sets that are consistent with market instruments with underlying a stock index, interest rates and inflation rates. In particular, calibrating using options with different strike prices and maturities will ensure that the generated economic scenarios are in line with current market conditions and will replicate the fat tails in the distributions. The options used during calibration are European calls on the stock index, zero-coupon inflation-indexed caps and floors, year-on-year inflation-indexed caps and floors and payer swaptions.

The data used for calibration includes:

1. The nominal zero-coupon rates for yearly maturities of 1-50 years:
   - 07/2004 - 11/2008: Zero-coupon rates from the public statistics website of DNB.
   - 12/2008 - present: bootstrapped swap data from Bloomberg to derive zero-coupon rates (EUSAx Curncy where x is the maturity).
2. The real zero-coupon rates were constructed based on the nominal zero-coupon rates and the Eurozone Harmonised Index of Consumer Prices (HICP) zero-coupon break-even inflation (EUSWIx Curncy where x is the maturity).
3. The stock index used is the MSCI World index (MSDEWIN)
4. Volatility index EURO STOXX 50 (V2TX)

5. The seasonally adjusted HICP index available publicly via DNB.
6. EURO STOXX calls with maturities 1,2 and 5 years and moneyness 80%, 90%, 100%, 105%, 110%, 120% (SX5E)
7. Eurozone payer swaptions with maturities 1-10, 12, 15, 20, 25, 30, 40 and 50 years (EUSAx Curncy with x the maturity)
8. Eurozone year-on-year inflation-indexed caps and floors with maturities 1,2,5 and 10 years. The caps included have strikes 2%, 3%, 4% (EUISCyx with strike y and maturity x) and the floors have strike -1% and 1% (EUISFyx with strike y and maturity x)
9. Eurozone zero-coupon inflation-indexed caps and floors with maturities 5,10 and 20 years. Caps included have strike 2%, 3%, 4% (EUIZCyx with strike y and x maturity) and floors with strike 1% (EUIZFyx with strike y and maturity x)
10. EURIBOR swaption implied volatilities
11. EURO STOXX option implied volatilities

Note that the Commission Parameters has used the euro-denominated MSCI World Index for equity returns, because this index extends beyond European markets, like investments made by Dutch pension funds. Moreover, note that no options exist on MSCI World Index, hence EURO STOXX options are used instead. The complete algorithm for calibrating under $\mathbb{P}$ and $\mathbb{Q}$ simultaneously is summarized below.

Algorithm 3.1: Joint calibration in $\mathbb{P}$ and $\mathbb{Q}$

Data: nominal and real zero-coupon bonds and options quoted on the market;
Initialization: Θ;
while *True* **do**
1. Reparametrize K and M as in Appendix A.1;
2. $\gamma_{ii} \to \exp(\gamma_{ii})$;
3. Update $\mathbb{E}v_\infty$ and $\mathbb{E}^{\mathbb{Q}}v_\infty$ such that the Feller condition holds in $\mathbb{P}$ and $\mathbb{Q}$;
4. Determine the market prices of risk λ_0 (not time-dependent yet) and Λ_1 using Equation 3.11, 3.12, 3.14 and 3.15;
5. Update $\mathbb{E}^{\mathbb{Q}}r_\infty$ as in $\mathbb{P}$-constraint in Equation 3.40;
6. Determine relevant $\phi(t,T)$ and $\Psi(t,T)$ for bonds as in Equation 3.18;
7. Calculate bond prices under the affine model as in Equation 3.22 and 3.27;
8. Determine $\mathbf{X}_{t_0}$ using Equation 3.32;
9. Determine f, f^R with subsequent $\tilde{\phi}$ and $\tilde{\phi}^R$ as in Section 3.4;
10. Generate paths in $\mathbb{Q}$ with time-dependent market prices of risk as in Section 3.7;
11. Use Monte Carlo to determine the option prices under the model as in Section 3.8;
12. Update $\mathbb{E}^{\mathbb{P}}r_\infty$ and $\mathbb{E}^{\mathbb{P}}\pi_\infty$ as in $\mathbb{P}$-constraints in Equation 3.38 and 3.39 respectively with time-dependent market prices of risk;
13. Update η_S and η_Π as in $\mathbb{P}$-constraints in Equation 3.41 and Equation 3.42 respectively under time-dependent market prices of risk;
14. Calculate the log-likelihood as in Equation 3.35;
15. Calculate the implied volatility mean squared errors as in Equation 3.44;

if *change in log-likelihood* $\leq$ *TOL as in Equation 3.45*, **then**
Exit;
else
Update Θ using BFGS and go to step 1;

Return: Θ

Thus far, we have only stated that the simulation of the consumer price index Π_t is based on the Eurozone Harmonized Index of Consumer Prices. However, Dutch pension funds primarily rely on the Dutch Consumer Price Index Π_t^{NL}, for which historical observations are recorded by the Statistics Netherlands (CBS). To simulate future scenarios of a Dutch inflation rate process, it is

assumed that the Dutch inflation rate $\Pi^{NL}_{t+\Delta t}/\Pi^{NL}_t$ equals $\Pi_{t+\Delta t}/\Pi_t$ plus a time-varying spread such that

$$\mathbb{E}^{\mathbb{P}}\left[\ln(\Pi^{NL}_{t+\Delta t}/\Pi^{NL}_t)\right] = \ln(1 + I^{NL}_{\lfloor t \rfloor+1})\Delta t, \tag{3.46}$$

where $I^{NL}_{\lfloor t \rfloor+1}$ is the forecast of the Dutch inflation rate as published by The Netherlands Bureau for Economic Policy Analysis (CPB [2024]).

3.6 HESTON'S STOCHASTIC VARIANCE PROCESS SIMULATIONS

For simulating paths in $\mathbb{P}$ and $\mathbb{Q}$, simulating the volatility in Heston's model is required. There exist analytical results for the Heston process, see Cox et al. [1985]. For example, the distribution of v_t conditional on v_s for $\Delta t = t - s > 0$ is known:

$$\mathbb{P}(v_t \leq x \mid v_s) = F_{\chi^2_d(\lambda)}\left(\frac{x}{\xi}\right),$$

where

$$\xi = \frac{w^2\left(1 - e^{-K_{vv}\Delta t}\right)}{4K_{vv}}, \quad \lambda = \frac{4v_s K_{vv} e^{-K_{vv}\Delta t}}{w^2(1 - e^{K_{vv}\Delta t})}, \quad d = \frac{4K_{vv}\mathbb{E}v_\infty}{w^2},$$

and $F_{\chi^2_d(\lambda)}(x)$ is the cumulative distribution function of the non-central chi-squared distribution:

$$F_{\chi^2_d(\lambda)}(x) = \sum_{i=0}^{\infty} \frac{e^{-\frac{\lambda}{2}}\left(\frac{\lambda}{2}\right)^i}{i!} \frac{\int_0^z x^{\frac{d}{2}} e^{-\frac{u}{2}} du}{\Gamma\left(i + \frac{k}{2}\right)},$$

with Gamma function Γ and probability density function:

$$f_{\chi^2_d(\lambda)}(x) = \frac{1}{2} e^{-\frac{x+\lambda}{2}} \left(\frac{x}{\lambda}\right)^{\frac{d}{4}-\frac{1}{2}} I_{\frac{d}{2}-1}\left(\sqrt{\lambda x}\right),$$

where I_d is the modified Bessel function of the first kind of order d. Also the mean and variance are known of the conditional distribution:

$$\begin{aligned}
\mathbb{E}[v_t|v_s] &= \mathbb{E}v_\infty + (v_s - \mathbb{E}v_\infty)e^{-K_{vv}\Delta t}, \\
\mathrm{Var}[v_t|v_s] &= \frac{v_s w^2 e^{-K_{vv}\Delta t}}{K_{vv}}\left(1 - e^{-K_{vv}\Delta t}\right) + \frac{\mathbb{E}v_\infty w^2}{2K_{vv}}\left(1 - e^{-K_{vv}\Delta t}\right)^2.
\end{aligned}$$

Although one may be inclined to simulate v_t using an Euler discretization scheme using standard normally distributed variables Z, the probability that v_t becomes negative is

$$\mathbb{P}(v_t < 0) = \mathbb{P}\left(Z < \frac{-v_s - K_{vv}\Delta t(\mathbb{E}v_\infty - v_s)}{w\sqrt{v_s}\sqrt{\Delta t}}\right) = \Phi\left(\frac{-v_s - K_{vv}\Delta t(\mathbb{E}v_\infty - v_s)}{w\sqrt{v_s \Delta t}}\right),$$

where $\Phi(\cdot)$ is the cumulative distribution function of the Normal distribution. Note that this probability is always non-zero for any time step Δt. If the variance process crosses the zero axis, it is possible to make the variance process positive again by absorbing or reflecting it with respect to the axis, however, this will create biased paths (Glasserman [2004]). Lord et al. [2010] show that by applying an ad hoc correction to the standard Euler scheme, their Full Truncation Euler scheme works empirically well:

$$v_t = v_s + K_{vv}(\mathbb{E}v_\infty - v_s^+) + wZ\sqrt{v_s^+ dt},$$

where $v_s^+ = \max(v_s, 0)$.

3.6.1 Broadie and Kaya Scheme

The first bias-free and exact discretization method for simulating from Heston's model was developed by Broadie and Kaya [2006]. The simulation is exact, in the sense that the process can be simulated for any future time, without discretizing the time-grid first. This is done as follows. Integrating the variance process gives:

$$v_t = v_s + \int_s^t K_{vv}(\mathbb{E}v_\infty - v_u)du + w\int_s^t \sqrt{v_u}dW_u.$$

Hence, Broadie and Kaya [2006] simulate from three stochastic variables:

1. $v_t \mid v_s$: sampling from $F_{\chi^2_d(\lambda)}\left(\frac{x}{\xi}\right)$ comes down to sampling from the Gamma distribution, which can be done via acceptance rejection methods, see Cheng and Feast [1979].

2. $\int_s^t v_u du \mid v_s, v_t$: Broadie and Kaya [2006] derive the characteristic function:

$$\begin{aligned}\phi(z, v_s, v_t) =& \mathbb{E}\left[\exp\left(iz\int_s^t v_u du\right) \mid v_s, v_t\right] \\ =& \frac{\gamma(z)e^{\frac{1}{2}(\gamma(z)-K_{vv})(t-s))}(1-\exp(-K_{vv}(t-s))}{K_{vv}\left(1-e^{-\gamma(z)(t-s)}\right)} \\ &\times \exp\left[\frac{v_s+v_t}{w^2}\left(\frac{K_{vv}\left(1+e^{-K_{vv}(t-s)}\right.}{1-e^{-K_{vv}(t-s)}} - \frac{\gamma(z)\left(1+e^{-\gamma(z)(t-s)}\right)}{1-e^{-\gamma(z)(t-s)}}\right)\right] \\ &\times \frac{I_{\frac{1}{2}d-1}\left[\sqrt{v_s v_t}4\gamma(z)e^{-\frac{\gamma(z)}{2}(t-s)}/w^2\left(1-e^{-\gamma(z)(t-s)}\right)\right]}{I_{\frac{1}{2}d-1}\left[\sqrt{v_s v_t}4\gamma(z)e^{-\frac{K_{vv}}{2}(t-s)}/w^2\left(1-e^{-K_{vv}(t-s)}\right)\right]},\end{aligned}$$

with $\gamma(z) = \sqrt{K_{vv}^2 - 2w^2 iz}$. Inverting the characteristic function to get the cumulative distribution function F like in Feller [1957] gives:

$$G(x, v_s, v_t) = \frac{2}{\pi}\int_0^\infty \frac{\sin(ux)}{x}\,\mathrm{Re}[\phi(u, v_s, v_t)]du,$$

which can be computed numerically. Simulating the required variable can be done via the inverse transform sampling method:

$$G\left(x_i, v_s, v_t\right) - U = 0,$$

where U is a standard uniform variable. The above can be solved using a root-searching algorithm like Newton-Raphson. Note that for each iteration of the root-searching algorithm, $G(x_i, v_s, v_t)$ has to be recomputed numerically.

3. $\int_s^t \sqrt{v_u}dW_u \mid \int_s^t v_u du$: since v_u and W_u are independent, this variable is distributed as $N\left(0, \int_s^t v_u\right)$.

The biggest bottleneck is step 2. If only path-independent options are used, with this scheme one can jump immediately to the maturity date T but with a high computational cost. Since in CP2022 path-dependent options are also included and option prices are evaluated using Monte Carlo, this scheme is exceptionally time-consuming. Another bottleneck is that sampling from the Gamma distribution in step 1 via direct inversion is relatively slow.

In an attempt to fix the bottlenecks from above, Andersen [2007] came up with an efficient scheme: Andersen's Quadratic-Exponential (QE) scheme, which is implemented in CP2022.

3.6.2 Andersen's Quadratic-Exponential Scheme

Instead of sampling v_t from the non-central chi-squared distribution, instead Andersen [2007] simulate from a related distribution with distribution parameters determined using moment-matching with the non-central chi-squared distribution. Let $\psi = \mathbb{E}[v_t|v_s]^2/\text{Var}[v_t|v_s]$. They argue that

1. When v_s is far from the origin, v_t is approximately a power function of a standard Gaussian variable Z:

$$v_t = a(b + Z)^2,$$

with parameters a and b such that the first two moments are matched with the exact moments of the non-central chi-squared distribution for $\psi \leq 2$:

$$\begin{aligned} b^2 &= 2\psi^{-1} - 1 + \sqrt{2\psi^{-1}}\sqrt{2\psi^{-1} - 1}, \\ a &= \frac{m}{1 + b^2}. \end{aligned}$$

2. When v_s is close to the origin, the probability density function of v_t is asymptotically:

$$\mathbb{P}(v_t \in [x, x + dx]) \approx \left(p\delta(0) + \beta(1 - p)e^{-\beta x}\right) dx, \quad x \geq 0,$$

with δ the Dirac delta function and parameters p and β. Then, using inverse transform sampling with standard uniform variate U,

$$v_t = \begin{cases} 0 & \text{if } 0 \leq u \leq p, \\ \beta^{-1} \log\left(\frac{1-p}{1-u}\right) & \text{if } p < u \leq 1. \end{cases}$$

The parameters p and β are chosen such that the first two moments of the distribution function above equals the exact moments of the non-central chi-squared distribution for $\psi \geq 1$:

$$\begin{aligned} p &= \frac{\psi - 1}{\psi + 1}, \\ \beta &= \frac{1 - p}{m} = \frac{2}{m(\psi + 1)}. \end{aligned}$$

Note that the parameters match the moments exactly only when $\psi \leq 2$ and $\psi \geq 1$ respectively. In $\psi \in [1, 2]$ where both methods overlap, only one method is chosen and is determined via critical level $\psi_c \in [1, 2]$. That is, if $\psi \leq \psi_c$, then method 1 is performed and method 2 is performed for $\psi > \psi_c$.

Since ψ depends on v_s and the Heston parameters, this scheme uses time-discretization and can be used for pricing path-dependent options. Although ψ has to be computed at each step, the computational cost is very cheap as many terms can be reused:

$$\psi = \frac{\gamma_1 v_s + \gamma_2}{(\gamma_3 v_s + \gamma_4)^2},$$

with

$$\gamma_1 = \frac{w^2 e^{-K_{vv}\Delta t}}{K_{vv}}\left(1 - e^{-K_{vv}\Delta t}\right), \gamma_2 = \frac{w^2 \mathbb{E}v_\infty}{2K_{vv}}\left(1 - e^{-K_{vv}\Delta t}\right)^2, \gamma_3 = e^{-K_{vv}\Delta t}, \gamma_4 = \mathbb{E}v_\infty\left(1 - e^{-K_{vv}\Delta t}\right),$$

which makes this scheme efficient.

3.7 MONTE CARLO SIMULATIONS UNDER $\mathbb{P}$ AND $\mathbb{Q}$

At the time of calibration t_0, N paths are simulated to create economic scenarios

$$\left\{\left(\mathbf{X}^{s,j}_{t_0+i\Delta t}, \mathbf{X}^{o,j}_{t_0+i\Delta t}, R^j_{t_0+i\Delta t}\right) = \left(v^j_{t_0+i\Delta t}, r^j_{t_0+i\Delta t}, \pi^j_{t_0+i\Delta t}, \ln S^j_{t_0+i\Delta t}, \ln \Pi^j_{t_0+i\Delta t}, R^j_{t_0+i\Delta t}\right)\right\}^{j=1\ldots N}_{i=0\ldots n},$$

using n time steps and $R^j_{t+\Delta t} = r^j_t \Delta t + R^j_t$ is the continuously-compounded spot interest rate used to approximate the discount factors when determining option prices using Monte Carlo. First, Heston's variance parameter is simulated using Andersen's QE scheme with iid uniformly distributed samples U^j_t. Simulating $\mathbb{P}$-scenarios is done based on the $\mathbb{P}$-dynamics in Equation 3.1 and using the Euler–Maruyama scheme:

$$v^j_{t+\Delta t} = f_{\text{Andersen}}\left(v^j_t, U^j_t\right), \quad \eta^j_t = \omega^{-1}\left(v^j_t \Delta t\right)^{-\frac{1}{2}}\left(v^j_{t+\Delta t} - v^j_t - K_{vv}\left(\mathbb{E}^{\mathbb{P}} v_\infty - v^j_t\right)\Delta t\right),$$

and

$$\begin{bmatrix} r^j_{t+\Delta t} - r^j_t \\ \pi^j_{t+\Delta t} - \pi^j_t \\ \ln\left(S^j_{t+\Delta t}/S^j_t\right) \\ \ln\left(\Pi^j_{t+\Delta t}/\Pi^j_t\right) \end{bmatrix} = \begin{bmatrix} 0_{4\times 1} & I_4 \end{bmatrix}\left(\begin{bmatrix} K\left(\mathbb{E}\mathbf{X}^s_\infty - \mathbf{X}^{s,j}_t\right) \\ \mu^o + K^o \mathbf{X}^{s,j}_t \end{bmatrix}\Delta t + \left(\Sigma^{\text{so},j}_t\right)^{\frac{1}{2}}\begin{bmatrix} \eta^j_t \\ \xi^j_t \end{bmatrix}\right),$$

with

$$\xi^j_t \overset{iid}{\sim} N\left(0_{4\times 1}, I_4\right), \quad \left(\Sigma^{\text{so},j}_t\right)^{\frac{1}{2}} = \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix}\left(\Gamma_0 + v^j_t \Gamma\right)^{\frac{1}{2}}\sqrt{\Delta t}.$$

The Dutch consumer price index is simulated as in Equation 3.46:

$$\begin{aligned} \ln\left(\Pi^{\text{NL},j}_{t+\Delta t}/\Pi^{\text{NL},j}_t\right) &= \ln\left(\Pi^j_{t+\Delta t}/\Pi^j_t\right) + H^j_t \\ H^j_t &= -\frac{1}{N}\sum_{j=1}^{N}\ln\left(\Pi^j_{t+\Delta t}/\Pi^j_t\right) + \ln\left(1 + I^{\text{NL}}_{\lfloor t\rfloor+1}\right)\Delta t \end{aligned}$$

Simulation in $\mathbb{Q}$ is analogous to the $\mathbb{P}$ case with the $\mathbb{Q}$-parameters and $\mathbb{Q}$-dynamics in Equation 3.29:

$$v^j_{t+\Delta t} = f_{\text{Andersen}}\left(v^j_t, U^j_t\right), \quad \eta^j_t = \omega^{-1}\left(v^j_t \Delta t\right)^{-\frac{1}{2}}\left(v^j_{t+\Delta t} - v^j_t - M_{vv}\left(\mathbb{E}^{\mathbb{Q}} v_\infty - v^j_t\right)\Delta t\right)$$

with

$$\begin{bmatrix} r^j_{t+\Delta t} - r^j_t \\ \pi^j_{t+\Delta t} - \pi^j_t \\ \ln\left(S^j_{t+\Delta t}/S^j_t\right) \\ \ln\left(\Pi^j_{t+\Delta t}/\Pi^j_t\right) \end{bmatrix} = \begin{bmatrix} 0_{4\times 1} & I_4 \end{bmatrix}\left(\begin{bmatrix} M\left(\mathbb{E}^{\mathbb{Q}}\mathbf{X}^s_\infty - \mathbf{X}^{s,j}_t\right) - \begin{bmatrix} 0 \\ f(t) \\ 0 \end{bmatrix} \\ \mu^o - \begin{bmatrix} \eta_S + 0 \\ \eta_\Pi + f^R(t) \end{bmatrix} + K^o \mathbf{X}^{s,j}_t \end{bmatrix}\Delta t + \left(\Sigma^{\text{so},j}_t\right)^{\frac{1}{2}}\begin{bmatrix} \eta^j_t \\ \xi^j_t \end{bmatrix}\right)$$

and

$$\xi^j_t \overset{iid}{\sim} N\left(0_{4\times 1}, I_4\right), \quad \left(\Sigma^{\text{so},j}_t\right)^{\frac{1}{2}} = \begin{bmatrix} \Sigma^{r\pi} \\ \Sigma^{S\Pi} \end{bmatrix}\left(\Gamma_0 + v^j_t \Gamma\right)^{\frac{1}{2}}\sqrt{\Delta t}.$$

3.8 MONTE CARLO METHOD FOR PRICING OPTIONS IN CP2022

By simulating paths as in Section 3.7, the prices of European call options on the stock index, payer swaptions, zero-coupon inflation-indexed caps and floors and year-on-year inflation-indexed caps and floors that are used for calibration, can be determined using Monte Carlo. The expected nominal and real discount rates at later times $T \geq t_0$ are the prices of nominal and real zero-coupon bonds:

$$P(T, T+\tau) = \mathbb{E}_T^{\mathbb{Q}}\left[e^{-\int_T^{T+\tau} r_u du}\right] = e^{\phi(T,T+\tau)+\Psi(\tau)'\mathbf{X}_T^s},$$
$$P^R(T, T+\tau) = \mathbb{E}_T^{\mathbb{Q}}\left[e^{-\int_T^{T+\tau} r_u du}\frac{\Pi_{T+\tau}}{\Pi_T}\right] = e^{\phi_R(T,T+\tau)+\Psi_R(\tau)'\mathbf{X}_T^s},$$

with simulated equivalents

$$\bar{P}\left(T, T+\tau, \mathbf{X}_T^{s,j}\right) = e^{\phi(T,T+\tau)+\Psi(\tau)'\mathbf{X}_T^{s,j}},$$
$$\bar{P}^R\left(T, T+\tau, \mathbf{X}_T^{s,j}\right) = e^{\phi_R(T,T+\tau)+\Psi_R(\tau)'\mathbf{X}_T^{s,j}}.$$

Option prices are determined using the Monte Carlo method, which involves simulating N paths of the underlying asset, as outlined in Section 3.7. The conditional expectations of the discounted payoffs are then approximated by the sample mean, a result that follows from the law of large numbers:

$$\begin{aligned}
\mathbf{C}(t_0, T, K) &= \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^T r_u du}(S_T - K)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^N\left(e^{-R_T^j}\left(e^{\ln S_T^j} - K\right)^+\right), \\
\mathbf{PSW}(t_0, T_a, T_b, K) &= \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^{T_a} r_u du}\left(1 - P(T_a, T_b) - K\sum_{T_k=T_a+1}^{T_b} P(T_a, T_k)\right)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^N e^{-R_{T_a}^j}\left(1 - \bar{P}\left(T_a, T_b, \mathbf{X}_{T_a}^{s,j}\right) - K\sum_{T_k=T_a+1}^{T_b}\bar{P}\left(T_a, T_k, \mathbf{X}_{T_a}^{s,j}\right)\right)^+, \\
\mathbf{YIC}(t_0, T, K) &= \sum_{T_k=t_0+1}^T \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^{T_k} r_u du}\left(\frac{\Pi_{T_k} - \Pi_{T_k-1}}{\Pi_{T_k-1}} - K\right)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^N\sum_{T_k=t_0+1}^T e^{-R_{T_k}^j}\left(e^{\ln\Pi_{T_k}^j - \ln\Pi_{T_k-1}^j} - 1 - K\right)^+, \\
\mathbf{YIF}(t_0, T, K) &= \sum_{T_k=t_0+1}^T \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^{T_k} r_u du}\left(K - \frac{\Pi_{T_k} - \Pi_{T_k-1}}{\Pi_{T_k-1}}\right)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^N\sum_{T_k=t_0+1}^T e^{-R_{T_k}}\left(K + 1 - e^{\ln\Pi_{T_k}^j - \ln\Pi_{T_k}^j - 1}\right)^+, \\
\mathbf{IC}(t_0, T, K) &= \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^T r_u du}\left(\frac{\Pi_T}{\Pi_{t_0}} - (K+1)^T\right)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^T\sum_{T_k=t_0+1}^T e^{-R_{T_k}}\left(e^{\ln\Pi_{T_k}^j - \ln\Pi_{T_k}^j - 1} - (K+1)^T\right)^+, \\
\mathbf{IF}(t_0, T, K) &= \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^T r_u du}\left((K+1)^T - \frac{\Pi_T}{\Pi_{t_0}}\right)^+\right] \\
&\approx \frac{1}{N}\sum_{j=1}^T\sum_{T_k=t_0+1}^T e^{-R_{T_k}}\left((K+1)^T - e^{\ln\Pi_{T_k}^j - \ln\Pi_{T_k}^j - 1}\right)^+.
\end{aligned}$$

3.9 VALIDATION OF MODEL ASSUMPTIONS AND SETTINGS

All results were obtained using Intel(R) Xeon(R) Platinum 8370C CPU @ 2.8GHz with 8 Gb of DDR4-3200 MHz memory. The code is written in Matlab R2024a. Additionally, we use the CP2022 parameters of Q2 2024 that are made publicly available by DNB each quarter (De Nederlandsche Bank [2022]) and can also be found in Appendix A.3.

