Applying stochastic volatility models in the risk-neutral and real-world probability measures

University of Pretoria

Faculty of Natural and Agricultural Sciences
Department of Actuarial Science

by

Alexis Jacques Levendis

11001993

Submitted in fulfilment of the requirements for the degree of
Philosophiae Doctor in Mathematical Science

Supervisor: Professor Eben Maré

30 June, 2023

Alexis Jacques Levendis

Applying stochastic volatility models in the risk-neutral and real-world probability measures

30 June, 2023

Supervisor: Professor Eben Maré

University of Pretoria
Department of Actuarial Science
Private Bag X20, Hatfield
Pretoria
0028

Abstract

Stochastic volatility models have become immensely popular since their introduction in 1993 by Heston. This is because their dynamics are more consistent with market behaviour compared to the standard Black-Scholes model. More specifically, stochastic volatility models can somewhat capture the asymmetric distribution often observed in daily equity returns. Numerous extensions to the stochastic volatility model of Heston have since been proposed, including jumps and stochastic interest rates. Due to their complex dynamics, numerical methods such as Monte Carlo simulation, the fast Fourier transform (FFT), and the efficient method of moments (EMM) are often required to calibrate and implement stochastic volatility models. In this thesis, we explore the application of stochastic volatility models to a variety of problems for which research is still in its infancy phase. We consider the pricing of embedded derivatives in the South African life insurance industry given the illiquid derivatives market; the pricing of rainbow and spread options that depend on two underlying assets; the calibration of stochastic volatility models with jumps to historical equity returns; and the use of stochastic volatility models in static hedging. Our findings suggest that stochastic interest rates are the dominant risk driver when pricing long-dated contingent claims; the FFT significantly outperforms Monte Carlo simulation in terms of efficiency; jumps are an important factor required to explain daily equity returns; and static hedging is a simple and effective way to replicate vanilla and exotic options.

List of Publications

- Levendis, A. and Maré, E. (2022a). An Economic Scenario Generator for Embedded Derivatives in South Africa. *South African Actuarial Journal*, 22(1), 79-118.
- Levendis, A. and Maré, E. (2022b). Efficient Pricing of Spread Options with Stochastic Rates and Stochastic Volatility. *Journal of Risk and Financial Management*, 15(11), 504.
- Venter, P., Levendis A. and Maré, E. (2022). Collateralised option pricing in a South African context: A Univariate GARCH approach. *Cogent Economics & Finance*, 10(1).
- Levendis, A. and Maré, E. (2023a). On the calibration of stochastic volatility models to estimate the real-world measure used in option pricing. *ORiON*, 39(1), 65-91.
- Levendis, A. and Maré, E. (2023b). Pricing Two-Asset Rainbow Options with the Fast Fourier Transform. *South African Statistical Journal*, 57(1).
- Levendis, A. and Maré, E. (2023c). Static Hedging of Vanilla and Exotic Options in South Africa. *Accepted and presented at the AFRIC 2023: Actuarial, Finance, Risk and Insurance Congress, Victoria Falls*.

Acknowledgement

My sincere gratitude goes to the following people, without whom this dissertation would not be possible:

- Professor Eben Maré for your enthusiasm, patience and motivation. My PhD journey has been extremely rewarding under your guidance. Thank you for being a great mentor and always being willing to share your wealth of knowledge.
- My wife, Annika Levendis, for your endless love and support. Thank you for motivating me to reach the finish line.
- My baby boy, Markus John Levendis, for bringing the biggest smile to my face. I hope that one day you will read this dissertation and that it will motivate you to achieve great things.
- My twin sister, Emily-Kate Olivier, for your enthusiasm and friendship. You are always so much fun to be around.
- My parents-in-law, Chris and Annelie Minnie, for your words of encouragement.
- My parents, John and Desirée Levendis, for being the best parents I could have asked for. I am proud to be your son.

God in Heaven

To Annika and Markus

Declaration

I declare that this thesis has been composed solely by me and that it has not been submitted, in whole or in part, in any previous application for a degree.