To validate the model, we perform several tests. A key feature of CP2022 is its affine property, which enables the direct derivation of the nominal and real interest-rate term structures. This property arises from the results in Theorem 3.3.1, under certain integrability assumptions. To verify these assumptions, we compare the nominal and real zero-coupon bond prices using Theorem 3.3.1 with those obtained using the Monte Carlo method. For the Monte Carlo simulations, we project $T = 100$ years into the future, using $N = 100,000$ simulations with a time step of $h = 1/12$ per year and use a 95% confidence interval.

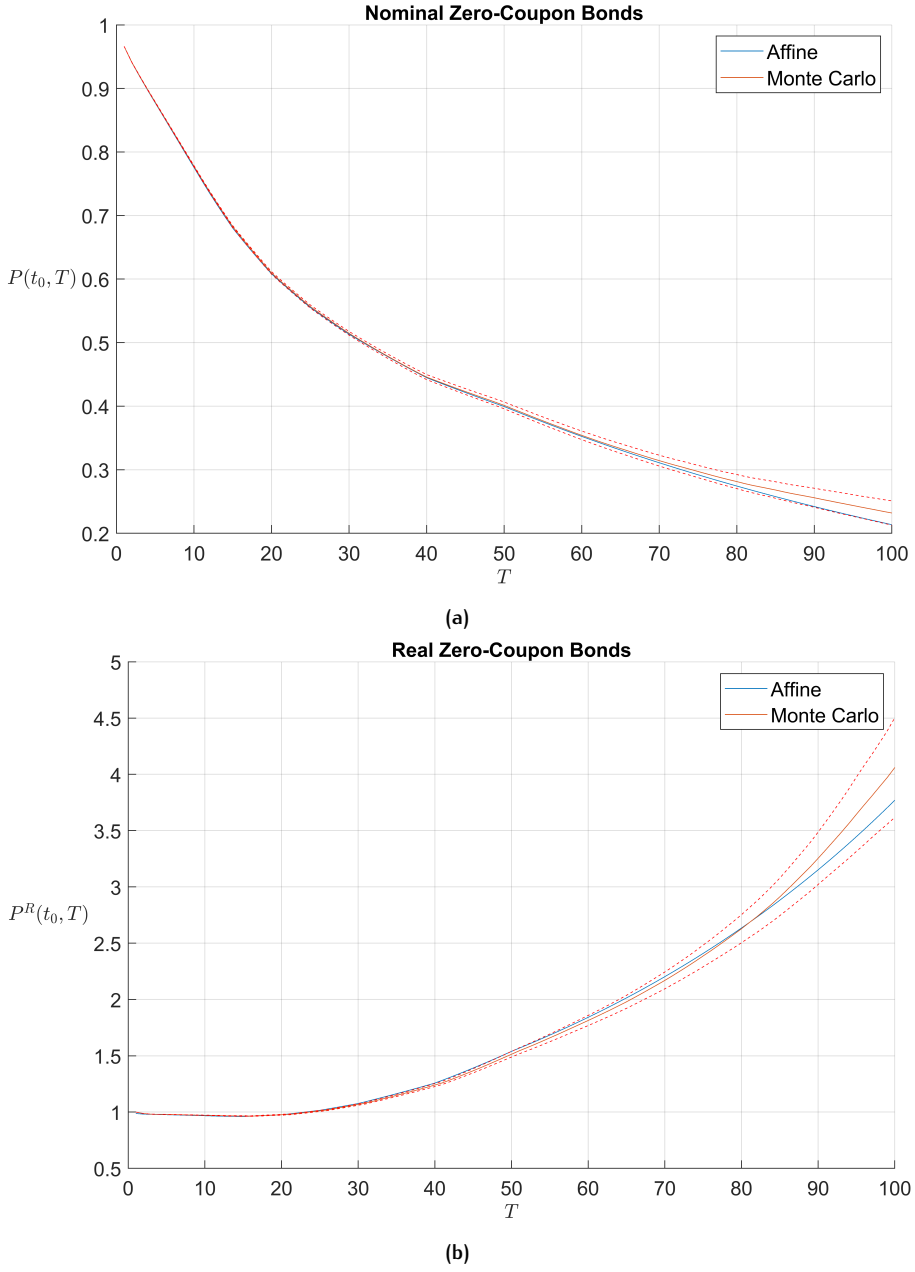

Figure 3.1: Nominal zero-coupon bond prices $P(t_0, T)$ (a) and real zero-coupon bond prices $P^R(t_0, T)$ (b) as a function of maturity T using Theorem 3.3.1 and Monte Carlo with a 95% confidence interval.

The results derived from the affine property show close agreement with the Monte Carlo simulations. Note that, in addition to simulation error, a discretization error is introduced when generating simulation paths using the Euler–Maruyama scheme, resulting in an additional error in Monte Carlo results. In the next example, we illustrate that increasing the number of time steps per year improves the accuracy of the Euler–Maruyama scheme, leading to convergence towards the true values.

To ensure that the model's term structure of interest rates aligns closely with the term structure observed in the market, adjustment functions $f(t)$ and $f^R(t)$ are defined. These functions minimize the discrepancies between the model and market data while also defining the transformation from $\mathbb{P}$ to $\mathbb{Q}$. Theoretically, if $\mathbb{Q}$ is an equivalent martingale measure, the discounted stock price is a martingale. To assess whether this martingale property also holds in CP2022, we perform Monte Carlo simulations using the same settings as described above. Almost all simulated values fall within the 95% confidence interval. In this scenario, the discretization error dominates the total numerical error. When using 100 time steps per year for the first 10 years, the results align with expectations.

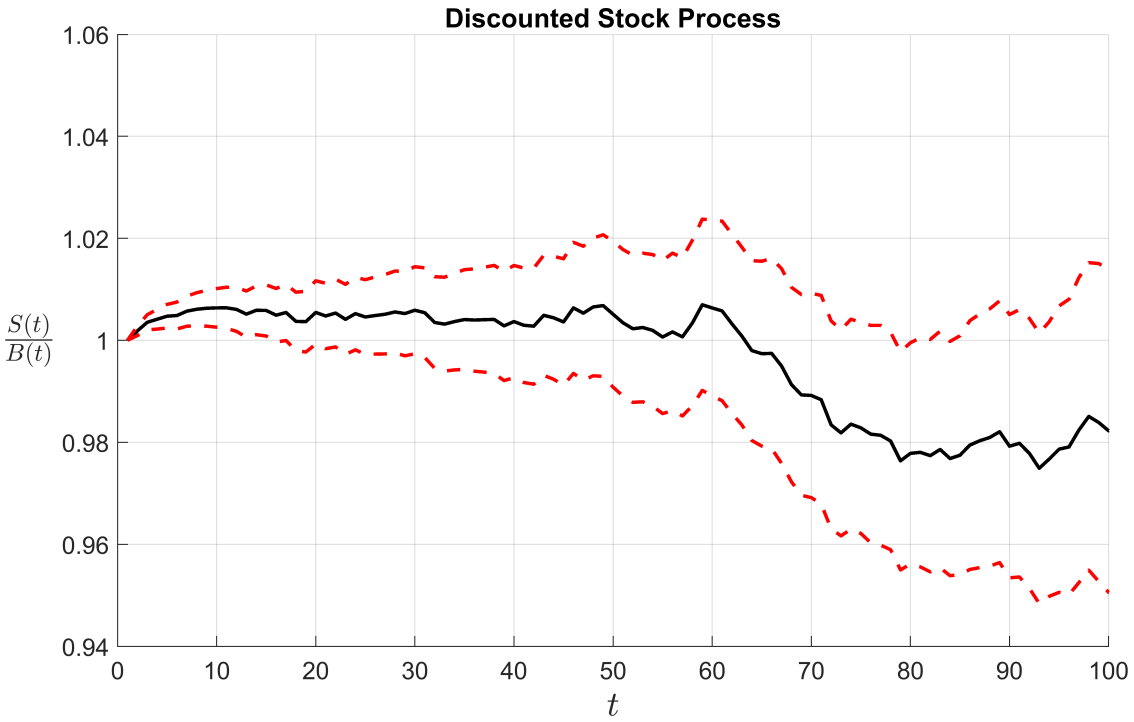

(a)

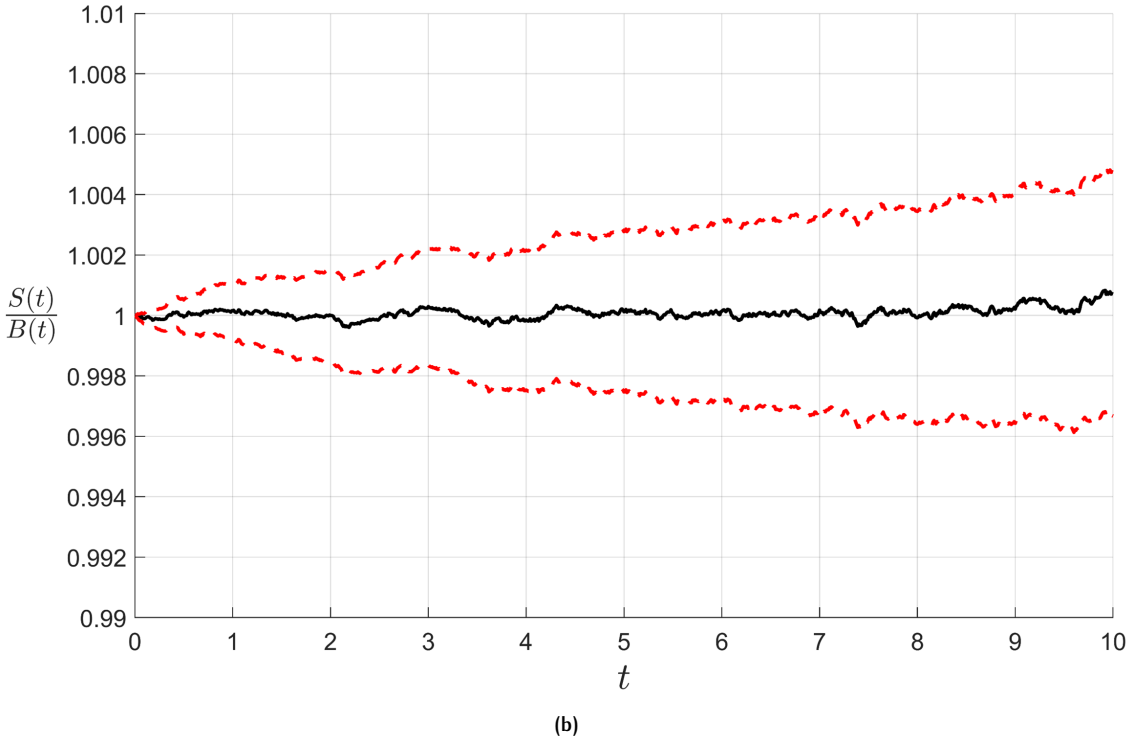

(b)

Figure 3.2: Discounted stock $S(t)/B(t)$ with $S(0)/B(0) = 1$ against time t using monthly time steps (a) and 100 time steps per year (b).

Additionally, we evaluate whether CP2022 can replicate the empirically observed skewness in the probability density functions of stock returns and the consumer price index under $\mathbb{P}$.

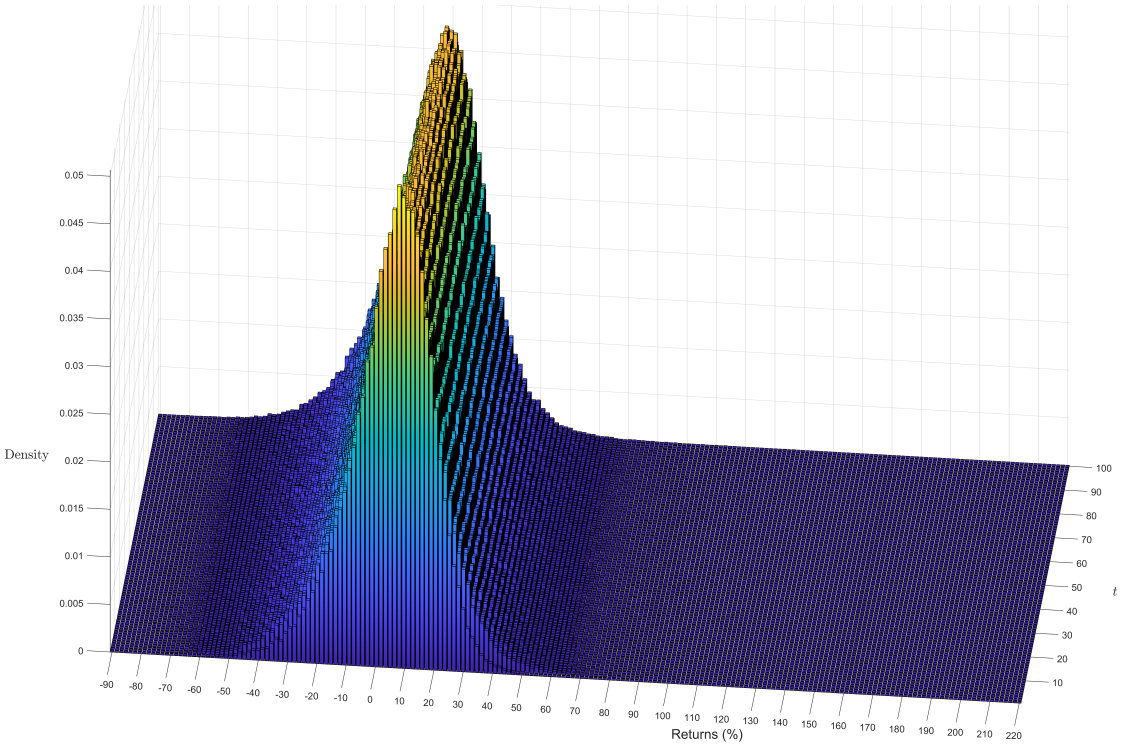

Figure 3.3: Empirical probability density function of stock returns $(S_t - S_{t-1}/S_t)$ projected 100 years in the future.

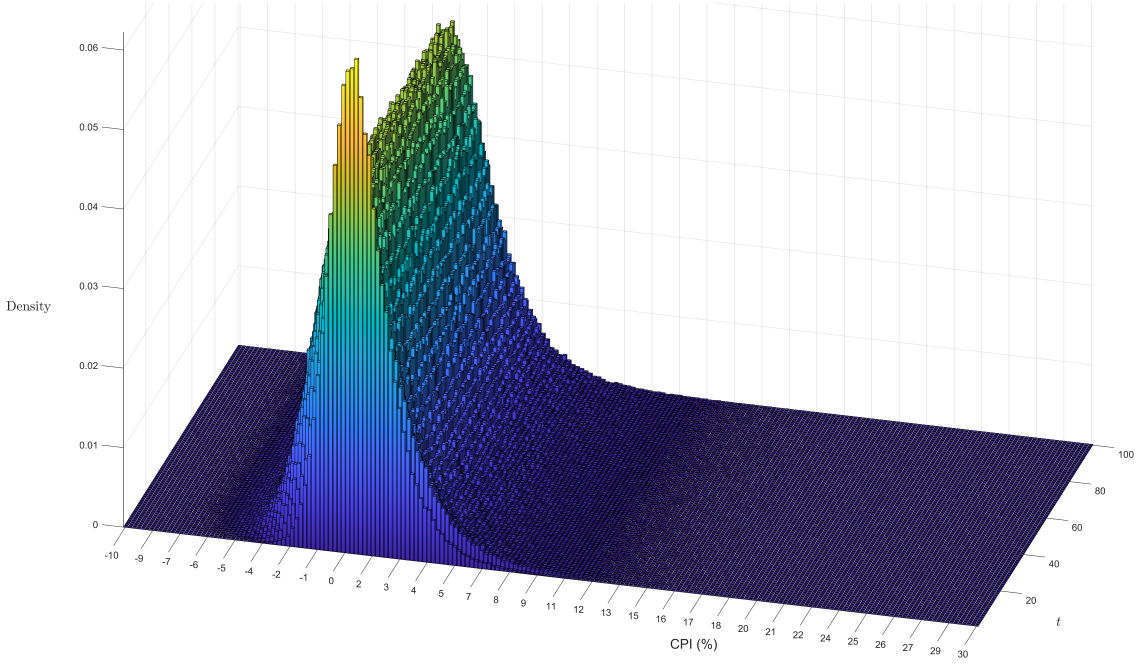

Figure 3.4: Empirical probability density function of CPI changes $(\Pi_t - \Pi_{t-1}/\Pi_t)$ projected 100 years in the future.

Stock returns are left-skewed, with more frequent and severe losses than gains, while inflation rates are right-skewed, as high inflation is more common than high deflation, as empirically observed. These empirical patterns are also visible in the tails of both distributions. Additionally, the long-term average stock returns and changes in CPI align with their imposed constraints of 5.4% and 2.0% as in Equation 3.41 and Equation 3.42 respectively.

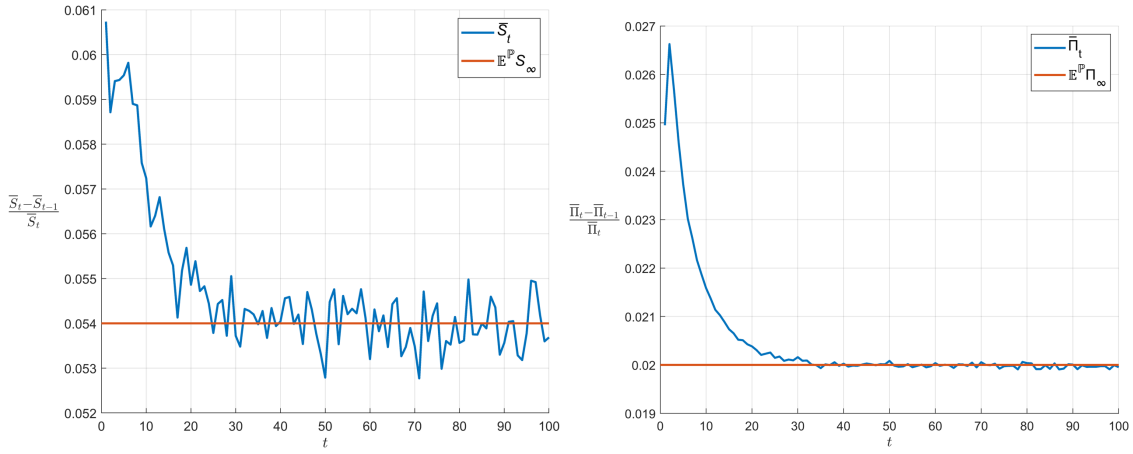

Figure 3.5: Geometric average of stock returns (a) and changes in CPI (b) projected 100 years in the future with their respective imposed constraints.

This page was intentionally left blank.

4 A FAST OPTION VALUATION METHOD IN CP2022

"Unless you're really, really sure that another set of basis functions is better, use Chebyshev polynomials."

John P. Boyd, 2001, from Boyd [2001].

In this chapter, we discuss the method of using Chebyshev polynomials as basis functions in the context of approximation theorem. Then, we illustrate a generalization of the Fourier-Cosine Series Expansion (COS) method, and accelerate option pricing in CP2022 by combining both methods. Finally, we will give an error analysis of the numerical methods used.

4.1 CHEBYSHEV POLYNOMIALS

Chebyshev polynomials, initially introduced by Chebyshev [1853], form a sequence of orthogonal polynomials that have a variety of applications in numerical analysis, such as in approximation theory, or spectral methods used to solve ordinary and partial differential equations. These polynomials can be defined recursively, which together with the orthogonal property, make them efficient to compute.

Definition 4.1.1 (Chebyshev polynomial)**.** *The Chebyshev polynomial $T_n(x)$ of the first kind is a polynomial of degree n, defined as*

$$T_n(\cos(\theta)) = \cos(n\theta), \quad x \in [-1, 1]$$

or, equivalently as

$$T_n(x) = \cos(n \arccos(x)).$$

We show the recurrence relation to obtain the first n Chebyshev polynomials of the first kind. First, note that

$$\begin{aligned} T_0(x) &= \cos(0) = 1, \\ T_1(x) &= \cos(\arccos(x)) = x. \end{aligned}$$

By making the substitution $x = \cos(t)$, we obtain:

$$\begin{aligned} T_1(t) &= \cos(\arccos(\cos(t))) = \cos(t), \\ T_n(t) &= \cos(n \arccos(\cos(t))) = \cos(nt), \\ T_{n-1}(t) &= \cos((n-1) \arccos(\cos(t))) = \cos((n-1)t), \\ T_{n+1}(t) &= \cos((n+1) \arccos(\cos(t))) = \cos((n+1)t). \end{aligned}$$

We use the following well-known trigonometric identities, also known as the angle sum formulas (see for example Bromwich [2005]):

$$\begin{aligned} \cos((n-1)x + x) &= \cos((n-1)x)\cos(x) - \sin((n-1)x)\sin(x), \\ \cos((n-1)x - x) &= \cos((n-1)x)\cos(x) + \sin((n-1)x)\sin(x). \end{aligned}$$

Adding the two identities above yields:

$$\cos(nx) + \cos((n-2)x) = 2\cos(x)\cos((n-1)x).$$

With respect to the Chebyshev polynomials, using the identity from above gives:

$$\cos((n-1)t) + \cos((n+1)t) = 2\cos(nt)\cos(t).$$

This yields the following recursive relationship for the Chebyshev polynomials of the first kind:

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x).$$

From the recurrence relation, combined with the initial polynomials $T_0(x)$ and $T_1(x)$, it is clear that the Chebyshev polynomials as defined in Definition 4.1.1 are indeed polynomials. The orthogonal property follows from the orthogonality of the cosine functions:

$$\int_0^\pi \cos(m\theta)\cos(n\theta)d\theta = \begin{cases} 0 & (m \neq n) \\ \pi/2 & (m = n \neq 0) \\ \pi & (m = n = 0) \end{cases}$$

Using the substitution $\cos\theta = x$ in the Chebyshev polynomials gives:

$$\int_{-1}^1 \frac{T_m(x)T_n(x)}{\sqrt{1-x^2}}dx = \begin{cases} 0 & (m \neq n) \\ \pi/2 & (m = n \neq 0) \\ \pi & (m = n = 0) \end{cases}$$

In other words, the Chebyshev polynomials of the first kind are orthogonal on $x \in [-1, 1]$ with the weighting function $(1-x^2)^{-1/2}$.