Pretoria, 30 June, 2023

Alexis Jacques Levendis

Contents

1 Introduction **1**

1.1 Background 1
1.2 Research Objectives 3
1.3 Thesis Structure 4

2 An Economic Scenario Generator for Embedded Derivatives in South Africa **7**

2.1 Introduction 7
2.2 Literature Review 9
2.3 The Heston-Hull-White Model 12
2.3.1 The Heston-Hull-White Characteristic Function 14
2.3.2 Heston-Hull-White Calibration 15
2.4 Stochastic Mortality 16
2.4.1 Cox-Ingersoll-Ross Model 17
2.4.2 AR(1)-ARCH(1) Model 19
2.5 Guaranteed Minimum Maturity and Death Benefits 20
2.6 Data 22
2.7 Numerical Results 28
2.7.1 Pricing 28
2.7.2 Hedging 40
2.8 Conclusion 44

3 Pricing Two-Asset Rainbow Options with the Fast Fourier Transform **45**

3.1 Introduction 45
3.2 Stock Price Dynamics for Two-Asset Options 47
3.2.1 Two-Factor Geometric Brownian Motion 48
3.2.2 Three-Factor Stochastic Volatility 48
3.3 Rainbow Options, Fourier Transform and Characteristic Functions 49
3.3.1 Best-of-2 and Worst-of-2 Call Options 49
3.3.2 The Result of Eberlein, Glau, and Papapantoleon 50
3.3.3 Two-Factor gBm Characteristic Function 51
3.3.4 Three-Factor Stochastic Volatility Characteristic Function 52

3.4 The Two-Dimensional FFT Method . . . 52
3.5 Numerical Results . . . 54
3.5.1 FFT and Two-Factor gBm . . . 54
3.5.2 Testing the Truncation Width . . . 56
3.5.3 FFT and Three-Factor Stochastic Volatility . . . 57
3.6 Conclusion . . . 59

4 Efficient Pricing of Spread Options with Stochastic Rates and Stochastic Volatility 60
4.1 Introduction . . . 60
4.2 The Two-Asset Heston–Hull–White Model . . . 62
4.3 The Result of Grzelak and Oosterlee . . . 64
4.4 The Two-Asset Heston–Hull–White Characteristic Function . . . 65
4.5 The Result of Hurd and Zhou (Extended) . . . 68
4.6 Numerical Results . . . 70
4.6.1 Implementation Testing . . . 71
4.6.2 Convergence . . . 71
4.6.3 Impact of Stochastic Interest Rates . . . 73
4.7 Conclusions . . . 74

5 Application of Stochastic Volatility Models in the Real-World Measure 75
5.1 Introduction . . . 75
5.2 Stochastic Volatility Models . . . 77
5.2.1 The Heston Stochastic Volatility Model . . . 77
5.2.2 The Bates Stochastic Volatility Jump Model . . . 79
5.2.3 Efficient Method of Moments . . . 82
5.3 Empirical Results . . . 84
5.3.1 S&P500 . . . 84
5.3.2 FTSE/JSE Top40 . . . 92
5.4 Volatility Targeting . . . 96
5.5 Conclusion . . . 103

6 Static Hedging of Vanilla and Exotic Options in a South African Context 105
6.1 Introduction . . . 105
6.2 Stochastic Volatility Double Jump Model . . . 106
6.2.1 SVJJ Dynamics . . . 106
6.2.2 SVJJ Characteristic Function . . . 107
6.3 Static Hedging . . . 108
6.3.1 Choie and Novometsky Optimisation . . . 108
6.3.2 Armstrong et al. Optimisation . . . 109
6.4 Results . . . 110

6.4.1 SVJJ $\mathbb{P}$-Measure Calibration 110
6.4.2 SVJJ $\mathbb{Q}$-Measure Calibration 112
6.4.3 Static Hedging Performance for Vanilla Call Option 115
6.4.4 Static Hedging Performance for Spread Call Option 121
6.5 Conclusion 126