We can represent any continuous function $f(x)$ with $x \in [-1, 1]$ by using Chebyshev polynomials as basis functions, as justified by the Weierstrass approximation theorem (Weierstrass [1885]):

$$f(x) = \sum_{n=0}^{\infty}{}' a_n T_n(x),$$

where in $\sum'$ the first term is halved, and the coefficients a_n can be obtained from the orthogonal property:

$$a_n = \frac{2}{\pi}\int_{-1}^1 \frac{f(x)T_n(x)dx}{\sqrt{1-x^2}}dx. \tag{4.1}$$

However, calculating these coefficients is computationally expensive and time-consuming when $f(x)$ is unknown. Instead, we use another orthogonal property. Since

$$\cos\left((2k-1)\frac{\pi}{2}\right) = 0, \qquad k \in \mathbb{N}_{>0}$$

the roots of $T_n(x)$ are

$$x_k = \cos\left(\frac{(2k-1)}{n}\frac{\pi}{2}\right) \equiv \cos(\theta_k). \tag{4.2}$$

Define

$$s_n(\theta) = \sum_{k=1}^{n}\cos\left(\left(k-\frac{1}{2}\right)\theta\right) = \mathrm{Re}\left(\sum_{k=0}^{n-1} e^{i(k+\frac{1}{2})\theta}\right) = \mathrm{Re}\left(e^{i\frac{\theta}{2}}\sum_{k=0}^{n-1} e^{ik\theta}\right). \tag{4.3}$$

Note that the right-hand side is a geometric series:

$$\sum_{k=0}^{n-1} z^k = \frac{1-z^n}{1-z},$$

with $z = e^{i\theta}$ and using

$$\sin(x) = \frac{e^{ix} - e^{-ix}}{2i},$$

yields

$$\begin{aligned}
\mathrm{Re}\left(e^{i\frac{\theta}{2}}\sum_{k=0}^{n-1} e^{ik\theta}\right) &= \mathrm{Re}\left(e^{i\frac{\theta}{2}}\frac{1-e^{in\theta}}{1-e^{i\theta}}\right) \\
&= \mathrm{Re}\left(\frac{1}{2i}\frac{2i}{e^{-i\frac{\theta}{2}} - e^{i\frac{\theta}{2}}}\right) - \mathrm{Re}\left(e^{in\theta}\frac{1}{2i}\frac{2i}{e^{-i\frac{\theta}{2}} - e^{i\frac{\theta}{2}}}\right) \\
&= \mathrm{Re}\left(\frac{1}{2i}\frac{1}{\sin\left(-\frac{\theta}{2}\right)}\right) - \mathrm{Re}\left(\frac{1}{2i}e^{in\theta}\frac{1}{\sin\left(-\frac{\theta}{2}\right)}\right) \\
&= \frac{\mathrm{Re}\left(\frac{e^{in\theta}}{2i}\right) - \mathrm{Re}\left(\frac{1}{2i}\right)}{\sin\left(\frac{\theta}{2}\right)} \\
&= \frac{\sin(n\theta)}{2\sin\left(\frac{\theta}{2}\right)}.
\end{aligned}$$

As a result,

$$s_n(\theta) = \sum_{k=1}^{n} \cos\left(\left(k - \frac{1}{2}\right)\theta\right) = \frac{\sin(n\theta)}{2\sin\left(\frac{\theta}{2}\right)}. \tag{4.4}$$

Note the following properties of $s_n(\theta)$. From the right-side of Equation 4.4,

$$s_n(\theta_l) = 0, \quad \theta_l = \frac{l\pi}{n}, \quad l \in (0, 2n),$$

and from the left-side of Equation 4.4,

$$s_n(0) = n.$$

Using the roots of T_n,

$$\begin{aligned}
a_{ij} &= \sum_{k=1}^{n} \cos(i\theta_k)\cos(j\theta_k) \\
&= \frac{1}{2}\sum_{k=1}^{n} \left(\cos((i+j)\theta_k) - \cos((i-j)\theta_k)\right) \\
&= \frac{1}{2}\left(s_n\left(\frac{i+j}{n}\pi\right) + s_n\left(\frac{i-j}{n}\pi\right)\right).
\end{aligned}$$

Using the above-mentioned properties of $s_n(\theta)$ yields

$$a_{ij} = \begin{cases} 0 & i \neq j,\ 0 < i, j \leq n-1 \\ \frac{1}{2}n & i = j,\ 0 < i \leq n-1 \\ n & i = j = 0 \end{cases}$$

and thus the following orthogonal property follows, using the roots x_k of T_n:

$$\sum_{k=1}^{n} T_i(x_k)T_j(x_k) = \begin{cases} 0 & i \neq j,\ 0 < i, j \leq n-1 \\ \frac{1}{2}n & i = j,\ 0 < i \leq n-1 \\ n & i = j = 0 \end{cases} \tag{4.5}$$

This orthogonal property yields an efficient scheme to approximate a function.

Lemma 4.1.2 (Chebyshev coefficients). *Suppose we represent a continuous function $f(x)$ with $x \in [-1,1]$ using Chebyshev polynomials of the first kind as $f(x) = \sum_{n=0}^{\infty}{}' a_n T_n(x)$ and truncate the series using N terms, that is, we approximate $f(x)$ by*

$$\tilde{f}(x) = \sum_{n=0}^{N-1}{}' a_n T_n(x).$$

Let x_k be the roots of T_N. Then

$$a_n = \frac{2}{N} \sum_{k=1}^{N} f(x_k) T_n(x_k).$$

Proof. Set $f(x_k) = \tilde{f}(x_k) = \sum_{n=0}^{N-1}{}' a_n T_n(x_k)$. Multiply both sides by $\frac{2}{N} T_m(x_k)$ and summing both sides yields

$$\frac{2}{N} \sum_{k=1}^{N} f(x_k) T_m(x_k) = \sum_{n=0}^{N-1}{}' a_n \left(\frac{2}{N} \sum_{k=1}^{N} T_n(x_k) T_m(x_k) \right) = a_m,$$

which follows from Equation 4.5. □

We can show the rate of convergence of Chebyshev polynomials when interpolating a function $f(x)$. To do this, we use the following preliminary result.

Lemma 4.1.3 (Series truncation error of an algebraically converging series). *The series truncation error of an algebraically converging series behaves like*

$$\sum_{k=N}^{\infty} \frac{1}{k^n} \sim \frac{1}{(n-1)(N-1)^{n-1}}.$$

Proof. See Bender and Orszag [2010]. □

Lemma 4.1.4 (Convergence rate Chebyshev polynomials). *Suppose $f(x)$ is m-times continuously differentiable and $f^{(m)}$ is bounded. Suppose we approximate the function by $\tilde{f}(x)$ with a truncated N Chebyshev series expansion. Then:*

$$|f(x) - \tilde{f}(x)| = \mathcal{O}\left(\frac{1}{N^m} \right).$$

Proof. First, remark that $|T_n(x)| \leq 1$. Then

$$|f(x) - \tilde{f}(x)| = \left| \sum_{n=0}^{\infty}{}' a_n T_n(x) - \sum_{n=0}^{N-1}{}' a_n T_n(x) \right| = \left| \sum_{n=N}^{\infty} a_n T_n(x) \right| \overset{(*)}{\leq} \sum_{n=N}^{\infty} |a_n|$$

where in $(*)$ we use the triangle inequality with the first remark. Using Equation 4.1 with $x = \cos(\theta)$ yields

$$a_n = \frac{2}{\pi} \int_0^{\pi} f(\cos(\theta)) \cos(n\theta) d\theta.$$

Integrating by parts gives

$$a_n = \frac{2}{n\pi} \int_0^{\pi} f'(\cos(\theta)) \sin(n\theta) \sin(\theta) d\theta.$$

Using the trigonometric identity $\sin(x)\sin(y) = \frac{1}{2}(\cos(x-y) - \cos(x+y))$ gives

$$a_n = \frac{1}{n\pi} \int_0^{\pi} f'(\cos(\theta))(\cos((n-1)\theta) - \cos((n+1)\theta)) d\theta.$$

Repeating m-times and bounding $f^{(m)}$ by $\max_{\theta\in[0,\pi)} f^{(m)}(\cos(\theta))$ shows

$$a_n \sim \mathcal{O}\left(\frac{1}{n^{m+1}}\right).$$

The result then follows from Lemma 4.1.3. □

Lemma 4.1.4 shows that using Chebyshev polynomials for interpolation has an algebraic convergence rate of order m. If f is infinitely differentiable with no singularities on the complex plane, Chebyshev polynomials exhibit geometric convergence for non-periodic infinitely differentiable functions with no singularities, and algebraic convergence otherwise (Boyd [2001]).

Chebyshev polynomials also exhibit exponential convergence when used in an integration scheme like the Clenshaw–Curtis quadrature integration rule (Trefethen [2019]). A few variations exist in the literature. Fejér quadrature approximates an integral by evaluating the roots of the Chebyshev polynomial, whereas Clenshaw–Curtis quadrature evaluates at the extreme points of the Chebyshev polynomials. We only consider the latter, since it is more accurate (Elliott [1965]).

Lemma 4.1.5 (Clenshaw–Curtis quadrature)**.** *The extreme points of $T_n(x)$ are*

$$x_k = \cos\left(\frac{k\pi}{n}\right), \quad k \in \{0, ..., n\}.$$

Then, using a cosine-series expansion of $f(\cos(\theta))$,

$$\int_{-1}^{1} f(x)dx = \int_0^{\pi} f(\cos(\theta))\sin(\theta)d\theta = \int_0^{\pi} \left(\sum_{k=0}^{\infty}{}'' w_k \cos(k\theta)\right)\sin(\theta)d\theta \approx \sum_{k=0}^{n} w_k f(x_k),$$

where

$$w_k = \frac{c_k}{n}\left(1 - \sum_{j=1}^{\lceil n/2\rceil} \frac{2 - \mathbb{1}_{\{j=n/2\}}}{4j^2-1}\cos(2jx_k)\right), \quad k \in \{0, ..., n\},$$

and

$$c_k = \begin{cases} 1 & \text{if } k = 0 \pmod n, \\ 2 & \text{else}. \end{cases}$$

Proof. See Waldvogel [2006]. □

Note that this scheme scales as $\mathcal{O}(n^2)$, but since the weights (including in the interpolating case) are equivalent to performing a Fast Fourier Transform (FFT), efficient FFT algorithms can be employed to reduce the complexity to $\mathcal{O}(n\log(n))$. For details on the connection between FFT and Clenshaw-Curtis quadrature, we refer to Waldvogel [2006]. For interpolating a function with Chebyshev polynomials, the coefficients in 4.1.2 can be obtained with a type-I Discrete Cosine Transform (DCT).

Up till now we have considered interpolation and integration on $[-1, 1]$. For interpolation and integration on any interval $[a, b]$, the linear transformation $\tilde{x} = \frac{b+a}{2} + \frac{b-a}{2}x$ is used and all previous results hold.

In particular, for interpolation, the roots of T_n are

$$\tilde{x}_k = \frac{b+a}{2} + \frac{b-a}{2}\cos\left(\frac{(2k-1)}{n}\frac{\pi}{2}\right),$$

and thus

$$\tilde{f}(\tilde{x}) = \sum_{n=0}^{N-1}{}'' a_n T_n(\tilde{x}),$$

with

$$a_n = \frac{2}{N}\sum_{k=1}^{N} f(\tilde{x}_k)T_n(x_k).$$

For integration, we have

$$\int_a^b f(x)dx = \frac{b-a}{2}\int_{-1}^{1} f(\tilde{x})\,d\tilde{x}.$$

4.2 FOURIER–COSINE SERIES EXPANSION METHOD

Since CP2022 is affine, the characteristic function can be obtained using Equation 3.18. As a result, efficient Fourier-based option pricing algorithms can be used.

We will illustrate how the original Fourier-cosine series expansion method (COS method) can be generalized when interest rates are stochastic, without changing measures or increasing the dimensionality of the option pricing problem. The typical approach is to switch to the equivalent martingale measure $\mathbb{Q}^T$ with the zero-coupon bond $P(t,T)$ as numéraire with Radon-Nikodym density $\frac{d\mathbb{Q}^T}{d\mathbb{Q}}\Big|_{\mathcal{F}_t} = \frac{P(t,T)B(t_0)}{P(t_0,T)B(t)}$. As a result, the price of a European call option on S_t equals

$$\mathbf{C}(t_0,T,K) = \mathbb{E}_{t_0}^{\mathbb{Q}}\left[e^{-\int_{t_0}^{T} r_u du}\left(S_T-K\right)^+\right] = P(t_0,T)\mathbb{E}_{t_0}^{\mathbb{Q}^T}\left[\left(S_T-K\right)^+\right].$$

This representation is convenient only when the Riccati ODEs admit an analytical solution for ϕ and Ψ such that $P(t_0,T)$ can be determined completely analytically. In our case, the Riccati ODEs must be solved numerically. Using this representation would require numerically solving extra ODEs to determine $P(t_0,T)$, which slows down the numerical computation of $\mathbf{C}(t_0,T,K)$. Therefore, an alternative approach will be used. First, we present the original COS method.

In a similar fashion to Chebyshev polynomials, any real function $f(x)$ on $[a,b]$ with finite support can be expressed as a cosine series (Boyd [2001]):

$$f(x) = \sum_{k=0}^{\infty}{}' A_k \cos\left(k\pi\frac{x-a}{b-a}\right), \tag{4.6}$$

and the Fourier coefficients A_k are determined using the orthogonal property of cosine functions:

$$\int_a^b f(x)\cos\left(m\pi\frac{x-a}{b-a}\right)dx \overset{(*)}{=} \sum_{k=0}^{\infty}\int_a^b A_k\cos\left(k\pi\frac{x-a}{b-a}\right)\cos\left(m\pi\frac{x-a}{b-a}\right)dx = \frac{b-a}{2}A_m,$$

where in $(*)$, Fubini's theorem is used. This gives directly the Fourier coefficients

$$A_k = \frac{2}{b-a}\int_a^b f(x)\cos\left(k\pi\frac{x-a}{b-a}\right)dx. \tag{4.7}$$

Assume f is an unknown probability density function, but its characteristic function Φ is known. Φ is defined on $\mathbb{R}$, but assume we truncate the characteristic function to $[a,b]$ such that $\tilde{\Phi}$ approximates Φ well:

$$\tilde{\Phi}(w) \equiv \int_a^b e^{iwx}f(x)dx \approx \int_{\mathbb{R}} e^{iwx}f(x)dx \equiv \Phi(w). \tag{4.8}$$

Comparing Equation 4.7 with Equation 4.8 and using Euler's identity shows that

$$A_k = \frac{2}{b-a}\,\mathrm{Re}\left(\tilde{\Phi}\left(\frac{k\pi}{b-a}\right)e^{-k\pi i\frac{a}{b-a}}\right). \tag{4.9}$$

We make the following approximations. Define

$$F_k = \frac{2}{b-a}\,\mathrm{Re}\left(\Phi\left(\frac{k\pi}{b-a}\right)e^{-k\pi i\frac{a}{b-a}}\right). \tag{4.10}$$

Suppose $[a,b]$ is chosen such that $\tilde{\Phi}(w) \approx \Phi(w)$ on $[a,b]$, and thus $F_k \approx A_k$. Truncating the Fourier-cosine expansion in Equation 4.6 to N terms provides the following approximation of the probability density function:

$$\tilde{f}(x) = \sum_{k=0}^{N-1}{}' F_k \cos\left(k\pi\frac{x-a}{b-a}\right). \tag{4.11}$$

For functions without singularities in the complex plane, the Fourier-cosine series exhibits exponential convergence (Boyd [2001]).

Similarly, the cumulative distribution function can be obtained by integrating Equation 4.11:

$$\tilde{F}(x) = \frac{1}{2}(x-a)F_0 + \sum_{k=1}^{N-1} F_k \frac{b-a}{k\pi}\sin\left(k\pi\frac{x-a}{b-a}\right). \tag{4.12}$$

4.3 PRICING EUROPEAN CALLS AND PUTS ON STOCK INDEX AND CPI IN CP2022

In this section, we illustrate a generalization of the COS method for pricing European options on the stock index and CPI in CP2022. Define

$$\xi(u,v,\mathbf{X}_t,t) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{u'\mathbf{X}_T + v'\int_t^T \mathbf{X}_s ds}\right] = e^{\phi_{uv}(t,T)+\Psi_{uv}(t,T)'\mathbf{X}_t} \tag{4.13}$$

as in Equation 3.18, with the model parameters as in Equation 3.25. We consider a European call and put with underlying S_t, maturity T and strike K. Let $w \in \mathbb{R}$ and define e_j as the 5×1 standard unit vector, where position j contains 1 and all other entries equal 0:

$$e_j = \begin{bmatrix}\mathbb{1}_{\{j=1\}} & \mathbb{1}_{\{j=2\}} & \mathbb{1}_{\{j=3\}} & \mathbb{1}_{\{j=4\}} & \mathbb{1}_{\{j=5\}}\end{bmatrix}'. \tag{4.14}$$

If we choose $v = -e_2$ and $u = iwe_4$, we can denote the characteristic function of $\ln(S_t)$ by

$$\xi(iwe_4, -e_2, \mathbf{X}_t, t) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_s ds + iw\ln(S_t)}\right]. \tag{4.15}$$

Using the martingale pricing approach, the option price $V(\mathbf{X}_t,t)$ conditioned on the initial state $\mathbf{X}_t$, is given by

$$\begin{aligned} V(\mathbf{X}_t,t) &= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_s}V(x,T)\right] \\ &= \int_{\mathbb{R}}\int_{\mathbb{R}} V(x,T)e^{-y}f_{(x,y)}(x,y)dxdy \\ &= \int_{\mathbb{R}} V(x,T)\left(\int_{\mathbb{R}} e^{-y}f_{(x,y)}(x,y)dy\right)dx, \end{aligned}$$

where $x = \ln(S_T/K), y = \int_t^T r_s ds$ and $f_{(x,y)}(x,y)$ is the conditional joint probability density function given the initial state $\mathbf{X}_t$. Using the characteristic function ξ as in Equation 4.15, we have

$$e^{-iw\ln K}\xi(iwe_4,-e_2,\mathbf{X}_t,t)=\int_{\mathbb{R}}\int_{\mathbb{R}}e^{iwx-y}f_{(x,y)}(x,y)dxdy=\int_{\mathbb{R}}e^{iwx}\left(\int_{\mathbb{R}}e^{-y}f_{(x,y)}(x,y)dy\right)dx$$
$$\equiv\int_{\mathbb{R}}e^{iwx}g(x)dx,$$

where we have defined $g(x)=\int_{\mathbb{R}}e^{-y}f_{(x,y)}(x,y)dy$. Now we can proceed as in Section 4.2. Assume $[a,b]$ is chosen such that the truncated $\tilde{\xi}$ approximates its infinite counterpart ξ well. Performing a Fourier-cosine series expansion of $g(x)$ shows that the Fourier coefficients are given by

$$A_k=\frac{2}{b-a}\,\mathrm{Re}\left(\tilde{\xi}\left(\frac{k\pi i}{b-a}e_4,-e_2,\mathbf{X}_t,t\right)e^{-k\pi i\frac{\ln(K)+a}{b-a}}\right),\tag{4.16}$$

and we define F_k as the infinite counterpart of A_k using ξ instead of $\tilde{\xi}$. Then, the option price is

$$\begin{aligned}V(\mathbf{X}_t,t)&\approx V_1(\mathbf{X}_t,t)\\&=\int_a^b V(x,T)g(x)dx\\&=\int_a^b V(x,T)\sum_{k=0}^{\infty}{}'A_k\cos\left(k\pi\frac{x-a}{b-a}\right)dx\\&\overset{(*)}{=}\sum_{k=0}^{\infty}{}'A_k\int_a^b V(x,T)\cos\left(k\pi\frac{x-a}{b-a}\right)dx\\&\equiv\sum_{k=0}^{\infty}{}'A_kB_k,\end{aligned}$$

where in $(*)$, Fubini's theorem is used. For European calls and puts, the payoff at maturity is $V(x,T)=\max(\theta K(e^x-1),0)$, where $\theta=1$ for a call and $\theta=-1$ for a put. In this case, B_k can be calculated analytically, using integration by parts where needed. Letting $[c,d]\subseteq[a,b]$,

$$\begin{aligned}\chi_k(c,d)&\equiv\int_c^d e^x\cos\left(k\pi\frac{x-a}{b-a}\right)dx\\&=\frac{1}{1+\left(\frac{k\pi}{b-a}\right)^2}\left(\cos\left(k\pi\frac{d-a}{b-a}\right)e^d-\cos\left(k\pi\frac{c-a}{b-a}\right)e^c\right.\\&\left.+\frac{k\pi}{b-a}\sin\left(k\pi\frac{d-a}{b-a}\right)e^d-\frac{k\pi}{b-a}\sin\left(k\pi\frac{c-a}{b-a}\right)e^c\right)\end{aligned}\tag{4.17}$$

and

$$\begin{aligned}\psi_k(c,d)&\equiv\int_c^d\cos\left(k\pi\frac{x-a}{b-a}\right)dx\\&=\begin{cases}\frac{b-a}{k\pi}\left(\sin\left(k\pi\frac{d-a}{b-a}\right)-\sin\left(k\pi\frac{c-a}{b-a}\right)\right), & k\neq 0\\(d-c), & k=0.\end{cases}\end{aligned}\tag{4.18}$$

There are three cases. For $a<0<b$:

$$B_k=\begin{cases}K(\chi_k(0,b)-\psi_k(0,b)), & \theta=1,\\K(-\chi_k(a,0)+\psi_k(a,0)), & \theta=-1,\end{cases}\tag{4.19}$$

while for $a<b<0$:

$$B_k=\begin{cases}0, & \theta=1,\\K(-\chi_k(a,b)+\psi_k(a,b)), & \theta=-1,\end{cases}\tag{4.20}$$

and for $0 < a < b$:

$$B_k = \begin{cases} K(\chi_k(a,b) - \psi_k(a,b)), & \theta = 1, \\ 0, & \theta = -1. \end{cases} \tag{4.21}$$

Afterwards, we truncate $V_1(\mathbf{X}_t, t)$ using N terms, which gives numerical approximation $V_2(\mathbf{X}_t, t)$. Finally, substituting F_k for A_k gives the final numerical approximation

$$V_3(\mathbf{X}_t, t) = \sum_{k=0}^{N-1}{}' F_k B_k. \tag{4.22}$$

In the original work by Fang and Oosterlee [2009a], the truncation range $[a, b]$ is based on the cumulants of ξ, and is set to

$$[a,b] = \left[\lambda_1 - L\sqrt{|\lambda_2| + \sqrt{|\lambda_4|}}, \lambda_1 + L\sqrt{|\lambda_2| + \sqrt{|\lambda_4|}}\right], \tag{4.23}$$

for a user-defined $L \in \mathbb{R}_{>0}$ and

$$\lambda_n = \frac{1}{i^n} \frac{\partial^n (\ln \xi(u, v, \mathbf{X}_t, t))}{\partial u^n}\bigg|_{u=0}. \tag{4.24}$$

Since ξ is determined numerically, λ_n cannot be obtained analytically and must instead be computed using finite differences. As this depends on u, which serves as the terminal condition for the ODEs and must be varied to compute finite differences, the number of ODEs to be solved increases and more computational effort is used. Alternatively, a simple integration range of the form $[-L\sqrt{T}, L\sqrt{T}]$ can be used, where L is large enough such that the truncation error from truncating the distribution is negligible. During calibration, model parameters change with each iteration. Therefore, an adaptive scheme, such as the first approach, is preferred for determining truncation ranges.

For calculating F_k, the characteristic function of CP2022 ξ is required. This is done by numerically solving the set of ODEs in Equation 3.19 and 3.20, at the terminal conditions determined by the COS method: $\phi(T,T) = 0$ and $\Psi(T,T) = \frac{k\pi i}{b-a} e_4$. Pricing zero-coupon inflation-indexed caps and floors is analogous to pricing European calls and puts on S_t, with the only difference being that we replace e_4 with e_5. The numerical algorithm used for solving the ODEs is the Dormand-Prince Runge-Kutta (4,5) method (Dormand and Prince [1980]), which is an algorithm based on the explicit five-stage Runge-Kutta method of order four. The ODEs that are numerically solved can be expressed as[1]

$$\frac{\partial}{\partial t}\tilde{\Psi}_{uv}(t,T)_i = -\frac{1}{2}\tilde{\Psi}_{uv}(t,T)'\Sigma G_i \Sigma' \tilde{\Psi}_{uv}(t,T) + \left(L'\tilde{\Psi}_{uv}(t,T)\right)_i + (e_2)_i, \quad \tilde{\Psi}_{uv}(T,T) = \frac{k\pi i}{b-a} e_4 \tag{4.25}$$

$$\tilde{\phi}_{uv}(t,T) = \int_t^T \left(\tilde{\Psi}_{uv}(s,T)'\tilde{\zeta}(s) + \frac{1}{2}\tilde{\Psi}_{uv}(s,T)'\Sigma G_0 \Sigma' \tilde{\Psi}_{uv}(s,T)\right) ds. \tag{4.26}$$

For calibration purposes, efficiently pricing options is crucial. While the COS method is highly efficient for option pricing, the majority of computation time is spent numerically solving the coupled ODEs. When pricing options with different maturities, the ODEs must be solved separately for each maturity, making the process time-consuming. Given m different maturities of the same type of European option and assuming all prices are computed using the same number

1 Although the solution ϕ is known, we deliberately treat it as an ODE to be solved rather than applying a numerical integration scheme directly to ϕ. This approach helps reduce numerical integration errors when employing a fast numerical integration scheme with a lower order of accuracy than the ODE solver used.

of cosine-series expansion terms N, the coupled ODEs in Equation 4.25 and 4.26 must be solved mN times. This results in significant computational costs, particularly as solving the ODEs with terminal conditions for long-maturity options further increases computation time. Instead, we can use a different method, which requires solving them a total of n times, where $n < N$.