7 Conclusion **128**

Bibliography **131**

List of Notation

μ	Annualised expected return
$S(t)$	Share price at time t
$v(t)$	Annualised variance at time t
$r(t)$	Annualised short-term interest rate at time t
$\mu(x+t,t)$	Force of mortality at time t and age $x+t$
$\varphi(\cdot)$	Standard normal density function
κ	Mean reversion speed of volatility
$\bar{v}$	Long-term mean of volatility
σ_v	Volatility of volatility
ρ	Correlation coefficient
λ_r	Mean reversion speed of interest rate
$\theta(t)$	Time-dependent mean of interest rate
η	Volatility of interest rate
γ	Mean reversion speed of mortality
ω	Long-term mean of mortality
ξ	Volatility of mortality
$N(t)$	Poisson process
λ	Jump intensity in Poisson process
J	Stock jump size
Z	Variance jump size
$W(t)$	Brownian motion
K	Option strike price
T	Option maturity/Length of time series
$\mathbb{Q}$	Risk-neutral probability measure
$\mathbb{P}$	Real-world probability measure
$\phi(\cdot)$	One-dimensional characteristic function
$\phi(\cdot,\cdot)$	Two-dimensional characteristic function
$S_x(0,t)$	Survival probability for a life aged x at inception up to time t
$\mathcal{N}(\mu,\sigma^2)$	Normal density function with mean μ and variance σ^2

List of Acronyms

EMM	Efficient method of moments
fBM	Fractional Brownian motion
FFT	Fast Fourier transform
FTSE	Financial Times Stock Exchange
GARCH	Generalised Autoregressive Conditional Heteroskedasticity
gBm	Geometric Brownian motion
GMDB	Guaranteed Minimum Death Benefit
GMMB	Guaranteed Minimum Maturity Benefit
JSE	Johannesburg Stock Exchange
ODE	Ordinary differential equation
PDE	Partial differential equation
SDE	Stochastic differential equation
SVJ	Stochastic volatility jump model
SVJJ	Stochastic volatility double jump model
S&P	Standard & Poor's
VaR	Value-at-Risk

Definitions and Terms

Calibration: The process of estimating model parameters from market data.

Convergence: Approaching a limit more and more closely as the number of terms of the series increases.

Efficient method of moments (EMM): A simulation-based method of parameter estimation used when the transition density function of a model does not exist.

Fast Fourier transform (FFT): An efficient mathematical algorithm that takes time domain data and maps it to its frequency domain.

Guaranteed Minimum Death Benefit (GMDB): A life insurance contract that guarantees the beneficiary a minimum amount if the policyholder dies before the annuity begins paying benefits.

Guaranteed Minimum Maturity Benefit (GMMB): A life insurance contract that guarantees the policyholder a minimum amount at maturity of the contract, protecting the policyholder from poor equity performance.

Market-consistent: Matching the market prices of tradable assets as closely as possible.

Monte Carlo simulation: A mathematical technique that simulates the range of possible outcomes for an uncertain event.

Rainbow option: An option that depends on two or more underlying assets. An example of a rainbow option is the "worst-of-2" call option that pays out based on the worst performing asset.

Replicating portfolio: A pool of assets designed to reproduce the market values of a pool of liabilities across a large number of stochastic scenarios.

Spread option: An option that depends on the difference (spread) between two underlying assets.

Static hedging: A strategy that seeks to replicate the value of an option without the need to constantly rebalance the portfolio.

Volatility targeting: A trading strategy that aims to keep the volatility of a portfolio constant by increasing (decreasing) equity exposure in times of low (high) volatility.

List of Figures

2.1 Historical FTSE/JSE Top40 closing prices . . . 23
2.2 Historical 3-month T-Bill rates in South Africa . . . 23
2.3 Survival probability for 50-year old male and female pensioner . . . 24
2.4 Log mortality rates for UK males . . . 25
2.5 Log mortality rates for UK females . . . 25
2.6 Change in log mortality rates for UK males . . . 26
2.7 Change in log mortality rates for UK females . . . 26
2.8 FTSE/JSE Top40 volatility surface on 16 November 2022 . . . 27
2.9 Heston-Hull-White model fit to market data on 16 November 2020 . . . 30
2.10 Heston-Hull-White volatility term-structure on 16 November 2020 . . . 31
2.11 Single sample path for each state variable . . . 32
2.12 CIR++ simulation for force of mortality . . . 33
2.13 Stochastic survival probability curves . . . 33
2.14 Mean survival probability curve versus input survival probability curve . . . 34
2.15 Calibrated AR(1)-ARCH(1) parameters for UK females . . . 35
2.16 AR(1)-ARCH(1) forecast for 85-year old UK female . . . 35
2.17 Standardised residuals versus standard normal density . . . 36
2.18 GMMB price for a 50-year old South African pensioner . . . 38
2.19 GMDB price for a 50-year old South African pensioner . . . 39
2.20 Calibrated parameters over time . . . 41
2.21 GMMB liability over time . . . 42
2.22 GMMB Δ over time . . . 42
2.23 Hedge performance over time . . . 43

3.1 Two-factor gBm FFT Price as a function of truncation width $\bar{u}$ with $N = 512$ . . . 56
3.2 FFT convergence under the three-factor stochastic volatility model . . . 58

4.1 Monte Carlo convergence using the parameters in Table 4.1 . . . 72
4.2 Impact of stochastic interest rates with $\eta = 0.05$, $\rho_{x_1,r} = 0.75$, and $\rho_{x_2,r} = 0.6$. . . . 73
4.3 Impact of stochastic interest rates with $\eta = 0.05$, $\rho_{x_1,r} = -0.75$, and $\rho_{x_2,r} = -0.6$. 74

5.1 Effect of σ_v on returns . . . 78

5.2 Effect of $\rho_{x,v}$ on returns . . . 79
5.3 Effect of λ on returns . . . 80
5.4 Effect of μ_S on returns . . . 81
5.5 Effect of σ_S on returns . . . 81
5.6 S&P500 daily historical closing prices . . . 84
5.7 S&P500 log returns . . . 85
5.8 S&P500 and normal density . . . 85
5.9 Comparison of densities under the Heston stochastic volatility model for S&P500 87
5.10 Comparison of Heston and S&P500 densities . . . 87
5.11 Comparison of densities under the Bates SVJ model for S&P500 . . . 90
5.12 Comparison of Bates and S&P500 densities . . . 91
5.13 FTSE/JSE Top40 historical closing prices . . . 92
5.14 FTSE/JSE Top40 log returns . . . 92
5.15 FTSE/JSE Top40 and normal density . . . 93
5.16 Comparison of densities under the Heston model for FTSE/JSE Top40 . . . 94
5.17 Comparison of densities under the Bates SVJ model for FTSE/JSE Top40 . . . 95
5.18 Bates SVJ equity forecast . . . 97
5.19 Weekly volatility forecast . . . 98
5.20 Equity weight based on 10% volatility target . . . 98
5.21 Cash weight based on 10% volatility target . . . 99
5.22 Portfolio value through time . . . 99
5.23 1-year returns . . . 100
5.24 1-year volatility . . . 100
5.25 3-year returns . . . 101
5.26 3-year volatility . . . 101
5.27 5-year returns . . . 102
5.28 5-year volatility . . . 102

6.1 FTSE/JSE Top40 closing prices . . . 110
6.2 SVJJ model versus FTSE/JSE Top40 distribution . . . 112
6.3 SVJJ fit to FTSE/JSE Top40 implied volatility surface . . . 113
6.4 Real-world FTSE/JSE Top40 distribution at $t = 0.25$ . . . 116
6.5 Replicating option quantities based on Choie and Novometsky optimisation . . . 117
6.6 Portfolio versus target option based on Choie and Novometsky optimisation . . . 118
6.7 Replicating option quantities based on Armstrong et al. optimisation . . . 118
6.8 Portfolio versus target option based on Armstrong et al. optimisation . . . 119
6.9 Real-world distributions for S_1 - S_2 at $t = 0.25$ . . . 121
6.10 Relationship between FTSE/JSE Top40 and $S_1 - S_2$ at $t = 0.25$ . . . 122
6.11 Replicating option quantities based on Choie and Novometsky optimisation . . . 123

6.12 Portfolio versus target option based on Choie and Novometsky optimisation . . . 124
6.13 Replicating option quantities based on Armstrong et al. optimisation 125
6.14 Portfolio versus target option based on Armstrong et al. optimisation 125