We denote the solutions to the ODEs by $\tilde{\phi}$ and $\tilde{\Psi}$ with time-dependent market prices of risk with model parameter $\tilde{\zeta}(t)$, and by ϕ and Ψ with constant market prices of risk with model parameter $\zeta(t)$.

First, we assume the market prices of risk are constant, $\lambda_0(t) = \lambda_0$ and

$$\zeta(t) = \begin{bmatrix} M\mathbb{E}^{\mathbb{Q}}\mathbf{X}_{\infty}^{s} \\ \mu^{o} - \begin{bmatrix} \eta_S \\ \eta_\Pi \end{bmatrix} \end{bmatrix}.$$

Since all coefficients in the ODEs are constant, including $\zeta(t)$, we have that $\phi_{uv}(t,T) = \phi_{uv}(\tau)$ and $\Psi_{uv}(t,T) = \Psi_{uv}(\tau)$, with $\tau = T - t$. The ODEs can be written in terms of time to maturity τ as

$$-\frac{\partial}{\partial\tau}\Psi_{uv}(\tau)_i = -\frac{1}{2}\Psi_{uv}(\tau)'\Sigma G_i\Sigma'\Psi_{uv}(\tau) + \left(L'\Psi_{uv}(\tau)\right)_i + (e_2)_i, \quad \Psi_{uv}(0) = \frac{k\pi i}{b-a}e_4, \tag{4.27}$$

$$\phi_{uv}(\tau) = \int_t^{t+\tau}\left(\Psi_{uv}(t+\tau-s)'\zeta(s) + \frac{1}{2}\Psi_{uv}(t+\tau-s)'\Sigma G_0\Sigma'\Psi_{uv}(t+\tau-s)\right)ds. \tag{4.28}$$

In our time-homogeneous ODEs, we only need to solve the equations once, up to the maturity of the longest-dated option. The solutions for shorter-maturity options are then directly available. This holds as long as the market prices of risk remain constant, which is not the case in $\mathbb{Q}$. However, the change in the market price of risk affects the ODEs only through

$$\tilde{\zeta}(t) = \zeta(t) - [0 \quad f(t) \quad 0 \quad 0 \quad f^R(t)]'\mathbb{1}_{\{t\geq t_0\}}.$$

Fortunately, $\tilde{\Psi}_{uv}$ remains unaffected by $\tilde{\zeta}(t)$, such that $\tilde{\Psi}_{uv}(\tau) = \Psi_{uv}(\tau)$. To obtain $\tilde{\phi}_{uv}(t,t+\tau)$, we subtract Equation 4.26 from Equation 4.28 to get

$$\tilde{\phi}_{uv}(t,t+\tau) = \phi_{uv}(t,t+\tau) - \int_t^{t+\tau}\Psi_{uv}(t+\tau-s)'[0 \quad f(s) \quad 0 \quad 0 \quad f^R(s)]'ds. \tag{4.29}$$

Using this approach, solving the ODEs in Equation 4.27 and 4.28 yields Ψ and ϕ, which can then be used to compute $\tilde{\phi}$ via Equation 4.29. Then, Ψ and $\tilde{\phi}$ can subsequently be used to determine the characteristic function ξ from Equation 4.13. In short, since our time-inhomogeneous affine model exhibits time dependence through only one ODE, for which the solution is known, we can transform the system of ODEs to a time-homogeneous form with an additional numerical integration scheme over the time-dependent component. This allows the characteristic function to be determined for all time points up to the terminal condition.

To obtain $\tilde{\phi}$, a numerical integration scheme is required. We consider two integration schemes. First, we use the trapezoidal rule with time steps that align with the discontinuities of $f(t)$ and $f^R(t)$, namely $h = 1/12$. The second scheme we consider is the Clenshaw-Curtis quadrature. Note that because $f(t)$ and $f^R(t)$ are discontinuous, the Clenshaw-Curtis quadrature will not converge exponentially anymore. Therefore, the interval of integration is partitioned at the discontinuities and each sub-interval is integrated separately. Using FFT, this scheme is still efficient. We will compare both integration methods in terms of accuracy and computational speed.

We have reduced the computational effort by decreasing the number of times the ODEs need to be solved from mN to N. Further efficiency gains can be achieved by solving the ODEs at $n < N$ points and then reconstructing the characteristic function using Chebyshev interpolation. Given that characteristic functions are uniformly continuous, Chebyshev interpolation is expected to

exhibit exponential convergence. We refer to this approach as the "Cheby-COS method". The algorithm is summarized below.

Algorithm 4.1: Fast option valuation of European calls and puts on stock index and CPI in CP2022 using the Cheby-COS method

Initialization: N_{Cheby}, N_{COS};

1. Solve the ODEs in Equation 4.27 and 4.28 at the roots of the Chebyshev polynomial $T_{N_{Cheby}}(x)$ for $\tau \in [0, T_{max}]$, where T_{max} is the longest maturity of the options.
2. Use a numerical integration scheme to obtain $\tilde{\phi}_{uv}(T_j, T_{max})$ for each maturity T_j as in Equation 4.29 and determine the characteristic function ξ from Equation 4.13 at the roots of the Chebyshev polynomial $T_{N_{Cheby}}(x)$.
3. Approximate ξ using N_{Cheby} Chebyshev polynomials as in Section 4.1.
4. Use the Chebyshev polynomial approximation of ξ to calculate option prices using the COS method with N_{COS} cosine-series expansion terms.

4.4 PRICING YEAR-ON-YEAR INFLATION-INDEXED CAPS AND FLOORS IN CP2022

In this section, we present a method for pricing year-on-year inflation-indexed caps and floors in CP2022 using the COS method. A cap can be interpreted as a call option on the inflation rate, while a floor corresponds to a put option on the same underlying. These options belong to the class of cliquet options, which is consists of a series of consecutive forward-starting options. The option price has the form

$$V(\mathbf{X}_t, t) = \sum_{T_j=t+1}^{T} \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^{T_j} r_u du} \left(\theta \left(e^{\ln \Pi_{T_j} - \ln \Pi_{T_{j-1}}} - 1 - K \right) \right)^+ \right], \tag{4.30}$$

where $\theta = 1$ for caps and $\theta = -1$ for floors. Since each term represents a forward-starting option, our objective is to value each term using the COS method. The first term in Equation 4.30 corresponds to a standard European option with maturity $T = 1$, so we can proceed as in Section 4.3. In Theorem 3.3.1, we have seen that the characteristic function is determined by the deterministic functions ϕ and Ψ that solve the Riccati ODEs. Additionally, for all $t \leq T_1 \leq T_2$, the characteristic function of the joint distribution $(\mathbf{X}_{T_1}, \mathbf{X}_{T_2})$ can be obtained without modifying the Riccati ODEs that need to be solved.

Theorem 4.4.1 (Transitional characteristic function of affine processes). *Let* $0 \leq t \leq T_1 \leq T_2$. *Assume the same conditions as in Theorem 3.3.1. Then*

$$\begin{aligned}\xi_\Delta(u_1,u_2,v_1,v_2,\mathbf{X}_t,T_1,T_2) &= \mathbb{E}_t\left[e^{u_1'\mathbf{X}_{T_1}+v_1'\int_t^{T_1}\mathbf{X}_s ds+u_2'\mathbf{X}_{T_2}+v_2'\int_t^{T_2}\mathbf{X}_s ds}\right]\\ &= e^{\phi(T_1,T_2,u_2,v_2)+\phi(t,T_1,u_1'+\Psi(T_1,T_2,u_2,v_2)',v_1+v_2)+\Psi(t,T_1,u_1'+\Psi(T_1,T_2,u_2,v_2)',v_1+v_2)'\mathbf{X}_t}.\end{aligned} \tag{4.31}$$

Proof.

$$\begin{aligned}\mathbb{E}_t\left[e^{u_1'\mathbf{X}_{T_1}+v_1'\int_t^{T_1}\mathbf{X}_s ds+u_2'\mathbf{X}_{T_2}+v_2'\int_t^{T_2}\mathbf{X}_s ds}\right] &\overset{\text{i}}{=} \mathbb{E}_t\left[\mathbb{E}_{T_1}\left[e^{u_1'\mathbf{X}_{T_1}+v_1'\int_t^{T_1}\mathbf{X}_s ds+u_2'\mathbf{X}_{T_2}+v_2'\int_t^{T_2}\mathbf{X}_s ds}\right]\right]\\ &\overset{\text{ii}}{=} \mathbb{E}_t\left[e^{u_1'\mathbf{X}_{T_1}+v_1'\int_t^{T_1}\mathbf{X}_s ds}\mathbb{E}_{T_1}\left[e^{u_2'\mathbf{X}_{T_2}+v_2'\int_t^{T_2}\mathbf{X}_s ds+v_2'\int_{T_1}^{T_2}\mathbf{X}_s ds}\right]\right]\\ &\overset{\text{ii}}{=} \mathbb{E}_t\left[e^{u_1'\mathbf{X}_{T_1}+(v_1'+v_2')\int_t^{T_1}\mathbf{X}_s ds}\mathbb{E}_{T_1}\left[e^{u_2'\mathbf{X}_{T_2}+v_2'\int_{T_1}^{T_2}\mathbf{X}_s ds}\right]\right]\\ &\overset{3.18}{=} \mathbb{E}_t\left[e^{u_1'\mathbf{X}_{T_1}+(v_1'+v_2')\int_t^{T_1}\mathbf{X}_s ds}e^{\phi(T_1,T_2,u_2,v_2)+\Psi(T_1,T_2,u_2,v_2)'\mathbf{X}_{T_1}}\right]\\ &= e^{\phi(T_1,T_2,u_2,v_2)}\mathbb{E}_t\left[e^{(u_1'+\Psi(T_1,T_2,u_2,v_2)')\mathbf{X}_{T_1}+(v_1'+v_2')\int_t^{T_1}\mathbf{X}_s ds}\right]\\ &\overset{3.18}{=}\end{aligned}$$

$$e^{\phi(T_1,T_2,u_2,v_2)+\phi(t,T_1,u_1'+\Psi(T_1,T_2,u_2,v_2)',v_1+v_2)+\Psi(t,T_1,u_1'+\Psi(T_1,T_2,u_2,v_2)',v_1+v_2)'\mathbf{X}_t},$$

where in step i we use the tower property of the conditional expectation, in step ii we use $\mathcal{F}_{T_1}$-measurability. □

To compute the transitional characteristic function of $(\ln\Pi_{T_j} - \ln\Pi_{T_j-1})$, which we denote by ξ_Δ, we set $u_2 = iwe_5, u_1 = -iwe_5, v_2 = -e_2$ and $v_1 = 0$ in Theorem 4.4.1. Following the approach in Section 4.3, we show how cliquet options are priced using the COS method. In the case of year-on-year options, the total price $V(\mathbf{X}_t,t)$ can be expressed as the sum of individual year-on-year option prices $V_j(\mathbf{X}_t,t)$ over successive yearly intervals. If we consider $V_j(\mathbf{X}_t,t)$ as a cap in Equation 4.30 with $\theta = 1$, conditioned on the initial state $\mathbf{X}_t$,

$$\begin{aligned}V_j(\mathbf{X}_t,t) &= \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^{T_j} r_u du}\left(e^{\ln\Pi_{T_j}-\ln\Pi_{T_j-1}}-1-K\right)^+\right]\\ &= \int_{\mathbb{R}}\int_{\mathbb{R}} V_j(x,T)e^{-y}f_{(x,y)}(x,y)dxdy\\ &= \int_{\mathbb{R}} V_j(x,T)\left(\int_{\mathbb{R}} e^{-y}f_{(x,y)}(x,y)dy\right)dx,\end{aligned}$$

where $x = (\ln(\Pi_{T_j}) - \ln(\Pi_{T_j-1}) - \ln(K+1)), y = \int_t^{T_j} r_s ds$ and $f_{(x,y)}(x,y)$ is the conditional joint probability density function given the initial state $\mathbf{X}_t$. Using the transitional characteristic function ξ_Δ, we have

$$\begin{aligned}e^{-iw\ln(K+1)}\xi_\Delta\left(-iwe_5, iwe_5, 0, -e_2, \mathbf{X}_t, T_{j-1}, T_j\right) &= \int_{\mathbb{R}}\int_{\mathbb{R}} e^{iwx-y}f_{(x,y)}(x,y)dxdy\\ &= \int_{\mathbb{R}} e^{iwx}\left(\int_{\mathbb{R}} e^{-y}f_{(x,y)}(x,y)dy\right)dx\\ &\equiv \int_{\mathbb{R}} e^{iwx}g(x)dx,\end{aligned}$$

with $g(x) = \int_{\mathbb{R}} e^{-y} f_{(x,y)}(x,y)dy$. If we perform a Fourier-cosine series expansion of $g(x)$, the Fourier coefficients are given by

$$F_k = \frac{2}{b-a} \operatorname{Re}\left(\xi_\Delta\left(-\frac{k\pi i}{b-a}e_5, \frac{k\pi i}{b-a}e_5, 0, -e_2, \mathbf{X}_t, T_{j-1}, T_j\right) e^{-k\pi i \frac{\ln(K+1)+a}{b-a}}\right). \tag{4.32}$$

Applying the COS method to each term in Equation 4.30 resembles the option pricing approach described in Section 4.3. However, computing the transitional characteristic function requires solving twice as many ODEs compared to the regular characteristic function. This process consists of two stages. In the first stage, we solve the ODEs with initial condition defined by u_2 to obtain $\phi(T_{j-1}, T_j, u_2, v_2)$ and $\Psi(T_{j-1}, T_j, u_2, v_2)$. In the second stage, these solutions are then used to solve the same set of ODEs but with initial condition defined by $-u_2 + \Psi(T_{j-1}, T_j, u_2, v_2)$. Repeating this process for each term in Equation 4.30 is very time-consuming. Fortunately, we can apply the same method outlined in Section 4.3 to avoid redundant ODE computations for options with different maturities. Since the ODEs remain unchanged, we have $\tilde{\Psi}(T_{j-1}, T_j, u, v) = \Psi(\tau, u_2, v_2)$ and $\phi(T_{j-1}, T_j, u, v) = \phi(\tau, u, v)$. We can obtain $\tilde{\phi}(T_{j-1}, T_j, u, v)$ via Equation 4.29 using a numerical integration method. To further reduce computational costs, we approximate the transitional characteristic function using Chebyshev polynomials. For year-on-year options, we have $T_j - T_{j-1} = 1$ for all j, which implies that $\Psi(T_j - 1, T_j, u_2, v_2) = \Psi(T_{max} - 1, T_{max}, u_2, v_2)$. The algorithm is summarized below.

Algorithm 4.2: Fast option valuation of cliquet options in CP2022 using the Cheby-COS method

Initialization: N_{Cheby}, N_{COS};

1. Solve the ODEs in Equation 4.27 and 4.28 at the roots of the Chebyshev polynomial $T_{N_{Cheby}}(x)$, with the initial condition defined by u_2 for $\tau \in [0, T_{max}]$, where T_{max} is the option with the longest maturity.
2. Use a numerical integration scheme to obtain $\tilde{\phi}(T_j, T_{max}, u_2, v_2)$ as in Equation 4.29 for each maturity T_j at the roots of the Chebyshev polynomial $T_{N_{Cheby}}(x)$.
3. Repeat step 1 and 2, by solving the same set of ODEs but with the initial condition defined by $-u_2 + \Psi(T_{max} - 1, T_{max}, u_2, v_2)$.
4. Use a numerical integration scheme to get $\tilde{\phi}(T_j, T_{max}, u_2, v_2)$ as in Equation 4.29 and determine the transitional characteristic function ξ_Δ in Theorem 4.4.1 at the roots of the Chebyshev polynomial $T_{N_{Cheby}}(x)$ for each subinterval $[T_{j-1}, T_j]$.
5. Approximate ξ_Δ using N_{Cheby} Chebyshev polynomials as in Section 4.1.
6. Use the Chebyshev polynomial approximation of ξ_Δ to calculate option prices using the COS method with N_{COS} cosine-series expansion terms.

4.5 ERROR ANALYSIS

Using a Fourier-based numerical method for pricing options introduces numerical errors, which can be identified in a similar fashion as done by Fang and Oosterlee [2009a]. Since the Chebyshev polynomials and Fourier-cosine series expansion are closely related, we introduce the theory behind the convergence of Fourier-cosine series.

4.5.1 Convergence of Fourier-Cosine Series Expansion Method

In this subsection, we discuss the theory on the convergence of Fourier-cosine series, which is based on Boyd [2001].

Definition 4.5.1 (Algebraic index of convergence). *The algebraic index of convergence n is the largest positive integer such that*

$$\lim_{k\to\infty} |A_k| k^n < \infty, \quad k \gg 1,$$

where A_k are the Fourier-cosines series expansion coefficients of the series. In other words, we say A_k decay asymptotically as

$$A_k \sim \mathcal{O}\left(\frac{1}{k^n}\right), \quad k \gg 1$$

with an algebraic index of convergence equal to n.

Definition 4.5.2 (Exponential index of convergence)**.** *If the algebraic index of convergence n is unbounded, that is, if A_k decreases faster than $1/k^n$ for any finite n, the series is said to have exponential convergence. Alternatively,*

$$A_k \sim \mathcal{O}\left(e^{-\gamma k^r}\right), \quad k \gg 1$$

where γ is a constant defined as the asymptotic rate of convergence, and $r > 0$ is the index of convergence.

If $r < 1$, the convergence is called subgeometric.
If $r = 1$, the convergence is either called supergeometric where

$$A_k \sim \mathcal{O}\left(k^{-n} e^{-k \ln(k)/j}\right),$$

for some $j > 0$, or geometric where

$$A_k \sim \mathcal{O}\left(k^{-n} e^{-\gamma k}\right).$$

Lemma 4.5.3 (Series truncation error of a geometrically converging series)**.** *If a series has geometric convergence, after truncating the series expansion with N terms, the error $E(N-1)$ behaves like*

$$E(N-1) \sim P(N) e^{-(N-1)v}.$$

We call $v > 0$ the asymptotic rate of convergence of the series, which satisfies

$$v = \lim_{N \to \infty} \left(-\log |E(N)| / N\right),$$

and $P(N)$ has slower growth than exponential growth.

Lemma 4.5.4 (Convergence of Fourier-cosine series)**.** *If $g(x)$ is infinitely continuously differentiable on $[a, b]$, then its Fourier-cosine series expansion on $[a, b]$ has geometric convergence. The constant γ as in Definition 4.5.2 is determined by the location in the complex plane of the singularities nearest to the expansion interval. Exponent n is determined by the type and strength of the singularity.*

If a function $g(x)$ is m-times continuously differentiable and if any of its derivatives is discontinuous, its Fourier-cosine series coefficients show algebraic convergence. Integration-by-parts shows that the algebraic index of convergence n, is at least as large as m, as in Lemma 4.1.4.

If $g(x)$ is discontinuous in $c \in [a, b]$, then the Fourier-cosine series converges to $\frac{1}{2}(g(c^-) + g(c^+))$.

Lemma 4.5.5 (Convergence of Chebyshev polynomials)**.** *Comparing series expansions via Chebyshev polynomials and a Fourier-cosine series expansion of a function $f(x)$, the coefficients of $f(x)$ in the Chebyshev series expansion are identical to the coefficients of Fourier-cosine series expansion of $f(cos(\theta))$. As a result, all results from Lemma 4.5.4 carry over. When $f(x)$ is non-periodic on $[a, b]$, Chebyshev-series expansion exhibit convergence **at least as fast as** Fourier-series expansion.*

All statements with their corresponding proofs can be found in Boyd [2001]. These results will be used to quantify the error sources.

4.5.2 Error Sources

There are 4 main error sources present in the COS method, which affects the final option prices. These errors include,

1. The integration range truncation error:

$$\epsilon_1 \equiv V(\mathbf{X}_t, t) - V_1(\mathbf{X}_t, t) = \int_{\mathbb{R} \setminus [a,b]} V(x, T) g(x) dx. \tag{4.33}$$

2. Truncation error of the Fourier-cosine series using N terms:

$$\epsilon_2 \equiv V_1(\mathbf{X}_t,t) - V_2(\mathbf{X}_t,t) = \sum_{k=N}^{\infty} \frac{2}{b-a} \operatorname{Re}\left(\xi\left(\frac{k\pi i}{b-a}e_4, -e_2, \mathbf{X}_t, t\right) e^{-k\pi i \frac{\ln(K)+a}{b-a}}\right) B_k. \quad (4.34)$$

3. Approximating A_k by F_k:

$$\epsilon_3 \equiv V_2(\mathbf{X}_t,t) - V_3(\mathbf{X}_t,t) = \sum_{k=0}^{N-1}{}' \frac{2}{b-a} \operatorname{Re}\left(\int_{\mathbb{R}\setminus[a,b]} e^{k\pi i \frac{y-\ln(K)-a}{b-a}} g(y)dy\right) B_k. \quad (4.35)$$

4. The error arising from numerically constructing the characteristic function $\tilde{\xi}_{\mathrm{n}}$

$$\begin{aligned}\epsilon_4 &\equiv V_3(\mathbf{X}_t,t) - V_4(\mathbf{X}_t,t)\\ &= \sum_{k=0}^{N-1}{}' \frac{2}{b-a} \operatorname{Re}\left(\left[\tilde{\xi}\left(\frac{k\pi i}{b-a}e_4, -e_2, \mathbf{X}_t, t\right) - \tilde{\xi}_{\mathrm{n}}\left(\frac{k\pi i}{b-a}e_4, -e_2, \mathbf{X}_t, t\right)\right] e^{-k\pi i \frac{\ln(K)+a}{b-a}}\right) B_k.\end{aligned} \quad (4.36)$$

The first three error sources have been identified by Fang and Oosterlee [2009a]. We summarize the results below.

Errors ϵ_1 and ϵ_3 are integration range truncation errors, which can be minimized if the truncated integration range $[a,b]$ is large enough. The following lemma is based on the works of Fang and Oosterlee [2009a].

Lemma 4.5.6. *Assuming the joint probability density function $f_{(x,y)}(x,y)$ with subsequent density function $g(y)$ are real, then for some constant $c > 0$,*

$$|\epsilon_3| \leq c \int_{\mathbb{R}\setminus[a,b]} g(y)dy. \quad (4.37)$$

Proof. Since $g(y)$ is real, positive and $|\cos(x)| \leq 1$,

$$\begin{aligned}|\epsilon_3| &= \left|\frac{2}{b-a}\sum_{k=0}^{N-1}{}'\left[\int_{\mathbb{R}\setminus[a,b]} \cos\left(k\pi\frac{y-\ln(K)-a}{b-a}\right) g(y)dy\right] B_k\right|\\ &\overset{(*)}{\leq} \frac{2}{b-a}\sum_{k=0}^{N-1}{}'\left|\int_{\mathbb{R}\setminus[a,b]} \cos\left(k\pi\frac{y-\ln(K)-a}{b-a}\right) g(y)dy\right| |B_k|\\ &\leq \frac{2}{b-a}\sum_{k=0}^{N-1}{}'\left[\int_{\mathbb{R}\setminus[a,b]} \left|\cos\left(k\pi\frac{y-\ln(K)-a}{b-a}\right)\right| |g(y)|\, dy\right] |B_k|\\ &\leq \frac{2}{b-a}\int_{\mathbb{R}\setminus[a,b]} g(y)dy \sum_{k=0}^{N-1}{}' |B_k|,\end{aligned}$$

where in $(*)$, the triangle inequality is used. The last part is a truncated summation of B_k, which in worst-case scenario exhibits algebraic convergence and can therefore be bounded. The result then follows.