List of Tables

2.1 Calibration details on 16 November 2020 29
2.2 CIR++ parameters for 50-year old male and female pensioners 32
2.3 Base parameters on 16/11/2020 37

3.1 FFT prices under two-factor gBm with $S_1(0) = 100,\ S_2(0) = 96,\ \delta_1 = 0.05,\ \delta_2 = 0.05,\ r = 0.1,\ \sigma_1 = 0.1,\ \sigma_2 = 0.2,\ \rho_{x_1,x_2} = 0.5,\ \epsilon_1 = -3,\ \epsilon_2 = -1,\ \alpha_1 = 0.75,\ \alpha_2 = 0.75,\ T = 1$. 55
3.2 Absolute difference between FFT and Stulz prices 55
3.3 MC and FFT prices under three-factor stochastic volatility with $S_1(0) = 100,\ S_2(0) = 96,\ \delta_1 = 0.05,\ \delta_2 = 0.05,\ r = 0.1,\ \sigma_1 = 1,\ \sigma_2 = 0.5,\ v(0) = 0.04,\ \kappa = 1,\ \bar{v} = 0.04,\ \sigma_v = 0.05,\ \rho_{x_1,x_2} = 0.5,\ \rho_{x_1,v} = -0.5,\ \rho_{x_2,v} = 0.25,\ \epsilon_1 = -3,\ \epsilon_2 = -1,\ \alpha_1 = 0.75,\ \alpha_2 = 0.75,\ T = 1,\ N = 512,\ \bar{u} = 100$. 58

4.1 $S_1(0) = 100, S_2(0) = 96, \delta_1 = 0.05, \delta_2 = 0.05, v(0) = 0.04, r(0) = 0.1, \sigma_1 = 1.0, \sigma_2 = 0.5, \kappa = 1, \bar{v} = 0.04, \sigma_v = 0.05, \lambda = 1, \theta = 0.1, \eta = 0, \rho_{x_1,x_2} = 0.5, \rho_{x_1,v} = -0.5, \rho_{x_1,r} = 0, \rho_{x_2,v} = 0.25, \rho_{x_2,r} = 0, \rho_{v,r} = 0, N = 256, \bar{u} = 40, \epsilon_1 = -3, \epsilon_2 = 1, T = 1$. 71
4.2 Convergence of FFT using the parameters in Table 4.1 72

5.1 Comparison of annualised Heston parameters for S&P500 86
5.2 Heston model daily statistics for the S&P500 88
5.3 S&P500 daily statistics over different periods 88
5.4 Calibrated Heston model parameters (annualised) for S&P500 89
5.5 Comparison of annualised Bates SVJ parameters for S&P500 90
5.6 Bates model daily statistics for the S&P500 91
5.7 Annualised Heston parameters for FTSE/JSE Top40 93
5.8 Heston model daily statistics for the FTSE/JSE Top40 94
5.9 Annualised Bates parameters for FTSE/JSE Top40 95
5.10 Bates model daily statistics for the FTSE/JSE Top40 96
5.11 Statistics for volatility targeting strategies over one year 100
5.12 Statistics for volatility targeting strategies over three years 101
5.13 Statistics for volatility targeting strategies over five years 102

6.1 SVJJ $\mathbb{P}$-parameters for FTSE/JSE Top40 111
6.2 SVJJ model daily statistics for the FTSE/JSE Top40 111
6.3 SVJJ $\mathbb{Q}$-parameters for FTSE/JSE Top40 113
6.4 MC and FFT European call option prices under SVJJ model with $S(0) = 100,\ r = 0.1,\ v(0) = 0.04,\ \beta = 1,\ \alpha = 0.04,\ \sigma_v = 0.05,\ \rho_{x,v} = -0.5,\ \lambda = 5,\ \mu_S = 0,\ \sigma_S = 0.01,\ \rho_J = -0.3,\ \mu_V = 0.02,\ N = 256,\ \bar{u} = 40,\ \epsilon_1 = -3,\ \epsilon_2 = 1,\ T = 1$. 114
6.5 Market information for European call option on 16 November 2020 115
6.6 Replicating option quantity table based on Choie and Novometsky optimisation . 117
6.7 SVJJ $\mathbb{P}$-parameters for S_1 and S_2 121
6.8 Market information for European spread call option on 16 November 2020 122
6.9 Replicating option quantity table based on Choie and Novometsky optimisation . 123

1 Introduction

> *The true investor welcomes volatility.*
>
> **— Warren Buffet**

1.1 Background

Modern financial markets have become increasingly sophisticated over the past couple of decades. Technological advancements and globalisation have made it possible to trade complex financial instruments at the touch of a button. However, one must not forget to appreciate the mathematical theory that underpins the prices of these instruments, first introduced by Black and Scholes (1973) and Merton (1973), and later extended by the likes of Heston (1993) and Bates (1996).

Nowadays, the flaws in the Black and Scholes (1973) model are well understood. For one, returns do not typically follow a normal distribution with constant volatility. Using the Black and Scholes (1973) model to price instruments with exotic-style payoffs is, therefore, simply incorrect.

The Heston (1993) stochastic volatility model has become a model of choice for the pricing of exotic equity options. The beauty of this model lies in its ability to produce skewed and fat-tailed distributions, which aligns with the behaviour of financial markets. Furthermore, the Heston (1993) model has a semi-closed form solution for European call and put options, which enables efficient calibration of model parameters.

The Bates (1996) stochastic volatility jump (SVJ) model is an extension of the Heston (1993) model that adds random jumps to the underlying equity dynamics. The advantage of the Bates (1996) SVJ model is its ability to produce even more pronounced skew and fat-tailed distributions than the Heston (1993) model.

Option pricing, based on the Heston (1993) and Bates (1996) SVJ models, has been a topic of extensive research since the models were first introduced. Many variants of these models now exist, for example, the Heston-Hull-White model introduced by Grzelak and Oosterlee (2011), where the equity component is described by the Heston (1993) model and the interest rate

component by the Hull and White (1990) model. More recently, the so-called "rough" Heston model was introduced by El Euch et al. (2019), which builds upon the concept of fractional Brownian motion (fBm) (see, Mandelbrot and van Ness, 1968).

The works of Heston (1993) and Bates (1996) should be celebrated due to the countless opportunities they have created in modern financial markets. Other notable contributions worth mentioning are the works of Carr and Madan (1999) and their application of the fast Fourier transform (FFT) to option pricing, and Duffie et al. (2000) on the application of Fourier transform methods and affine models.

The reason that the contributions made by Carr and Madan (1999) and Duffie et al. (2000) are so significant is due to the fact that many exotic options do not have closed-form solutions under stochastic volatility dynamics. Duffie et al. (2000) introduced a general framework that allows one to write the characteristic function of a stochastic volatility model in explicit form, which can then be used in the FFT algorithm of Carr and Madan (1999) to obtain the option value.

The vast majority of literature on stochastic volatility models, including the works by Heston (1993) and Bates (1996), tends to focus almost exclusively on the pricing of options in developed economies under the so-called risk-neutral probability measure. The application of stochastic volatility models under the real-world probability measure, and in developing economies, is often neglected.

Andersen et al. (2002) used a statistical inference technique called the efficient method of moments (EMM) to calibrate stochastic volatility models to historical equity returns rather than option prices, which opens doors to new opportunities.

South Africa can be seen as a developing market compared to major economies such as the United States and United Kingdom. This often means that simplified models are used to value complex financial instruments, simply due to a lack of market data. Consequently, there is a limited body of research on stochastic volatility models applied to the South African market.

Although stochastic volatility models have been around for a long time, there is significant potential to increase their scope from their initial intended use being single-asset option pricing. For example, stochastic volatility models can be used by insurance companies, through so-called economic scenario generators, to value options embedded in life insurance contracts. Furthermore, stochastic volatility models can been used to value multi-asset options and forecast future equity performance.

As financial markets move towards artificial intelligence (AI) and machine learning, one must not forget the foundation on which they are built. As the philosopher Matshona Dhliwayo said: *"A tree's beauty lies in its branches, but its strength lies in its roots."*

1.2 Research Objectives

The central theme of this thesis is on the application of stochastic volatility models and numerical methods under different probability measures.