For the European options, using that

$$\sum_{k=1}^{N}\frac{1}{k^2} < \sum_{k=1}^{\infty}\frac{1}{k^2} = \frac{\pi^2}{6}, \qquad \sum_{k=1}^{N}\frac{1}{k} < 1+\ln(N),$$

we get

$$\begin{aligned}\sum_{k=0}^{N-1}{}'|\chi_k(c,d)| &< \frac{1}{2}(e^d+e^c) + (e^d+e^c)\sum_{k=1}^{N-1}\left|\left(\frac{b-a}{k\pi}\right)^2 + \frac{b-a}{k\pi}\right|\\ &< \frac{1}{2}(e^d+e^c) + \frac{1}{\pi}(e^d+e^c)(b-a)\left(\frac{1}{6}(b-a)\pi + 1 + \ln(N-1)\right),\end{aligned}$$

and

$$\sum_{k=0}^{N-1}{}'|\psi_k(c,d)| \leq \frac{1}{2}(d-c) + 2\sum_{k=1}^{N-1}\frac{b-a}{k\pi} < \frac{1}{2}(d-c) + \frac{2}{\pi}(b-a)(1+\ln(N-1)).$$

Therefore, for the European options we have

$$|\epsilon_3| \leq \frac{2K}{b-a}\left(\sum_{k=0}^{N-1}{}'|\chi_k(a,b)| + |\psi_k(a,b)|\right)\int_{\mathbb{R}\setminus[a,b]} g(y)dy \equiv c\int_{\mathbb{R}\setminus[a,b]} g(y)dy.$$

□

The series truncation error ϵ_2 can be bound using the theory from Section 4.5.1. If $g(x)$ is infinitely continuously differentiable with non-zero derivatives,

$$|\epsilon_2| \leq P(N)e^{-(N-1)v},$$

which follows from Lemma 4.5.3. Otherwise, if the density has discontinuous derivatives or singularities,

$$|\epsilon_2| \leq \frac{C}{N^m},$$

where m is the lowest order derivative with a singularity or discontinuity, see Lemma 4.5.4 and Lemma 4.1.3.

To analyse ϵ_4, we need to quantify the numerical error arising from numerically constructing the characteristic function. We analyse each error source.

Step I

The ODEs are solved numerically with solutions $\phi_{uv}(t,T)$ and $\Psi_{uv}(t,T)$, whose accuracy can be controlled by a predetermined tolerance level TOL. The difference between the true characteristic function $\xi(u,v,\mathbf{X}_t,t)$ and approximation $\tilde{\xi}(u,v,\mathbf{X}_t,t)$ is

$$\begin{aligned}
\epsilon_I &\equiv \tilde{\xi}(u,v,\mathbf{X}_t,t) - \xi(u,v,\mathbf{X}_t,t)\\
&= e^{\phi_{uv}(t,T)+\Psi_{uv}(t,T)'\mathbf{X}_t} - e^{\phi_{uv}(t,T)+\epsilon+(\Psi'_{uv}(t,T)+\boldsymbol{\epsilon}')\mathbf{X}_t}\\
&= (1-e^{\epsilon+\boldsymbol{\epsilon}'\mathbf{X}_t})\xi(u,v,\mathbf{X}_t,t)\\
&\overset{(*)}{\approx} (\epsilon+\boldsymbol{\epsilon}'\mathbf{X}_t)\xi(u,v,\mathbf{X}_t,t),
\end{aligned}$$

where in $(*)$ we assume the errors are small enough to justify a first-order Taylor expansion.

Step II

The numerical integration scheme of a function with error ϵ can be bound using the triangle inequality. Denote I_N by the numerical integration approximation of $f(x)$ and $\tilde{I}_N$ by the numerical integration approximation of $f(x)+\epsilon$. Then

$$|\epsilon_{II}| \equiv \left|\tilde{I}_N - \int_a^b f(x)dx\right| \leq \left|I_N - \int_a^b f(x)dx\right| + \left|\tilde{I}_N - I_N\right|.$$

Case 1 : trapezoidal rule.

The first term is well-known for the trapezoidal rule, and for some $c \in (a,b)$ equals $\frac{(b-a)^3}{12N^2}|f''(c)|$. The second term equals

$$\left|\tilde{I}_N - I_N\right| \leq \frac{1}{N}\sum_{i=1}^{N}|\epsilon_i| \leq \|\epsilon\|_\infty.$$

Case 2 : Clenshaw-Curtis quadrature.

The first term is also well-known, which converges geometrically for infinitely differentiable functions with no singularities, and algebraically otherwise. For the second term we use Lemma 4.1.5 and without loss of generality we assume N is even.

$$\begin{aligned}
|\tilde{I}_N - I_N| &\leq \frac{b-a}{2}\sum_{k=0}^{N}|\epsilon_k w_k| \\
&\leq \frac{b-a}{2}\|\epsilon\|_\infty \sum_{k=0}^{N}|w_k| \\
&= \frac{b-a}{2}\|\epsilon\|_\infty \sum_{k=0}^{N}\left|\frac{c_k}{N}\left(1-\sum_{j=1}^{\lceil N/2\rceil}\frac{2-\mathbb{1}_{\{j=N/2\}}}{4j^2-1}\cos(2jx_k)\right)\right| \\
&\leq \frac{b-a}{N}\|\epsilon\|_\infty \sum_{k=0}^{N}\left|1+\sum_{j=1}^{N/2}\frac{2}{4j^2-1}\right| \\
&= \frac{b-a}{N}\|\epsilon\|_\infty \sum_{k=0}^{N}\left|1+\sum_{j=1}^{N/2}\left(\frac{1}{2j-1}-\frac{1}{2j+1}\right)\right| \\
&\overset{(*)}{=} \frac{b-a}{N}\|\epsilon\|_\infty \sum_{k=0}^{N}\left|2-\frac{1}{N+1}\right| \\
&= (b-a)\|\epsilon\|_\infty\left(2+\frac{1}{N}\right),
\end{aligned}$$

where in $(*)$, we use the following telescoping series:

$$\sum_{j=1}^{N/2}\left(\frac{1}{2j-1}-\frac{1}{2j+1}\right) = 1-\frac{1}{N+1}.$$

Step III

Similar to step I, replacing ϕ by $\tilde{\phi}$ yields error:

$$\epsilon_{III} = (1-e^{\epsilon})\xi(u,v,\mathbf{X}_t,t) \approx \epsilon\xi(u,v,\mathbf{X}_t,t).$$

Step IV

Denote f by the true function, $\tilde{f}_\epsilon$ by the Chebyshev interpolated function using measurements with errors ϵ_i and $\tilde{f}$ by the exact Chebyshev interpolated function with N terms. Then, by the triangle inequality

$$|\tilde{f}_\epsilon(x) - f(x)| \leq |\tilde{f}(x) - f(x)| + |\tilde{f}_\epsilon(x) - \tilde{f}(x)|.$$

The first error is well-known, and converges geometrically for infinitely differentiable functions with no singularities, and algebraically otherwise. For the second term, using Lemma 4.1.2, the error comes from the Chebyshev coefficients:

$$|\tilde{a}_n - a_n| = \frac{2}{N}\sum_{k=1}^{N}\epsilon_k T_n(x_k) \leq \frac{2}{N}\|\epsilon\|_\infty \sum_{k=1}^{N}|T_n(x_k)| \leq 2\|\epsilon\|_\infty.$$

As a result, for the second term we have

$$|\tilde{f}_\epsilon(x) - \tilde{f}(x)| \leq \sum_{n=0}^{N-1}{}'2\|\epsilon\|_\infty |T_n(x)| \leq 2N\|\epsilon\|_\infty.$$

The final numerical error in the characteristic function is determined by iteratively substituting each error into the corresponding expressions from steps I to IV. Denote the final error in the numerical approximation of the characteristic function by $\hat{\epsilon}$. We then have the following result.

Lemma 4.5.7. *Suppose $\hat{\epsilon}$ is the numerical error arising from approximating the characteristic function which can be bounded and B_k is defined as in Equation 4.19. Then, for some constant $c > 0$,*

$$|\epsilon_4| \leq c\|\hat{\epsilon}\|_\infty.$$

Proof.

$$\begin{aligned}
|\epsilon_4| &= \frac{2}{b-a}\left|\sum_{k=0}^{N-1}{}' \operatorname{Re}\left(\hat{\epsilon}e^{-k\pi i\frac{\ln(K)+a}{b-a}}\right) B_k\right| \\
&\overset{(*)}{\leq} \frac{2}{b-a}\sum_{k=0}^{N-1}{}'\left|\operatorname{Re}\left(\hat{\epsilon}e^{-k\pi i\frac{\ln(K)+a}{b-a}}\right)\right| |B_k| \\
&\leq \frac{2}{b-a}\sum_{k=0}^{N-1}{}' |\hat{\epsilon}|\,|B_k|\,,
\end{aligned}$$

where in $(*)$, the triangle inequality is used. Since $\hat{\epsilon}$ is bounded, we are left with a truncated summation of B_k, which in worst-case scenario exhibits algebraic convergence and can be bounded, see for example Lemma 4.5.6. □

The numerical computation of the transitional characteristic function introduces an additional source of numerical error compared to the regular characteristic function. Obtaining the transitional characteristic function involves two stages: first, solving the ODEs with initial conditions defined by u_2 to obtain $\phi(T_{j-1}, T_j, u_2, v_2)$ and $\Psi(T_{j-1}, T_j, u_2, v_2)$, and then using these solutions to solve the ODEs with the new initial condition defined by $-u_2 + \Psi(T_{j-1}, T_j, u_2, v_2)$. This second stage introduces a potential source of numerical error due to the accuracy of the numerical solution $\Psi(T_{j-1}, T_j, u_2, v_2)$ obtained in the first stage. Any errors in this solution propagate into the second stage, affecting the accuracy of the final results. While the precise impact of this error on the overall numerical solution is difficult to quantify, it can be mitigated by using an accurate ODE solver or enforcing strict tolerance levels within the solver.

5 RESULTS

In this chapter, we present the numerical results of option prices under CP2022, computed using the Monte Carlo method, as well as the COS method and Cheby-COS methods. We will first consider European call and put options on the stock index and CPI, followed by year-on-year inflation-indexed caps and floors. In particular, we verify our implementation by first comparing our results with Monte Carlo simulations. We then conduct error analysis experiments by varying the number of cosine terms N_{COS} in the COS method, the number of Chebyshev polynomials N_{Cheby} used to approximate the characteristic function, and the number of quadrature points n used per year for each integration method. Finally, we compare option prices computed using different configurations and compare our option pricing method with the Monte Carlo approach when calibrating CP2022.

All results were obtained using Intel(R) Xeon(R) Platinum 8370C CPU @ 2.8GHz with 8 Gb of DDR4-3200 MHz memory. The code is written in Matlab R2024a. Additionally, we use the CP2022 parameters of Q2 2024 that are made publicly available by DNB each quarter (De Nederlandsche Bank [2022]) and can also be found in Appendix A.3.

5.1 EUROPEAN CALLS AND PUTS ON STOCK INDEX AND CPI

First, to verify our implementation, we compare the regular COS method with the Monte Carlo method used in production for pricing European options on the stock index and CPI. The Monte Carlo method employs monthly time steps and 100,000 simulations—this configuration is referred to as MC I. However, we cannot take these as reference values, since the discretization error when generating simulation paths using the Euler–Maruyama scheme dominates the total error in Monte Carlo, while the COS method is expected to exhibit exponential convergence. Hence, the Monte Carlo reference values uses 1,000,000 simulations with 10,000 time steps per year and enable antithetic variates as a variance reduction technique (MC II). Due to hardware and time constraints, for MC II we only simulate up to five years into the future instead of twenty years. To enable a more accurate comparison for maturities exceeding five years, we reran MC I with 120 time steps per year and 100,000 simulations, and to this configuration as MC I*.

For the remainder of this chapter, for the COS method, we shall use the truncated integration interval $[a, b]$ as defined in Equation 4.23 with $L = 10$ and approximate Equation 4.24 using central differences of second order. Assume that a truncated interval $[a_i, b_i]$ is used for year-on-year options on $[T_i, T_{i+1}]$. Additionally, when pricing European options on S_t and Π_t simultaneously with n maturities, we define the common truncated integration interval $[a, b]$ per option type as

$$\begin{aligned} a &= \min\{a_i : i = 1, \dots n\}, \\ b &= \max\{b_i : i = 1, \dots n\}. \end{aligned} \tag{5.1}$$

Furthermore, based on the strikes used during calibration for each option, we use $K = \{0.8, 0.81, \dots, 1.3\}$ for European call options on S_t, $K = \{0, 0.001, \dots, 0.05\}$ for zero-coupon inflation-indexed caps, floors, and year-on-year inflation-indexed caps, and $K = \{-0.025, -0.024, \dots, 0.025\}$ for year-on-year inflation-indexed floors. The ODEs are solved using the Dormand-Prince Runge-Kutta (4,5) method (Dormand and Prince [1980]) and we fix the relative error tolerance to 10^{-6}.

In Figure 5.1, we illustrate the discretization error in the results of MC I by comparing the COS method with identical parameters, first against MC I and then against MC II using a 95% confidence interval for a zero-coupon inflation-indexed cap with a maturity of 1 year.

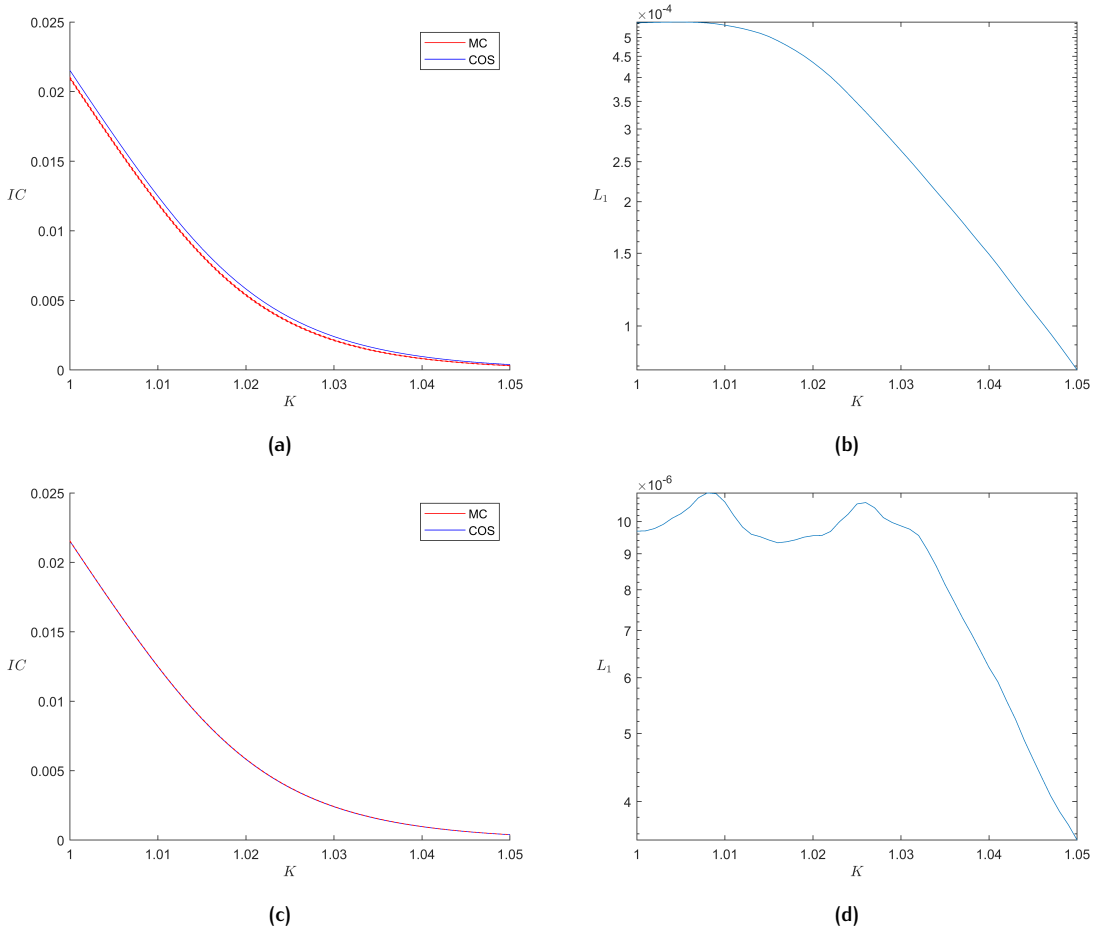

Figure 5.1: Zero-coupon inflation-indexed cap prices with a maturity of 1 year using MC I and the COS method (a) with absolute difference (b), and using MC II and the COS method (c) with absolute difference (d). Monte Carlo results plotted with a 95% confidence interval.

Note that the discretization error dominates the numerical error in MC I, since prices computed by the COS method first do not fall inside the confidence interval, but do fall inside the interval in MC II. Hence, we shall use MC II as Monte Carlo reference values. Next, to verify the implementation of the COS method, we compare the COS method with MC II when pricing European calls (C), zero-coupon inflation-indexed caps (IC) and floors (IF), by varying the cosine-series expansion terms N_{COS}. We take the values obtained from MC II as reference and define the maximum absolute difference and relative difference over the strikes as:

$$\begin{aligned} L_\infty &= \max_K \left| V_{MC}(T,K) - V_{COS}(T,K) \right|, \\ L_\infty^{rel} &= \max_K \left| \frac{V_{MC}(T,K) - V_{COS}(T,K)}{V_{MC}(T,K)} \right|. \end{aligned} \tag{5.2}$$

N_{COS}	C								
	$T=1$			$T=2$			$T=5$		
	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)
16	2.0e-2	1.5e1	0.06	1.5e-1	3.3e0	0.07	1.7e0	9.3e0	0.07
32	1.7e-3	3.7e-1	0.09	1.1e-2	5.5e-1	0.10	7.7e-2	5.9e-1	0.15
64	9.2e-5	2.5e-2	0.20	2.9e-4	5.8e-3	0.23	4.2e-4	4.1e-3	0.30
128	7.6e-5	9.5e-3	0.47	1.8e-4	4.3e-3	0.53	2.6e-4	2.0e-3	0.75
256	7.6e-5	9.5e-3	1.12	1.8e-4	4.3e-3	1.25	2.6e-4	2.0e-3	1.71
512	7.6e-5	9.5e-3	2.80	1.8e-4	4.3e-3	3.22	2.6e-4	2.0e-3	4.51

Table 5.1: Maximum absolute and relative differences in European call prices on S_t computed using MC II and the COS method for different number of cosine-series expansion terms N_{COS}.

N_{COS}	IC								
	$T=1$			$T=2$			$T=5$		
	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)
16	1.2e-2	2.2e1	0.06	7.8e-3	2.7e0	0.12	1.3e-2	4.6e0	0.12
32	6.0e-3	5.2e0	0.10	2.3e-3	2.3e0	0.19	2.0e-3	1.6e0	0.15
64	2.1e-3	1.8e0	0.16	2.5e-4	3.1e-1	0.33	2.9e-5	2.4e-2	0.32
128	3.8e-4	3.5e-1	0.33	2.5e-5	7.8e-3	0.66	6.9e-6	3.5e-3	0.73
256	2.4e-5	3.1e-2	0.72	2.1e-5	6.4e-3	1.27	6.9e-6	3.5e-3	1.61
512	1.1e-5	9.0e-3	1.83	2.1e-5	6.4e-3	2.84	6.9e-6	3.5e-3	4.10

Table 5.2: Maximum absolute and relative differences in zero-coupon inflation-indexed cap prices computed using MC II and the COS method for different number of cosine-series expansion terms N_{COS}.

N_{COS}	IF								
	$T=1$			$T=2$			$T=5$		
	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)
16	1.7e-2	2.3e2	0.08	7.3e-3	2.1e1	0.07	7.0e-3	2.3e1	0.10
32	6.4e-3	4.9e1	0.11	1.8e-3	1.6e1	0.13	1.2e-3	6.7e0	0.15
64	1.9e-3	7.1e0	0.16	2.2e-4	2.8e0	0.22	3.4e-5	1.4e-1	0.38
128	3.3e-4	5.4e0	0.34	2.3e-5	3.4e-2	0.45	2.8e-5	2.0e-3	0.72
256	1.8e-5	1.2e-1	0.94	2.1e-5	6.5e-3	1.04	2.8e-5	2.0e-3	1.57
512	7.7e-6	2.8e-3	1.86	2.1e-5	6.5e-3	2.64	2.8e-5	2.0e-3	4.01

Table 5.3: Maximum absolute and relative differences in zero-coupon inflation-indexed floor prices computed using MC II and the COS method for different number of cosine-series expansion terms N_{COS}.

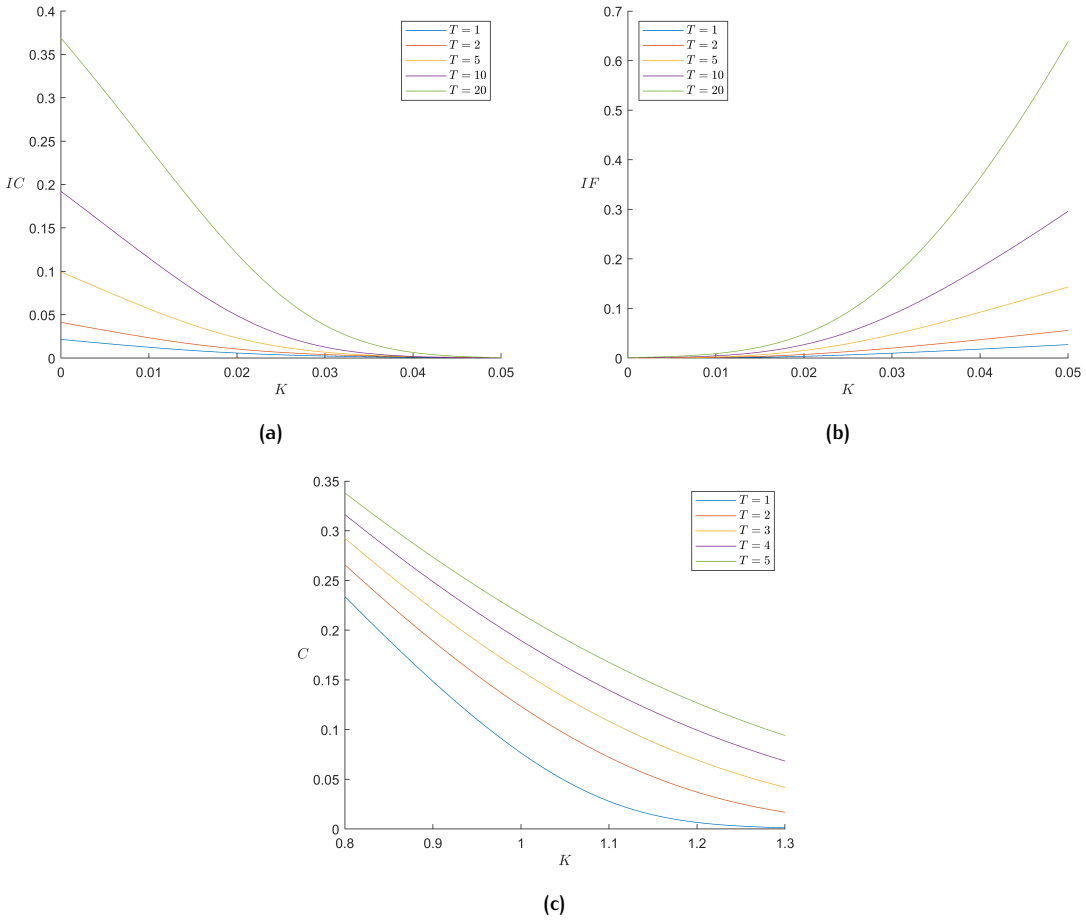

Figure 5.2: Illustration of inflation-indexed cap (a) and floor (b) prices, and European call option prices on S_t (c), for various strike prices K using the COS method.

The COS method matches the prices computed by Monte Carlo, with the difference between the two methods being within the Monte Carlo sampling error. The total CPU time of the COS method is dominated by the time required to numerically solve the system of ODEs to obtain the characteristic function. Despite this, the COS method is able to accurately price the options within a few seconds, whereas MC II took over 6 hours, showing that the COS method dominates the Monte Carlo method in both computational speed and accuracy.

The COS method is expected to achieve exponential convergence. To verify this numerically, we compute reference option prices using the COS method with a conservatively large number of cosine-series expansion terms, say $N_{COS} = 2048$. We define the maximum absolute and relative difference over all strikes by L_∞^{COS} and $L_\infty^{rel,COS}$ respectively:

$$\begin{aligned} L_\infty^{COS} &= \max_K \left| V_{COS(N=2048)}(T,K) - V_{COS}(T,K) \right|, \\ L_\infty^{rel,COS} &= \max_K \left| \frac{V_{COS(N=2048)}(T,K) - V_{COS}(T,K)}{V_{COS(N=2048)}(T,K)} \right|. \end{aligned} \tag{5.3}$$

N_{COS}	C					
	$T=1$		$T=2$		$T=5$	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
16	2.0e-2	1.5e1	1.5e-1	3.4e0	1.7e0	9.4e0
32	1.7e-3	3.8e-1	1.1e-2	5.5e-1	7.7e-2	5.9e-1
64	3.5e-5	1.6e-2	1.5e-4	9.1e-3	2.1e-4	2.1e-3
128	4.4e-8	3.0e-5	5.4e-8	3.1e-6	2.8e-9	2.5e-8
256	1.3e-13	9.5e-11	2.0e-14	1.1e-12	2.8e-16	1.3e-15
512	1.7e-16	1.1e-15	1.4e-16	1.1e-15	0	0

Table 5.4: Maximum absolute difference between the COS method varying amount of cosine-series expansion terms N_{COS} and the COS method with $N_{COS} = 2048$ as reference values for European calls on stock index.

N_{COS}	IC									
	$T=1$		$T=2$		$T=5$		$T=10$		$T=20$	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
16	1.2e-2	2.2e1	7.8e-3	2.7e0	1.3e-2	4.6e0	2.0e-2	1.0e2	1.0e-2	1.6e1
32	6.0e-3	5.2e0	2.3e-3	2.3e0	2.0e-3	1.6e0	6.4e-4	2.4e0	1.1e-5	1.9e-2
64	2.1e-3	1.8e0	2.5e-4	3.1e-1	2.4e-5	2.2e-2	5.9e-7	1.7e-3	2.7e-14	3.2e-11
128	3.9e-4	3.4e-1	4.4e-6	2.9e-3	4.8e-9	3.9e-6	9.7e-16	5.0e-12	0	0
256	1.5e-5	2.2e-2	8.9e-10	8.5e-7	5.6e-17	1.2e-15	2.8e-17	4.2e-16	0	0
512	5.2e-8	8.4e-5	6.9e-17	2.0e-15	1.7e-18	4.0e-16	0	0	0	0
1024	8.8e-14	1.3e-10	1.1e-19	1.7e-16	0	0	0	0	0	0

Table 5.5: Maximum absolute difference between the COS method varying amount of cosine-series expansion terms N_{COS} and the COS method with $N_{COS} = 2048$ as reference values for inflation-indexed cap prices.

N_{COS}	IF									
	$T=1$		$T=2$		$T=5$		$T=10$		$T=20$	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
16	1.7e-2	2.3e2	7.3e-3	2.1e1	7.0e-3	2.3e1	5.0e-3	8.2e0	6.4e-4	1.1e-1
32	6.4e-3	4.9e1	1.8e-3	1.6e1	1.1e-3	6.7e0	1.9e-4	2.2e-1	3.8e-7	6.6e-5
64	1.9e-3	7.1e0	2.2e-4	2.8e0	2.6e-5	1.4e-1	1.4e-7	1.6e-4	1.1e-16	2.3e-14
128	3.3e-4	5.4e0	3.8e-6	3.8e-2	5.0e-9	3.2e-5	1.4e-16	2.2e-13	0	0
256	1.8e-5	1.2e-1	6.7e-10	9.3e-6	1.2e-16	2.6e-15	1.1e-16	7.6e-16	0	0
512	5.6e-8	6.1e-4	4.9e-17	1.8e-15	8.7e-19	4.1e-16	0	0	0	0
1024	7.3e-14	1.1e-9	1.1e-19	4.3e-16	0	0	0	0	0	0

Table 5.6: Maximum absolute difference between the COS method varying amount of cosine-series expansion terms N_{COS} and the COS method with $N_{COS} = 2048$ as reference values for inflation-indexed floor prices.

We observe exponential convergence, as theoretically predicted. Next, we investigate the impact of the number of quadrature points used per integration scheme employed when integrating $f(t)$ and $f^R(t)$ as in Equation 4.29 in the final option prices. Note that the interval of integration is partitioned at the discontinuities and each sub-interval with $h = 1/12$ is integrated separately using n quadrature points. We consider inflation-indexed cap prices by using reference values obtained with $N_{COS} = 2048$ and denote the maximum absolute and relative difference over all strikes by L_∞^{COS} and $L_\infty^{rel,COS}$ respectively.

n	IC											
	$T=1$				$T=2$				$T=5$			
	tpz rule		CC rule		tpz rule		CC rule		tpz rule		CC rule	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
1	4.8e-9	6.9e-7	4.8e-9	6.9e-7	1.9e-9	4.7e-8	1.9e-9	4.7e-8	1.8e-9	3.6e-8	1.8e-9	3.6e-8
2	1.2e-9	1.7e-7	4.5e-15	3.3e-12	4.8e-10	1.2e-8	7.2e-15	1.1e-11	4.5e-10	9.1e-9	5.8e-15	1.6e-11
4	3.0e-10	4.3e-8	1.1e-16	1.9e-13	1.2e-10	2.9e-9	3.4e-15	1.6e-12	1.1e-10	2.3e-9	1.4e-16	3.3e-13

Table 5.7: Maximum absolute difference between the COS method for varying amount of quadrature points n per sub-interval using trapezoidal integration rule and Clenshaw-Curtis quadrature, and the COS method with $N_{COS} = 2048$ as reference values for inflation-indexed cap prices.

Within the relative error tolerance level of the ODE solver, which we set to 10^{-12} in this case, we observe the expected exponential convergence for Clenshaw-Curtis quadrature and second-order convergence for the trapezoidal rule in Table 5.7. Notably, using just one quadrature point per sub-interval already yields accurate results.

Each term in the cosine-series expansion requires numerically solving the system of ODEs. Since the majority of the computation time in the COS method is spent on solving these ODEs, we aim to reduce this cost by interpolating the characteristic function using Chebyshev polynomials and sampling from the interpolated function instead. Assume that in the Cheby-COS method, we fix $N_{COS} = 512$. We investigate the impact of the number of Chebyshev polynomials N_{Cheby} used to interpolate the characteristic function on the accuracy of the option prices. We compare the Cheby-COS method with the COS method with $N_{COS} = 512$, and define the maximum absolute difference and relative difference over the strikes as:

$$\begin{aligned} \hat{L}_\infty &= \max_K \left| V_{COS}(T,K) - V_{Cheby-COS}(T,K) \right|, \\ \hat{L}_\infty^{rel} &= \max_K \left| \frac{V_{COS}(T,K) - V_{Cheby-COS}(T,K)}{V_{COS}(T,K)} \right|. \end{aligned} \tag{5.4}$$

N_{Cheby}	C					
	$T=1$		$T=2$		$T=5$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	4.1e-1	3.3e2	1.4e0	8.5e1	8.9e0	9.5e1
16	4.6e-2	3.7e1	1.8e-1	1.0e1	5.1e-1	5.4e0
32	1.8e-3	1.4e0	7.3e-3	4.3e-1	6.6e-2	7.1e-1
64	2.8e-6	2.3e-3	1.1e-5	6.8e-4	7.3e-5	7.8e-4
128	9.4e-11	5.5e-8	3.6e-9	2.2e-7	1.2e-7	1.2e-6

Table 5.8: Maximum absolute and relative differences in European call option prices on S_t, which are computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby}.

N_{Cheby}	IC									
	$T = 1$		$T = 2$		$T = 5$		$T = 10$		$T = 20$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	4.5e-2	1.2e2	1.7e-1	2.7e2	6.6e-1	1.7e3	1.3e0	6.7e3	2.0e0	3.5e3
16	2.5e-4	6.5e-1	1.3e-2	2.2e1	9.4e-2	2.5e2	4.4e-1	2.3e3	2.0e0	3.5e3
32	1.8e-8	4.6e-5	2.1e-5	3.6e-2	4.2e-4	1.1e0	7.3e-2	3.8e2	9.4e-1	1.6e3
64	9.0e-11	1.7e-7	4.6e-10	7.9e-7	7.3e-9	1.8e-5	6.7e-7	2.5e-4	2.1e-3	3.2e0
128	4.3e-11	1.1e-7	1.9e-10	3.2e-7	2.5e-9	5.4e-6	1.7e-8	1.9e-5	8.5e-8	1.5e-4

Table 5.9: Maximum absolute and relative differences in inflation-indexed cap prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby}.

N_{Cheby}	IF									
	$T = 1$		$T = 2$		$T = 5$		$T = 10$		$T = 20$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	1.4e-2	2.7e2	4.1e-3	3.0e1	1.8e-1	1.2e3	4.5e-1	5.9e2	1.2e0	2.6e2
16	1.0e-4	1.9e0	5.7e-3	7.8e1	1.1e-2	6.5e0	1.1e-1	2.3e2	1.9e-1	6.5e1
32	3.7e-10	7.3e-6	1.1e-5	1.8e-1	2.3e-4	1.6e0	3.0e-3	7.2e0	1.8e-2	1.6e1
64	1.0e-10	1.8e-6	9.5e-11	1.4e-6	1.3e-9	9.1e-6	3.6e-7	3.7e-4	2.9e-4	1.4e-1
128	3.5e-12	4.0e-8	2.6e-11	2.1e-7	6.0e-10	2.7e-6	2.8e-9	5.4e-6	9.8e-10	6.2e-7

Table 5.10: Maximum absolute and relative differences in inflation-indexed floor prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby}.

In Section 4.5, we demonstrated that Chebyshev polynomials exhibit exponential convergence when approximating functions. Moreover, we showed that when each error source in the COS method is minimized, by using conservative values for $[a, b]$ and N_{COS}, Chebyshev polynomials also exhibit exponential convergence to option prices up to the numerical error introduced from the ODE solver. This behaviour can also be seen in Table 5.8, 5.9, and 5.10, which shows that Chebyshev polynomials are a reliable method for enhancing computational speed when the COS method is used with a characteristic function that requires numerically solving a large system of ODEs.

Finally, we compare option prices computed by the COS method and the Cheby-COS method across multiple maturities simultaneously, using a numerical integration scheme for Equation 4.29. Specifically, we fix $N_{COS} = 512$ and $N_{Cheby} = 128$, and compare the trapezoidal rule (tpz) with 12 steps per year and Clenshaw-Curtis quadrature (CC) with 48 quadrature points per year in terms of speed and accuracy. Option prices calculated using Monte Carlo are reported with a 95% confidence interval.

T	$C(K=0.8)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.2339 ± 0.0008	0.2340 ± 0.0003	0.2339	0.2339	0.2339	0.2339
2	0.2658 ± 0.0012	0.2660 ± 0.0004	0.2659	0.2659	0.2659	0.2659
5	0.3381 ± 0.0021	0.3383 ± 0.0007	0.3381	0.3381	0.3381	0.3381
T	$C(K=1.2)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.0066 ± 0.0002	0.0066 ± 0.0001	0.0066	0.0066	0.0066	0.0066
2	0.0372 ± 0.0005	0.0372 ± 0.0002	0.0371	0.0371	0.0371	0.0371
5	0.1268 ± 0.0014	0.1271 ± 0.0005	0.1269	0.1269	0.1269	0.1269
Total CPU time (s)	130	$2.2 \cdot 10^4$	4.0	5.2	1.1	1.3

Table 5.11: European call option prices for different strikes and maturities obtained by Monte Carlo with 95% confidence interval, COS method, Cheby-COS method, employing the trapezoidal rule or Clenshaw-Curtis quadrature as a numerical integration scheme.

T	MC I	
	$C(K=0.8)$	$C(K=1.3)$
1	0.2353 ± 0.0008	0.0063 ± 0.0002
2	0.2681 ± 0.0012	0.0372 ± 0.0005
5	0.3400 ± 0.0021	0.1271 ± 0.0014

Table 5.12: European call option prices obtained from MC I with 95% confidence interval. CPU time is 11 s.

T	$IC(K = 0.02)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
5	0.0234 ± 0.0002	0.0233 ± 0.0001	0.0233	0.0233	0.0233	0.0233
10	0.0491 ± 0.0004	–	0.0491	0.0491	0.0491	0.0491
20	0.1193 ± 0.0013	–	0.1203	0.1203	0.1203	0.1203
T	$IC(K = 0.04)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
5	0.00165 ± 0.00006	0.00165 ± 0.00002	0.00164	0.00164	0.00164	0.00164
10	0.00191 ± 0.00008	–	0.00193	0.00193	0.00193	0.00193
20	0.00636 ± 0.00031	–	0.00652	0.00651	0.00651	0.00651
Total CPU time (s)	130	$2.2 \cdot 10^4$	5.8	11.0	1.5	2.9

Table 5.13: Inflation-indexed cap prices for different strikes and maturities obtained by Monte Carlo with 95% confidence interval, COS method, Cheby-COS method, employing the trapezoidal rule or Clenshaw-Curtis quadrature as a numerical integration scheme.

T	$IF(K = 0.01)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
5	0.00227 ± 0.00005	0.00228 ± 0.00002	0.00228	0.00228	0.00228	0.00228
10	0.00455 ± 0.00010	–	0.00457	0.00457	0.00457	0.00457
20	0.00912 ± 0.00017	–	0.00906	0.00906	0.00906	0.00906
Total CPU time (s)	130	$2.2 \cdot 10^4$	5.8	11.0	1.5	2.9

Table 5.14: Inflation-indexed floor prices for different maturities obtained by Monte Carlo with 95% confidence interval, COS method, Cheby-COS method, employing the trapezoidal rule or Clenshaw-Curtis quadrature as a numerical integration scheme.

T	MC I		
	$IC(K = 0.02)$	$IC(K = 0.04)$	$IF(K = 0.01)$
5	0.0230 ± 0.0002	0.00158 ± 0.00006	0.00218 ± 0.00005
10	0.0489 ± 0.0004	0.00189 ± 0.00009	0.00436 ± 0.00010
20	0.1201 ± 0.0013	0.00638 ± 0.00031	0.00638 ± 0.00031

Table 5.15: Inflation-indexed cap and floor prices prices obtained from MC I with 95% confidence interval. CPU time is 11 s.

Since the CPU time of the COS method scales linearly with N_{COS}, the Cheby-COS method is approximately four times faster for both numerical integration schemes, as it requires solving four times fewer ODEs compared to the COS method. Additionally, the numerical integration scheme employed had little impact in the final option prices, though that is also determined by the smoothness of the numerical solution ϕ and Ψ, which change during calibration. All variations of the COS and Cheby-COS method fall within the confidence intervals of the Monte Carlo estimates, whereas the option prices computed using MC I are dominated by the discretization error. Again, we see that the COS method and Cheby-COS method dominate the Monte Carlo method in both accuracy and computational speed.

5.2 YEAR-ON-YEAR INFLATION-INDEXED CAPS AND FLOORS

In this section, we compare year-on-year inflation-indexed cap (YIC) and floor (YIF) prices using different methods. First, we verify the implementation of the COS method by comparing results with Monte Carlo reference values obtained from MC II.

N_{COS}	YIC								
	$T=1$			$T=2$			$T=5$		
	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)
16	1.2e-2	2.2e1	0.12	2.1e-2	9.0e0	0.12	5.3e-2	7.2e0	0.28
32	6.0e-3	5.2e0	0.21	9.2e-3	1.7e0	0.24	2.8e-2	2.4e0	0.45
64	2.1e-3	1.8e0	0.33	2.6e-3	8.6e-1	0.48	8.9e-3	3.3e-1	0.94
128	3.8e-4	3.5e-1	0.56	3.5e-4	8.0e-2	0.90	1.5e-3	1.3e-1	1.92
256	2.4e-5	3.1e-2	1.18	2.9e-5	6.5e-3	2.05	5.7e-5	4.5e-3	3.92
512	1.1e-5	9.0e-3	2.73	2.1e-5	4.5e-3	4.68	9.8e-6	1.4e-3	8.84

Table 5.16: Maximum absolute and relative differences in year-on-year inflation-indexed cap prices computed using MC II and the COS method for different number of cosine-series expansion terms N_{COS}.

N_{COS}	YIF								
	$T=1$			$T=2$			$T=5$		
	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)	L_∞	L_∞^{rel}	CPU time (s)
16	1.7e-2	4.0e5	0.08	2.8e-2	3.1e3	0.12	8.9e-2	4.0e2	0.28
32	6.4e-3	2.1e5	0.14	9.6e-3	2.2e3	0.22	3.2e-2	5.8e1	0.45
64	1.9e-3	8.1e4	0.26	2.4e-3	5.7e2	0.39	8.2e-3	6.1e1	0.93
128	3.3e-4	1.7e4	0.50	3.2e-4	9.3e1	0.85	1.2e-3	4.8e0	1.91
256	1.8e-5	2.5e2	1.09	1.3e-5	2.6e0	1.97	6.0e-5	2.9e-1	3.93
512	1.1e-6	1.4e0	2.70	7.2e-6	1.2e-1	4.43	9.2e-6	2.7e-2	8.91

Table 5.17: Maximum absolute and relative differences in year-on-year inflation-indexed floor prices computed using MC II and the COS method for different number of cosine-series expansion terms N_{COS}.

To verify the exponential convergence of the COS method as theoretically predicted, we analyse its convergence by varying the number of cosine-series expansion terms N_{COS}. Reference values are obtained using $N_{COS} = 2048$ for year-on-year inflation-indexed caps and floors.

N_{COS}	YIC							
	$T = 1$		$T = 2$		$T = 5$		$T = 10$	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
16	1.2e-2	2.2e1	2.1e-2	9.1e0	5.3e-2	7.2e0	9.6e-2	5.7e0
32	6.0e-3	5.2e0	9.2e-3	1.7e0	2.8e-2	2.4e0	4.9e-2	2.0e0
64	2.1e-3	1.8e0	2.7e-3	8.6e-1	9.0e-3	3.3e-1	1.5e-2	2.6e-1
128	3.9e-4	3.4e-1	3.6e-4	7.8e-2	1.5e-3	1.3e-1	2.0e-3	8.1e-2
256	1.5e-5	2.2e-2	8.1e-6	2.5e-3	4.9e-5	4.2e-3	5.2e-5	1.9e-3
512	5.2e-8	8.4e-5	1.2e-8	4.4e-6	3.5e-7	3.2e-5	3.9e-7	1.5e-5
1024	8.8e-14	1.3e-10	2.2e-15	1.2e-12	4.1e-12	5.6e-10	5.1e-12	2.3e-10

Table 5.18: Maximum absolute difference between the COS method varying amount of cosine-series expansion terms N_{COS} and the COS method with $N_{COS} = 2048$ as reference values for year-on-year inflation-indexed cap prices.

N_{COS}	YIF							
	$T = 1$		$T = 2$		$T = 5$		$T = 10$	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
16	1.7e-2	2.3e5	2.8e-2	2.7e3	8.9e-2	3.9e2	1.6e-1	1.1e2
32	6.4e-3	1.2e5	9.5e-3	2.0e3	3.2e-2	5.7e1	5.6e-2	1.2e1
64	1.9e-3	4.8e4	2.4e-3	5.1e2	8.2e-3	6.0e1	1.3e-2	1.4e1
128	3.3e-4	9.9e3	3.1e-4	8.3e1	1.2e-3	4.6e0	1.6e-3	8.1e-1
256	1.8e-5	1.5e2	9.5e-6	2.2e0	5.2e-5	3.1e-1	5.7e-5	4.9e-2
512	5.6e-8	7.7e-1	1.4e-8	3.3e-3	2.9e-7	1.0e-3	3.3e-7	1.8e-4
1024	7.3e-14	2.0e-6	2.1e-15	4.2e-10	5.0e-12	2.5e-8	6.6e-12	5.2e-9

Table 5.19: Maximum absolute difference between the COS method varying amount of cosine-series expansion terms N_{COS} and the COS method with $N_{COS} = 2048$ as reference values for year-on-year inflation-indexed floor prices.

Like in Section 5.1, we can observe the expected exponential convergence. Note that convergence is slowest when considering relative errors for YIF with maturity $T = 1$, since the option prices under the given model parameters and for the considered strikes, are very small, with the smallest being of order 1e-8.

We compare the integration scheme employed by comparing year-on-year inflation-indexed cap prices with references values obtained with $N_{COS} = 2048$.

n	YIC											
	$T = 2$				$T = 5$				$T = 10$			
	tpz rule		CC rule		tpz rule		CC rule		tpz rule		CC rule	
	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$	L_∞^{COS}	$L_\infty^{rel,COS}$
1	6.4e-9	3.1e-7	6.4e-9	3.1e-7	9.0e-9	1.3e-7	9.0e-9	1.3e-7	3.6e-9	2.0e-8	3.6e-9	2.0e-8
2	1.6e-9	7.7e-8	1.2e-14	7.8e-12	2.3e-9	3.2e-8	6.7e-14	1.1e-11	9.0e-10	5.1e-9	3.5e-14	1.7e-12
4	4.0e-10	1.9e-8	7.4e-18	4.5e-15	5.7e-10	8.1e-9	1.4e-17	1.7e-15	2.3e-10	1.3e-9	1.7e-16	9.1e-15

Table 5.20: Maximum absolute difference between the COS method varying amount of quadrature points n per sub-interval using trapezoidal integration rule and Clenshaw-Curtis quadrature, and the COS method with $N_{COS} = 2048$ as reference values for year-on-year inflation-indexed cap prices.

We observe the same convergence rates as earlier seen with inflation-indexed caps in Section 5.1. Unlike the European options considered, the choice of the numerical integration scheme has a greater impact here, as the transitional characteristic function involves using a numerical integration of two parameters, $\tilde{\phi}(t, T_1, u, v)$ and $\tilde{\phi}(T_1, T_2, u, v)$, instead of just one parameter. Additionally, for year-on-year options, the total price can be viewed as the sum of individual year-on-year option prices over successive yearly intervals, with our method evaluating these options at each interval. As a consequence, as maturity increases, numerical errors accumulate for year-on-year options with longer maturities.

Below, we illustrate option prices for various strike prices using the COS method and MC II with maturities of 1, 2 and 5 years. For maturity of 10 years, we compare the COS method with MC I*. As a result, due to the lower accuracy of MC I*, the differences between the two methods will be larger. Monte Carlo results are shown with a 95% confidence interval.

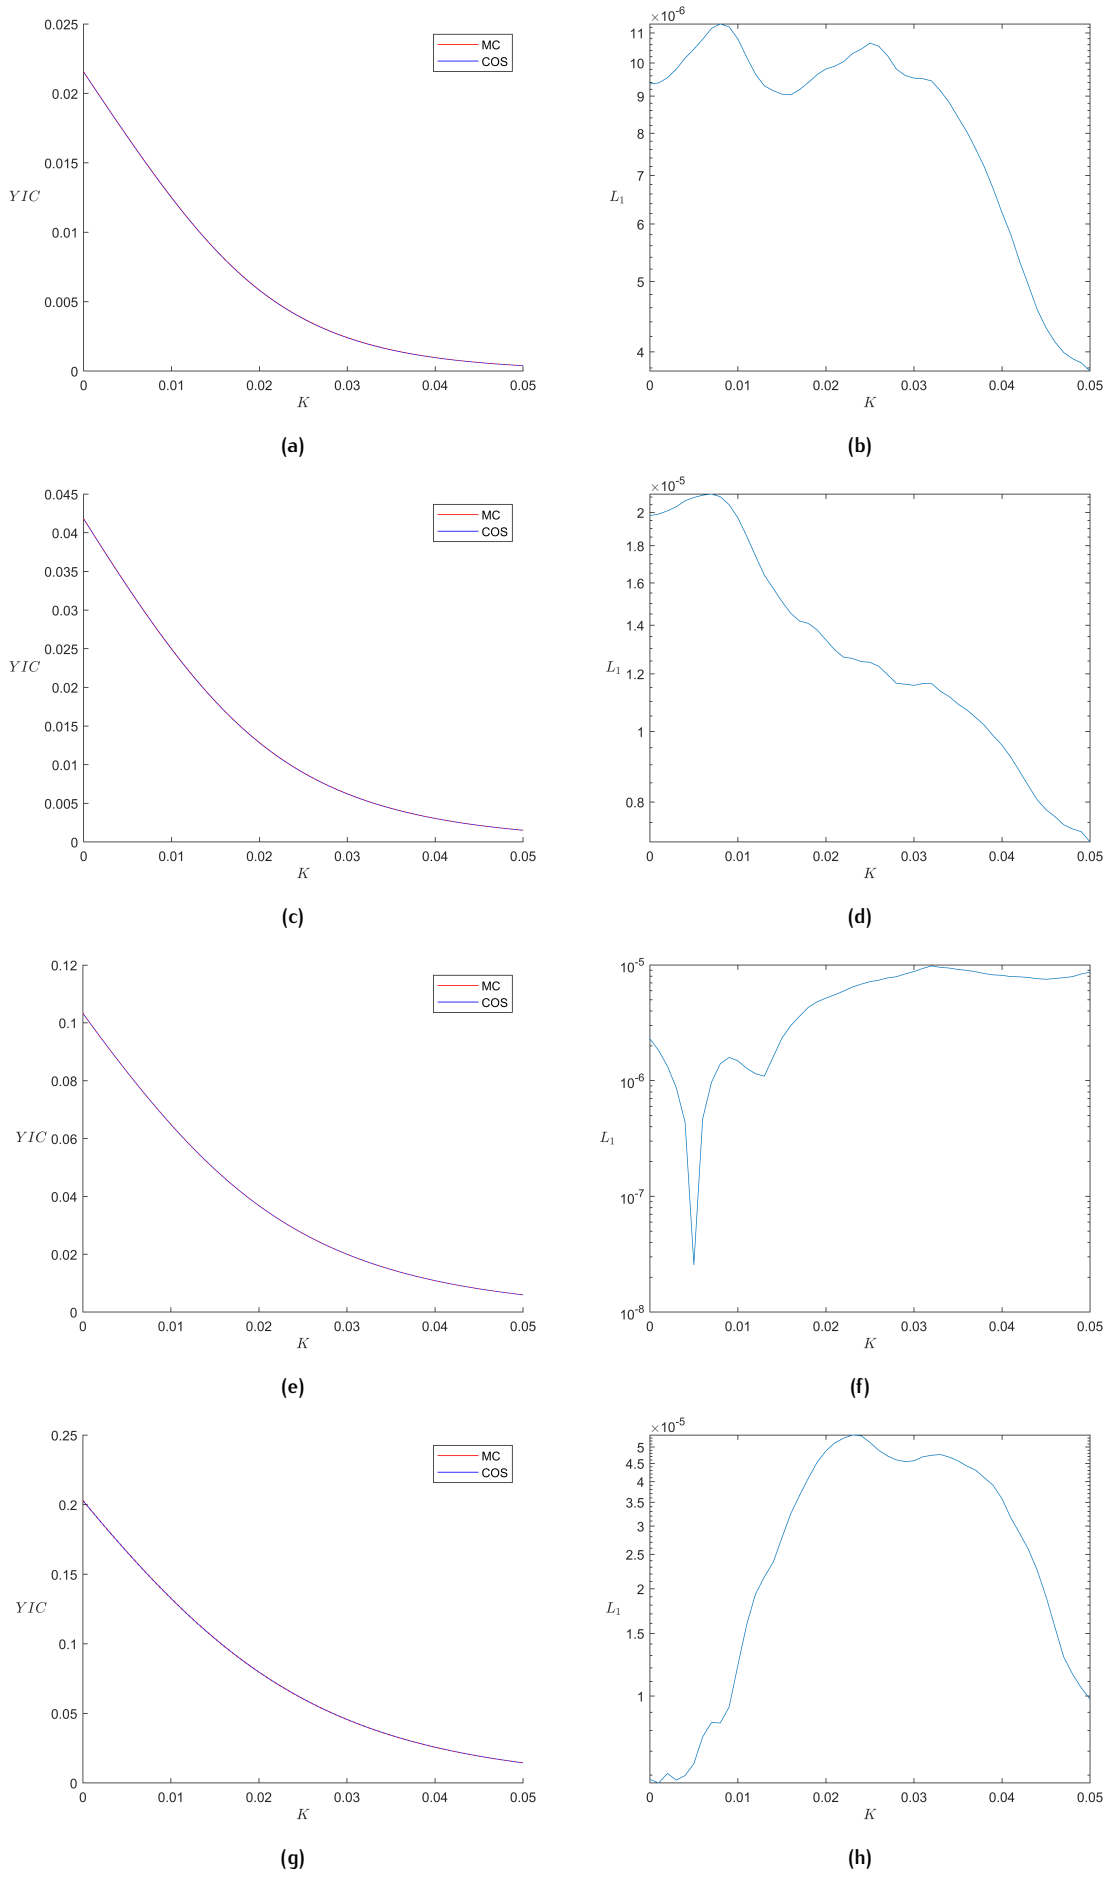

Figure 5.3: Year-on-year inflation-indexed cap prices computed by the COS method and MC II with 95% confidence interval, and absolute difference with maturities of 1 year (a,b), 2 years (c,d), 5 years (e,f) and 10 years (g,h), where only for 10 years MC I* was used instead of MC II.

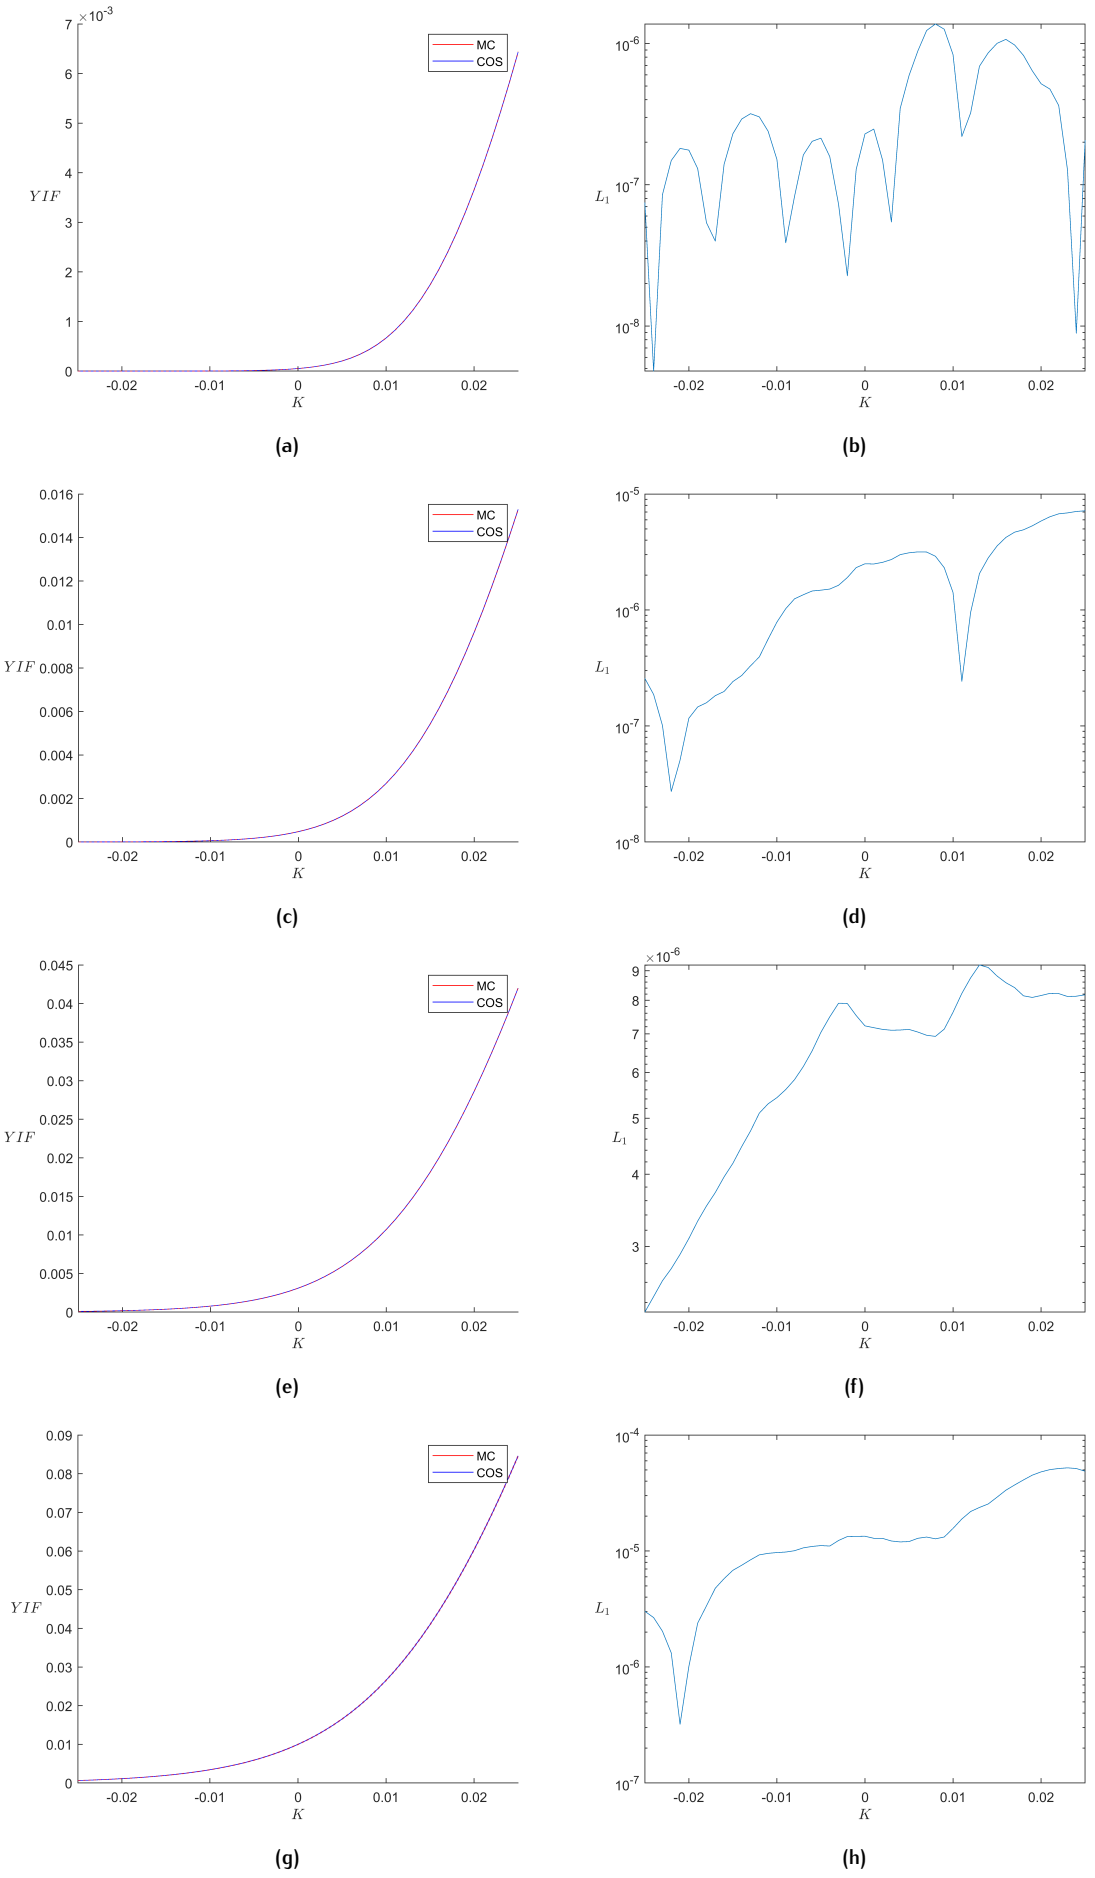

Figure 5.4: Year-on-year inflation-indexed floor prices computed by the COS method and MC II with 95% confidence interval, and absolute difference with maturities of 1 year (a,b), 2 years (c,d), 5 years (e,f) and 10 years (g,h), where only for 10 years MC I* was used instead of MC II.

Compared to the European options considered, the COS method now takes approximately twice as long. This is expected, as computing the transitional characteristic function at each step requires solving twice as many ODEs. Again, we analyse the effect of approximating the transitional characteristic function by Chebyshev polynomials and observe that relatively few polynomials are required in the Cheby-COS method to reach accurate results relative to the COS method with $N_{COS} = 512$.

N_{Cheby}	YIC							
	$T = 1$		$T = 2$		$T = 5$		$T = 10$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	4.5e-2	1.2e2	5.7e-2	3.8e1	1.7e-1	2.8e1	3.4e-1	2.3e1
16	2.5e-4	6.5e-1	6.0e-4	3.9e-1	2.3e-3	3.9e-1	5.7e-3	4.0e-1
32	1.8e-8	4.6e-5	2.3e-7	1.5e-4	6.9e-7	1.2e-4	1.7e-6	1.2e-4
64	9.0e-11	1.7e-7	1.4e-9	9.3e-7	4.1e-9	6.3e-7	4.5e-8	2.7e-6

Table 5.21: Maximum absolute and relative differences in year-on-year inflation-indexed cap prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby}.

N_{Cheby}	YIF							
	$T = 1$		$T = 2$		$T = 5$		$T = 10$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	1.4e-2	9.2e5	4.1e-2	2.0e4	7.0e-2	7.5e2	1.6e-1	2.4e2
16	1.0e-4	6.4e3	1.2e-5	6.1e0	3.6e-4	3.6e0	8.3e-4	1.2e0
32	4.4e-10	2.9e-2	1.8e-7	9.3e-2	4.1e-8	4.1e-4	1.2e-7	1.7e-4
64	9.6e-11	5.7e-3	2.9e-9	1.4e-3	2.0e-9	1.8e-5	3.8e-9	3.9e-6

Table 5.22: Maximum absolute and relative differences in year-on-year inflation-indexed floor prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby}.

When setting N_{Cheby} = 64, the dominant source of numerical error arises from the chosen tolerance level of the ODE solver, which primarily affects the smallest option prices. By repeating the experiment with a stricter error tolerance of 10^{-12} in the ODE solver, we once again observe the exponential convergence of Chebyshev polynomials in the final option prices.

N_{Cheby}	YIC							
	$T=1$		$T=2$		$T=5$		$T=10$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	4.5e-2	1.2e2	5.5e-2	3.6e1	1.7e-1	2.9e1	3.4e-1	2.3e1
16	3.6e-4	9.3e-1	5.5e-4	3.7e-1	2.3e-3	3.9e-1	5.7e-3	4.0e-1
32	2.5e-8	6.4e-5	1.7e-7	1.1e-4	4.8e-7	8.0e-5	1.7e-6	1.2e-4
64	2.6e-14	6.7e-7	5.7e-14	3.5e-11	1.5e-13	2.5e-11	2.7e-13	1.8e-11

Table 5.23: Maximum absolute and relative differences in year-on-year inflation-indexed cap prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby} with an error tolerance set to 10^{-12} in the ODE solver.

N_{Cheby}	YIF							
	$T=1$		$T=2$		$T=5$		$T=10$	
	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$	$\hat{L}_\infty$	$\hat{L}_\infty^{rel}$
8	1.8e-2	4.8e6	4.2e-2	2.0e4	6.5e-2	7.0e2	1.1e-1	1.7e2
16	1.3e-4	3.6e4	4.4e-5	2.3e1	3.9e-4	4.0e0	8.7e-4	1.3e0
32	4.7e-9	1.3e0	2.1e-7	1.1e-1	6.2e-8	6.5e-4	9.3e-8	1.4e-4
64	1.2e-14	3.3e-6	4.0e-14	2.0e-8	6.0e-14	6.4e-10	1.3e-13	2.0e-10

Table 5.24: Maximum absolute and relative differences in year-on-year inflation-indexed floor prices computed using the COS method and the Cheby-COS method for different number of Chebyshev polynomials N_{Cheby} with an error tolerance set to 10^{-12} in the ODE solver.

We fix $N_{COS} = 512$ and $N_{Cheby} = 64$, and price year-on-year inflation-indexed cap and floor prices across all maturities simultaneously in the COS method and Cheby-COS, employing the two numerical integration schemes previously considered. Option prices estimated by Monte Carlo are again reported with a 95% confidence interval.

T	$YIC(K = 0.02)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.00582 ± 0.00006	0.00583 ± 0.00002	0.00582	0.00582	0.00582	0.00582
2	0.01288 ± 0.00011	0.01284 ± 0.00003	0.01283	0.01283	0.01283	0.01283
5	0.03684 ± 0.00020	0.03679 ± 0.00006	0.03678	0.03678	0.03678	0.03678
10	0.07966 ± 0.00034	–	0.07970	0.07970	0.07970	0.07970
T	$YIC(K = 0.04)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.00098 ± 0.00003	0.00097 ± 0.00001	0.00097	0.00097	0.00097	0.00097
2	0.00309 ± 0.00006	0.00305 ± 0.00002	0.00304	0.00304	0.00304	0.00304
5	0.01089 ± 0.00012	0.01087 ± 0.00004	0.01086	0.01086	0.01086	0.01086
10	0.02569 ± 0.00018	–	0.02573	0.02573	0.02573	0.02573
Total CPU time (s)	130	$2.2 \cdot 10^4$	7.0	15.3	1.1	2.2

Table 5.25: Year-on-year inflation-indexed cap prices for different maturities obtained by Monte Carlo with 95% confidence interval, COS method, Cheby-COS method, employing the trapezoidal rule or Clenshaw-Curtis quadrature as a numerical integration scheme.

T	$YIF(K = -0.01)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.0000022 ± 0.0000007	0.0000022 ± 0.0000002	0.0000019	0.0000023	0.0000019	0.0000023
2	0.0000590 ± 0.0000043	0.0000600 ± 0.0000014	0.0000587	0.0000592	0.0000587	0.0000592
5	0.0007752 ± 0.0000199	0.0007748 ± 0.0000063	0.0007797	0.0007802	0.0007797	0.0007802
10	0.0034243 ± 0.0000527	–	0.0034335	0.0034340	0.0034335	0.0034340
T	$YIF(K = 0.01)$					
	MC I*	MC II	COS (tpz)	COS (CC)	Cheby-COS (tpz)	Cheby-COS (CC)
1	0.000668 ± 0.000014	0.000666 ± 0.000004	0.000665	0.000665	0.000665	0.000665
2	0.002708 ± 0.000032	0.002701 ± 0.000010	0.002700	0.002700	0.002700	0.002700
5	0.010731 ± 0.000083	0.010705 ± 0.000026	0.010712	0.010712	0.010712	0.010712
10	0.026540 ± 0.000176	–	0.026556	0.026556	0.026556	0.026556
Total CPU time (s)	130	$2.2 \cdot 10^4$	7.0	15.3	1.1	2.2

Table 5.26: Year-on-year inflation-indexed floor prices for different maturities obtained by Monte Carlo with 95% confidence interval, COS method, Cheby-COS method, employing the trapezoidal rule or Clenshaw-Curtis quadrature as a numerical integration scheme.

T	MC I			
	$YIC(K = 0.02)$	$YIC(K = 0.04)$	$YIF(K = -0.01)$	$YIF(K = 0.01)$
1	0.00538 ± 0.00006	0.00082 ± 0.00002	0.0000020 ± 0.0000006	0.00067 ± 0.00001
2	0.01229 ± 0.00010	0.00283 ± 0.00005	0.0000523 ± 0.0000041	0.00264 ± 0.00003
5	0.03628 ± 0.00020	0.01062 ± 0.00012	0.0007234 ± 0.0000194	0.01037 ± 0.00008
10	0.07936 ± 0.00033	0.02557 ± 0.00019	0.0033516 ± 0.0000564	0.02602 ± 0.00017

Table 5.27: Year-on-year inflation-indexed cap and floor prices prices obtained from MC I with 95% confidence interval. CPU time is 11 s.

We observe that the Cheby-COS method produces results similar to its COS counterpart. However, since it requires solving eight times fewer ODEs, the CPU time is also reduced by approximately a factor of eight. Note that the only differences appear in $YIF(K = -0.01)$ on the order of 10^{-7}, which arise due to the integration scheme employed. The prices computed by the COS method closely align with those from the Monte Carlo method, and the Cheby-COS method offers the additional benefit of significantly reducing CPU time with minimal accuracy loss. In particular, the Cheby-COS method outperforms MC I in both accuracy and computational speed.

5.3 CALIBRATION OF CP2022

In this section, we examine the choice of option pricing methodology in the calibration process. Specifically, we calibrate separately on European call options on a stock index, zero-coupon inflation-indexed caps and floors, and year-on-year inflation-indexed caps and floors. All strikes and maturities considered per option during calibration can be found in Section 3.5.3. We compare two different methods used for pricing options during calibration. The first method is a Monte Carlo method with 100,000 simulations and monthly time steps with antithetic variates as a variance reduction technique. The second method utilizes the Cheby-COS method combined with the trapezoidal integration rule with monthly time steps, to price all options simultaneously. The configurations used per option type are shown below.

	Cheby-COS method	
	N_{COS}	N_{Cheby}
C	192	32
IC / IF	96	32
YIC / YIF	256	32

Table 5.28: Cheby-COS configuration used for calibration.

We initialize the model with the parameters from Q2 2024 and calibrate it for Q3 2024 using a fixed number of optimization steps. We apply this approach to both methods and compare their computational speed and the implied volatility mean squared error (e^2) as in Equation 3.44. We compare the accuracy of both methods by benchmarking the results against reference values obtained with the COS method with $N_{COS} = 2048$ and define the relative implied volatility mean squared error e^2_{rel} per method as

$$e^2_{rel} = \left| \frac{e^2_{ref} - e^2}{e^2_{ref}} \right|.$$

	e^2		e^2_{rel}		Total CPU time (s)	
	MC	Cheby-COS	MC	Cheby-COS	MC	Cheby-COS
C	5.4e-5	6.4e-5	1.4e-1	8.6e-3	1.3e4	2.6e2
IC	3.4e-5	3.5e-5	2.1e-2	4.1e-3		3.2e2
IF	3.5e-5	3.9e-5	1.0e-1	2.9e-3		
YIC	1.8e-5	2.3e-5	2.0e-1	2.0e-2		4.6e2
YIF	9.0e-6	8.5e-6	2.7e-2	8.2e-3		

Table 5.29: Comparison of the Monte Carlo-based approach and the Cheby-COS method for option valuation during calibration. The implied volatility mean squared error (e^2) reported with relative errors e^2_{rel} compared to the COS method with $N_{COS} = 2048$, along with the corresponding computation time in seconds.

We make the following observations. Within the model's framework, the Cheby-COS method yields more accurate option prices compared to the Monte Carlo approach, based on relative errors in Table 5.29. Additionally, the computation time required by the Cheby-COS method is significantly lower than that of the Monte Carlo method. While the Monte Carlo approach offers the flexibility to price any option directly through simulations, the Cheby-COS method requires a distinct implementation for different option payoffs. Nevertheless, the total computation time for the Cheby-COS method remains lower than that of the Monte Carlo method. Furthermore, computation time can be further reduced by leveraging parallel computing, where the option prices for each option type are computed simultaneously using the Cheby-COS method. While the Monte Carlo method can be parallelized by sampling smaller batches of paths, the discretization error remains the main bottleneck, which cannot be solved by parallel computing.

This is because, although distinct simulation paths can be computed in parallel, the step-by-step computation of each individual path remains sequential and cannot be accelerated through parallelization.

The accuracy of option pricers also impacts model calibration. During optimization, the constraints in Equation 3.43 on e^2 are enforced, which are incorporated into the objective function via a penalty method to be maximized. These thresholds are scaled based on the average volatility of the market instruments used for calibration. Since the accuracy of option pricers affects e^2, it also determines whether a large penalty is applied, influencing how model parameters are adjusted. For example, if the threshold for C would be 6.0e-5, the Monte Carlo error estimate would remain below this threshold. However, the more accurate Cheby-COS method reveals that the true error exceeds the threshold, resulting in a significant penalty being added to the objective function to be maximized. Consequently, this could cause the optimization process to converge to different optima.

This page was intentionally left blank.

6 CONCLUSION AND FUTURE RESEARCH

CONCLUSION

In this thesis, we have developed an efficient method for calculating option prices in the CP2022 model, an affine five-state economic framework used by De Nederlandsche Bank to project economic scenarios in the future and employed by all Dutch pension funds for regulatory risk management. Our work is based on a Fourier method called the COS method (Fang and Oosterlee [2009a]), which relies on the availability of the characteristic function of the state factors in CP2022. We have shown how the COS method can be generalized when interest rates are stochastic without changing to the T-forward measure. The option pricing problem can be reduced to the one-dimensional COS method by performing a cosine series expansion of the discounted probability density function. This is achieved by incorporating the integrated process inside the characteristic function, which combines the dynamics of both the underlying and integrated state factors. The characteristic function is obtained numerically by solving a system of Riccati ordinary differential equations (ODEs), a process that can be computationally intensive. There are two main ways to make the process more efficient. Since our time-inhomogeneous affine model exhibits time dependence in only one ODE, for which the solution is known, we can transform the system of ODEs to a time-homogeneous form with an additional numerical integration scheme over the time-dependent component. This allows the characteristic function to be determined for all time points up to the terminal condition. Further speed-up is achieved by reducing the number of times the ODEs need to be solved, through interpolation of the solution using Chebyshev polynomials—a technique combined with the COS method we refer to as the Cheby-COS method.

We compare our methods with the Monte Carlo method and demonstrate that both approaches outperform the Monte Carlo method in terms of computational speed and accuracy, when applied to European options. Additionally, we derive the transitional characteristic function, defined at two future time points, which is necessary for pricing cliquet options. Since this requires solving twice as many times the ODEs as before, both of the aforementioned improvements are crucial for efficiently pricing cliquet options in CP2022. If CP2022 is calibrated separately for each option, our methods outperform Monte Carlo in both accuracy and computation time. Additionally, the improved accuracy of the computed option prices may lead to different final model parameters after optimization.

FUTURE RESEARCH

While the Monte Carlo method offers the flexibility of pricing all option types simultaneously, both of our methods require pricing each option type separately. Computation time can be further reduced by parallelizing the workload for each option type. Currently, calibration is performed using all options simultaneously. Alternatively, calibration could be performed in batches, starting with a selected subset of options and gradually add the remaining options as the calibration improves. Using this approach, our option pricing approach could serve as an alternative to Monte Carlo simulations and calibration could be accelerated further if options with the shortest maturities are considered first. However, this approach may risk convergence to a local optimum, which could result in suboptimal results across the entire set of options. We leave this approach for future research to explore further.

In a recent study (Junike [2024]), the number of Fourier-cosine series expansion terms used in the COS method was adaptively selected based on the truncation range to ensure accuracy within a predefined tolerance. Since model parameters change during calibration, the truncation range will also vary accordingly. Rather than fixing a high number of expansion terms to accommodate potentially large truncation ranges, computation time can be reduced by adapting the number of terms to the current truncation range for the desired accuracy, since additional expansion terms involve solving more ODEs numerically.

Finally, in our research we have not considered pricing swaptions. While our method for pricing cliquet options can be extended to swaptions by using a reiterated tower law for conditional expectations, this would significantly increase the number of times the ODEs need to be solved. The works of Schrager and Pelsser [2006] shows great promise for future research on pricing swaptions in the CP2022 model. In their work, they derive the dynamics of the swap rate and transition from the risk-neutral measure to the swap measure in an affine model. They assert that the additional stochastic drift in the state factors under the swap measure, as well as the stochastic diffusion of the swap rate, are low-variance martingales. To restore the model's affine structure, they approximate these factors by their time-zero values. By utilizing the characteristic function of the swap rate, the same Fourier methods can be applied for efficient option valuation.

BIBLIOGRAPHY

Andersen, L. B. (2007). Efficient simulation of the heston stochastic volatility model. *Available at SSRN 946405*.

Aparicio, F. M. and Estrada, J. (2001). Empirical distributions of stock returns: European securities markets, 1990-95. *The European Journal of Finance*, 7(1):1–21.

Ari, M. A., Mulas-Granados, M. C., Mylonas, M. V., Ratnovski, M. L., and Zhao, W. (2023). *One hundred inflation shocks: Seven stylized facts*. International Monetary Fund.

Bank of England (2023). Transition from libor to risk-free rates. https://www.bankofengland.co.uk/markets/transition-to-sterling-risk-free-rates-from-libor. Accessed on 02-09-2024.

Bender, C. M. and Orszag, S. A. (2010). *Advanced mathematical methods for scientists and engineers I*. Springer, New York, NY.

Bianchetti, M., Kucherenko, S., and Scoleri, S. (2015). Pricing and risk management with high-dimensional quasi-monte carlo and global sensitivity analysis. *Wilmott*, 2015(78):46–70.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3(1):167–179.

Boyd, J. P. (2001). *Chebyshev and Fourier spectral methods*. Courier Corporation.

Boyle, P. P., Evnine, J., and Gibbs, S. (1989). Numerical evaluation of multivariate contingent claims. *The Review of Financial Studies*, 2(2):241–250.

Brand, C., Bielecki, M., and Penalver, A. (2018). The natural rate of interest: estimates, drivers, and challenges to monetary policy JEL Classification: E52, E43. Occasional Paper Series 217, European Central Bank.

Brennan, M. J. and Xia, Y. (2002). Dynamic asset allocation under inflation. *The journal of finance*, 57(3):1201–1238.

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

Broadie, M. and Kaya, Ö. (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations research*, 54(2):217–231.

Broadie, M. and Yamamoto, Y. (2003). Application of the fast gauss transform to option pricing. *Management Science*, 49(8):1071–1088.

Bromwich, T. J. I. (2005). *An introduction to the theory of infinite series*, volume 335. American Mathematical Soc.

Carr, P. and Madan, D. (1999). Option valuation using the fast fourier transform. *Journal of computational finance*, 2(4):61–73.

Chebyshev, P. L. (1853). *Théorie des mécanismes connus sous le nom de parallélogrammes*. Imprimerie de l'Académie impériale des sciences.

Chen, D. H. and van Wijnbergen, S. J. (2020). Redistributive consequences of abolishing uniform contribution policies in pension funds. *De Economist*, 168(3):305–341.

Cheng, R. C. H. and Feast, G. M. (1979). Some simple gamma variate generators. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(3):290–295.

Commission Parameters (2014). Advies commissie parameters. *Publicatie Rijksoverheid*.

Commission Parameters (2019). Advies commissie parameters. *Publicatie Rijksoverheid*.

Commission Parameters (2022a). Advies commissie parameters. *Publicatie Rijksoverheid*.

Commission Parameters (2022b). Technical appendix: Specification of the cp2022 model. Technical report, Government of the Netherlands.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of financial Economics*, 7(3):229–263.

CPB (2024). Centraal economisch plan 2024. *CPB*.

Dai, Q. and Singleton, K. (2003). Term structure dynamics in theory and reality. *The Review of financial studies*, 16(3):631–678.

Dai, Q. and Singleton, K. J. (2000). Specification analysis of affine term structure models. *The Journal of Finance*, 55(5):1943–1978.

De Graaf, C. S., Feng, Q., Kandhai, D., and Oosterlee, C. W. (2014). Efficient computation of exposure profiles for counterparty credit risk. *International Journal of Theoretical and Applied Finance*, 17(04):1450024.

de Jong, F. (2000). Time series and cross-section information in affine term-structure models. *Journal of Business & Economic Statistics*, 18(3):300–314.

De Nederlandsche Bank (2022). DNB publiceert definitieve scenariosets bij Wet toekomst pensioenen. https://www.dnb.nl/voor-de-sector/open-boek-toezicht/sectoren/pensioenfondsen/dnb-publiceert-definitieve-scenariosets-bij-wet-toekomst-pensioenen/. [Accessed 21-07-2024].

De Nederlandsche Bank (2024). Pension funds (macroeconomic). https://www.dnb.nl/en/statistics/dashboards/pension-funds/pension-funds-macroeconomic/. [Accessed 11-08-2024].

Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(3):463–520.

Delbaen, F. and Schachermayer, W. (1998). The fundamental theorem of asset pricing for unbounded stochastic processes. *Mathematische Annalen*, 312(2):215–250.

Dennis, J. E. and Schnabel, R. B. (1983). *Numerical methods for unconstrained optimization and nonlinear equations*. Prentice-Hall series in computational mathematics. Prentice Hall, Old Tappan, NJ.

Detlefsen, K. and Härdle, W. K. (2006). Calibration risk for exotic options. *Available at SSRN 2894410*.

Dimson, E., Marsh, P., and Staunton, M. (2022). Credit suisse global investment returns yearbook 2022 summary edition. *Zurich: Credit Suisse Research Institute*.

Dormand, J. R. and Prince, P. J. (1980). A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26.

Draper, N. (2014). A financial market model for the netherlands. *Netspar*.

Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.

EIOPA (2020). Background document on the opinion on the 2020 review of solvency ii. *EIOPA-BoS-20/750*.

Elliott, D. (1965). Truncation errors in two chebyshev series approximations. *Mathematics of Computation*, 19(90):234–248.

European Central Bank (2024). Long-term interest rate statistics for eu member states. https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/long_term_interest_rates/html/index.en.html. Accessed: 2024-10-06.

Eurostat (2024). Pension systems in the eu — statistics explained. https://ec.europa.eu/eurostat/statistics-explained/index.php?oldid=514007. Accessed: 2024-08-25.

Ezzine, A., Alla, A., and Raissi, N. (2024). A legendre–galerkin spectral method for option pricing under regime switching models. *Results in Applied Mathematics*, 24:100505.

Fang, F., Jönsson, H., Oosterlee, C. W., and Schoutens, W. (2009). Fast valuation and calibration of credit default swaps under lévy dynamics. *Available at SSRN 1628672*.

Fang, F. and Oosterlee, C. W. (2009a). A novel pricing method for european options based on fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, 31(2):826–848.

Fang, F. and Oosterlee, C. W. (2009b). Pricing early-exercise and discrete barrier options by fourier-cosine series expansions. *Numerische Mathematik*, 114(1):27–62.

Fang, F. and Oosterlee, C. W. (2011). A fourier-based valuation method for bermudan and barrier options under heston's model. *SIAM Journal on Financial Mathematics*, 2(1):439–463.

FCA (2024). About libor transition. https://www.fca.org.uk/markets/libor-transition. Accessed on 02-09-2024.

Feller, W. (1957). *An introduction to probability theory and its applications.* Wiley & Sons.

Feng, L. and Linetsky, V. (2008). Pricing discretely monitored barrier options and defaultable bonds in lévy process models: a fast hilbert transform approach. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 18(3):337–384.

Feng, Q. and Oosterlee, C. W. (2016). Wrong way risk modeling and computation in credit valuation adjustment for european and bermudan options. *Available at SSRN 2852819*.

Filipović, D. (2005). Time-inhomogeneous affine processes. *Stochastic Processes and their Applications*, 115(4):639–659.

Foulon, S. et al. (2010). Adi finite difference schemes for option pricing in the heston model with correlation. *International Journal of Numerical Analysis & Modeling*, 7(2):303–320.

Gaudenzi, M. and Zanette, A. (2011). Pricing cliquet options by tree methods. *Computational Management Science*, 8:125–135.

Glasserman, P. (2004). *Monte Carlo methods in financial engineering*, volume 53. Springer.

Heath, D., Jarrow, R., and Morton, A. (1992a). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105.

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

Higham, D. J. (2004). *An Introduction to Financial Option Valuation: Mathematics, Stochastics and Computation*. Cambridge University Press.

Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.

Hou, D. and Skeie, D. R. (2014). Libor: Origins, economics, crisis, scandal, and reform. *FRB of New York Staff Report*, (667).

Hull, J. (2017). *Options, Futures, and Other Derivatives, global edition*. Pearson Higher Education, 9 edition.

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

James, J. and Webber, N. (2000). *Interest rate modelling*. Wiley-Blackwell Publishing Ltd.

Joslin, S., Le, A., and Singleton, K. J. (2013). Why gaussian macro-finance term structure models are (nearly) unconstrained factor-vars. *Journal of Financial Economics*, 109(3):604–622.

Junike, G. (2024). On the number of terms in the cos method for european option pricing. *Numerische Mathematik*, 156(2):533–564.

Junike, G. and Stier, H. (2024). From characteristic functions to multivariate distribution functions and european option prices by the damped cos method.

Kloeden, P. E. and Platen, E. (1992). *Numerical Solution of Stochastic Differential Equations*. Springer Berlin, Heidelberg.

Koijen, R. S. J., Nijman, T. E., and Werker, B. J. M. (2009). When Can Life Cycle Investors Benefit from Time-Varying Bond Risk Premia? *The Review of Financial Studies*, 23(2):741–780.

Korn, R., Temoçin, B. Z., and Wenzel, J. (2017). Applications of the central limit theorem for pricing cliquet-style options. *European Actuarial Journal*, 7:465–480.

Kruse, S. and Nögel, U. (2005). On the pricing of forward starting options in heston's model on stochastic volatility. *Finance and Stochastics*, 9:233–250.

Lee, R. W. et al. (2004). Option pricing by transform methods: extensions, unification and error control. *Journal of Computational Finance*, 7(3):51–86.

Leitao, Á., Oosterlee, C. W., Ortiz-Gracia, L., and Bohte, S. M. (2018). On the data-driven cos method. *Applied Mathematics and Computation*, 317:68–84.

Liu, S., Borovykh, A., Grzelak, L. A., and Oosterlee, C. W. (2019). A neural network-based framework for financial model calibration. *Journal of Mathematics in Industry*, 9(1):9.

Lord, R. and Kahl, C. (2006). Optimal fourier inversion in semi-analytical option pricing. Technical report, Tinbergen Institute Discussion Paper.

Lord, R., Koekkoek, R., and Dijk, D. V. (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10(2):177–194.

Mason, J. C. and Handscomb, D. C. (2002). *Chebyshev Polynomials*. Chapman & Hall/CRC.

Mercer, C. (2024). Mercer cfa institute global pension index 2024. *CFA institute, Monash University*.

Ministry of Social Affairs and Employment (2022). Wet toekomst pensioenen. https://www.eerstekamer.nl/wetsvoorstel/36067_wet_toekomst_pensioenen. [Accessed 29-10-2024].

Muns, S. (2015). A financial market model for the netherlands: A methodological refinement. Technical report, CPB.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modeling of yield curves. *Journal of business*, pages 473–489.

Øksendal, B. (2003). *Stochastic Differential Equations*. Springer Berlin Heidelberg, Berlin, Heidelberg.

Ortiz-Gracia, L. and Oosterlee, C. W. (2013). Robust pricing of european options with wavelets and the characteristic function. *SIAM Journal on Scientific Computing*, 35(5):B1055–B1084.

Ortiz-Gracia, L. and Oosterlee, C. W. (2014). Efficient var and expected shortfall computations for nonlinear portfolios within the delta-gamma approach. *Applied Mathematics and Computation*, 244:16–31.

Protter, M. and Morrey, C. (2012). *Intermediate Calculus*. Undergraduate Texts in Mathematics. Springer New York.

Ruijter, M. J. and Oosterlee, C. W. (2012). Two-dimensional fourier cosine series expansion method for pricing financial options. *SIAM Journal on Scientific Computing*, 34(5):B642–B671.

Ruijter, M. J., Oosterlee, C. W., and Aalbers, R. F. (2013). On the fourier cosine series expansion method for stochastic control problems. *Numerical Linear Algebra with Applications*, 20(4):598–625.

Schrager, D. F. and Pelsser, A. A. (2006). Pricing swaptions and coupon bond options in affine term structure models. *Mathematical Finance*, 16(4):673–694.

Shreve, S. E. et al. (2004). *Stochastic calculus for finance II: Continuous-time models*, volume 11. Springer.

Stoxx (2024). Stoxx strategy indexguide. https://www.stoxx.com/document/Indices/Common/Indexguide/stoxx_strategy_guide.pdf. [Accessed 22-07-2024].

Trefethen, L. N. (2019). *Approximation theory and approximation practice, extended edition*. SIAM.

Van Beek, M., Mandjes, M., Spreij, P., and Winands, E. (2020). Regime switching affine processes with applications to finance. *Finance and Stochastics*, 24:309–333.

van Haastrecht, A., Lord, R., Pelsser, A., and Schrager, D. (2009). Pricing long-dated insurance contracts with stochastic interest rates and stochastic volatility. *Insurance: Mathematics and Economics*, 45(3):436–448.

van Haastrecht, A. and Pelsser, A. (2011). Generic pricing of fx, inflation and stock options under stochastic interest rates and stochastic volatility. *Quantitative Finance*, 11(5):665–691.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188.

von Sydow, L., Josef Höök, L., Larsson, E., Lindström, E., Milovanović, S., Persson, J., Shcherbakov, V., Shpolyanskiy, Y., Sirén, S., Toivanen, J., et al. (2015). Benchop–the benchmarking project in option pricing. *International Journal of Computer Mathematics*, 92(12):2361–2379.

von Sydow, L., Milovanović, S., Larsson, E., In't Hout, K., Wiktorsson, M., Oosterlee, C. W., Shcherbakov, V., Wyns, M., Leitao, A., Jain, S., et al. (2019). Benchop–slv: The benchmarking project in option pricing–stochastic and local volatility problems. *International Journal of Computer Mathematics*, 96(10):1910–1923.

Waldvogel, J. (2006). Fast construction of the fejér and clenshaw–curtis quadrature rules. *BIT Numerical Mathematics*, 46:195–202.

Weierstrass, K. (1885). Über die analytische darstellbarkeit sogenannter willkürlicher functionen einer reellen veränderlichen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin*, 2(633-639):364.

Windcliff, H., Forsyth, P., and Vetzal, K. (2006). Numerical methods and volatility models for valuing cliquet options. *Applied Mathematical Finance*, 13(4):353–386.

Wolfe, P. (1969). Convergence conditions for ascent methods. *SIAM Review*, 11(2):226–235.

World Bank (2023). Inflation, consumer prices (annual Accessed: 2024-10-06.

Zhang, B. and Oosterlee, C. W. (2013). Efficient pricing of european-style asian options under exponential lévy processes based on fourier cosine expansions. *SIAM Journal on Financial Mathematics*, 4(1):399–426.

This page was intentionally left blank.

A APPENDIX

A.1 A PARAMETRIZATION OF K AND M THAT ENSURES POSITIVE EIGENVALUES

The mapping

$$\begin{bmatrix} x_1 & 0 & 0 \\ x_2 & x_3 & x_4 \\ x_5 & x_6 & x_7 \end{bmatrix} \mapsto \begin{bmatrix} e^{x_7} & 0 & 0 \\ x_6 & x_1 & x_2 \\ x_5 & x_* & e^{x_3} - x_1 \end{bmatrix}, \quad x_* = (x_2)^{-1} \left(x_1 \left(e^{x_3} - x_1 \right) - \frac{e^{2x_3}}{1 + e^{x_4}} \right)$$

creates matrices with positive real eigenvalues $\forall x_i \in \mathbb{R}$ and it is invertible. After the parametrization, the eigenvalues equal $\lambda_1 = e^{x_7}$ and $\lambda_{2,3} = \frac{1}{2} e^{x_3} \left(1 \pm 1/\sqrt{1 + e^{-x_4}} \right)$.

A.2 VEGA VALUES FOR OPTIONS

The implied volatilities σ quoted on the market follow from pricing options under Black's model.

$$\mathbf{C}(t_0, T, K) = S_{t_0} \Phi\left(d_+\right) - K P\left(t_0, T\right) \Phi\left(d_-\right),$$

$$d_\pm = \frac{\ln\left(\frac{S}{KP(t_0,T)}\right) \pm \frac{1}{2}\sigma^2 \left(T - t_0\right)}{\sigma\sqrt{T - t_0}},$$

$$\mathbf{PSW}(t_0, T_a, T_b, K) = \left((s_{ab} - K)\Phi\left(\frac{s_{ab} - K}{\sigma\sqrt{T_a - t_0}}\right) + \sigma\sqrt{T_a - t_0}\, \varphi\left(\frac{s_{ab} - K}{\sigma\sqrt{T_a - t_0}}\right) \right) \sum_{T_k = T_a + 1}^{T_b} P(t_0, T_k),$$

$$s_{ab} = \frac{P\left(t_0, T_a\right) - P\left(t_0, T_b\right)}{\sum_{T_k = T_a + 1}^{T_b} P\left(t_0, T_k\right)},$$

$$\mathbf{YIC}(t_0, T, K) = \sum_{T_k = t_0 + 1}^{T} P\left(t_0, T_k\right) \left(F_k \Phi\left(d_{k+}\right) - (1 + K)\Phi\left(d_{k-}\right) \right),$$

$$\mathbf{YIF}(t_0, T, K) = \sum_{T_k = t_0 + 1}^{T} P\left(t_0, T_k\right) \left(-F_k \Phi\left(-d_{k+}\right) + (1 + K)\Phi\left(-d_{k-}\right) \right),$$

$$d_{k\pm} = \frac{\ln\left(\frac{F_k}{1+K}\right) \pm \frac{1}{2}\sigma^2}{\sigma}, \quad F_k = \frac{P^R\left(t_0, T_k\right) / P\left(t_0, T_k\right)}{P^R\left(t_0, T_{k-1}\right) / P\left(t_0, T_{k-1}\right)},$$

and

$$\mathbf{IC}(t_0,T,K) = P(t_0,T)\left(F\Phi\left(\tilde{d}_+\right) - (1+K)^T\Phi\left(\tilde{d}_-\right)\right),$$
$$\mathbf{IF}(t_0,T,K) = P(t_0,T)\left(-F\Phi\left(-\tilde{d}_+\right) + (1+K)^T\Phi\left(-\tilde{d}_-\right)\right),$$
$$\tilde{d}_\pm = \frac{\ln\left(\frac{F}{(1+K)^T}\right) \pm \frac{1}{2}\sigma^2(T-t_0)}{\sigma\sqrt{T-t_0}}, \quad F = \frac{P^R(t_0,T)}{P(t_0,T)}.$$

The vega values $\mho$ by definition equal

$$\frac{\partial \mathbf{C}(t_0,T,K)}{\partial\sigma} = S_{t_0}\varphi(d_+)\sqrt{T-t_0},$$
$$\frac{\partial \mathbf{PSW}(t_0,T_a,T_b,K)}{\partial\sigma} = \varphi\left(\frac{s_{ab}-K}{\sigma\sqrt{T_a-t_0}}\right)\sqrt{T_a-t_0}\sum_{T_k=T_a+1}^{T_b} P(t_0,T_k),$$
$$\frac{\partial \mathbf{YIC}(t_0,T,K)}{\partial\sigma} = \frac{\partial \mathbf{YIF}(t_0,T,K)}{\partial\sigma} = \sum_{T_k=t_0+1}^{T} P(t_0,T_k)F_k\varphi(d_{k+}),$$
$$\frac{\partial \mathbf{IC}(t_0,T,K)}{\partial\sigma} = \frac{\partial \mathbf{IF}(t_0,T,K)}{\partial\sigma} = P^R(t_0,T)\varphi\left(\tilde{d}_+\right)\sqrt{T-t_0}.$$

A.3 INITIAL PARAMETERS

Table A.1: List of parameters for the sets published in Q2 2024.

Parameter	Value
$\mathbb{E}^{\mathbb{P}} v_\infty$	7.24e-2
$\mathbb{E}^{\mathbb{P}} r_\infty$	7.11e-3
$\mathbb{E}^{\mathbb{P}} \pi_\infty$	4.34e-3
$\mathbb{E}^{\mathbb{Q}} v_\infty$	1.24e-1
$\mathbb{E}^{\mathbb{Q}} r_\infty$	7.06e-2
$\mathbb{E}^{\mathbb{Q}} \pi_\infty$	1.18e-2
K_{vv}	2.21
K_{vr}	3.72e-1
$K_{v\pi}$	3.20e-1
K_{rr}	2.84e-1
$K_{r\pi}$	-4.83e-2
$K_{\pi r}$	-3.47e-1
$K_{\pi\pi}$	2.57e-1
M_{vv}	1.31
M_{vr}	2.46e-1
$M_{v\pi}$	2.15e-1
M_{rr}	2.79e-2
$M_{r\pi}$	9.62e-4
$M_{\pi r}$	-6.67e-2
$M_{\pi\pi}$	6.68e-2
ω	5.66e-1
σ_{vr}	1.28e-1
$\sigma_{v\pi}$	8.13e-2
σ_{r1}	-2.99e-3
$\sigma_{\pi 1}$	1.43e-3
σ_{r2}	-4.63e-3
$\sigma_{\pi 2}$	-3.16e-3
γ_{11}	8.73e1
γ_{22}	2.30e-67
γ_{33}	2.40e5
γ_{44}	8.67e3
η_s	6.01e-2
η_π	1.55e-2
σ_{S1}	-4.36-1
σ_{S2}	8.61e-3
σ_{S3}	4.72e-3
σ_{S4}	9.21e-4
σ_{S5}	-3.59e-4
$\sigma_{\Pi 1}$	6.93e-3
$\sigma_{\Pi 2}$	4.81e-4
$\sigma_{\Pi 3}$	-2.23e-3
$\sigma_{\Pi 4}$	0
$\sigma_{\Pi 5}$	-3.90e-4
v_0	1.65e-2
r_0	-1.43e-3
π_0	5.26e-3
$\ln S_0$	0
$\ln \Pi_0$	0

This page was intentionally left blank.