Contingent claims are typically priced under the risk-neutral $\mathbb{Q}$-probability measure due to the principle of no arbitrage. The real-world $\mathbb{P}$-probability measure is generally used to forecast a random variable such as an equity price, but receives considerably less attention in the literature. Therefore, in this thesis, our ultimate goal is to link stochastic volatility models under the $\mathbb{P}$- and $\mathbb{Q}$-measures, which will open doors to new opportunities.

To achieve this, we consider five practical problems, where each problem is presented as a chapter. The first three chapters focus solely on the application of stochastic volatility models and numerical methods in the risk-neutral $\mathbb{Q}$-measure. The fourth chapter deals with stochastic volatility models and numerical methods in the real-world $\mathbb{P}$-measure. The fifth and final chapter links the ideas from the $\mathbb{P}$- and $\mathbb{Q}$-measures by applying a stochastic volatility model to a portfolio risk management problem.

The research questions addressed in this thesis are the following:

1. **What are the important risk factors to consider when pricing long-dated options**?

   This question is addressed in Chapters 2 and 4. In Chapter 2, Section 2.7.1, the prices of long-dated insurance liabilities such as Guaranteed Minimum Maturity Benefits (GMMBs) and Guaranteed Minimum Death Benefits (GMDBs) are compared under different model assumptions, i.e., constant versus stochastic volatility and interest rates. In Chapter 4, Section 4.6.3, the prices of long-dated spread options are compared assuming constant versus stochastic interest rates.

2. **Does the FFT lead to significant time saving, but similar accuracy, compared to Monte Carlo simulation for numerical evaluation of option values?**

   This question is addressed in Chapters 3 and 4. In Chapter 3, Section 3.5.3, the convergence of the FFT under the three-factor stochastic volatility model for "worst-of-2" call options is discussed. In Chapter 4, Section 4.6.2, the convergence and execution time of the FFT under the two-asset Heston-Hull-White model is compared with the results obtained from a Monte Carlo simulation for spread options.

3. **Are jumps an important risk factor to consider when modelling equity returns?**

   This question is addressed in Chapters 5 and 6. In Chapter 5, Sections 5.3.1 and 5.3.2, the Heston (1993) and Bates (1996) stochastic volatility jump (SVJ) models are calibrated to historical S&P500 and FTSE/JSE Top40 returns. The distribution generated by the models

is then compared with the historical returns distribution. In Chapter 6, Section 6.4.1, the stochastic volatility double jump (SVJJ) model of Duffie et al. (2000) is calibrated to historical FTSE/JSE Top40 returns. The model distribution and historical returns distribution are then compared.

4. **Is volatility targeting an effective trading strategy?**

   Volatility targeting is discussed in Chapter 5, Section 5.4. We discuss the general concept of volatility targeting and perform simulations to test the outcome of various volatility targeting strategies over different investment horizons. Returns and volatility are assessed for 10%, 15%, and 20% volatility targets over 1, 3, and 5 years. The results are compared with an equity-only holding strategy.

5. **Can we hedge long-dated options with a combination of other derivatives and cash?**

   This question is addressed in Chapter 6 and links the ideas from the previous chapters. Hedging is a portfolio risk management problem, and, therefore, falls within the realm of the real-world $\mathbb{P}$-probability measure. On the other hand, derivative pricing is linked to the risk-neutral $\mathbb{Q}$-probability measure. Therefore, in Sections 6.4.1 and 6.4.2, we calibrate the SVJJ model of Duffie et al. (2000) to historical FTSE/JSE Top40 returns and the FTSE/JSE Top40 implied volatility surface, which enables us to forecast FTSE/JSE Top40 prices ($\mathbb{P}$-measure), and price options ($\mathbb{Q}$-measure). Two static hedging programs are then applied to hedge a long-dated vanilla European call option and European spread call option in Sections 6.4.3 and 6.4.4.

These questions stem from problems in the financial services industry where stochastic volatility is often observed. These are: 1) pricing derivatives embedded in life insurance contracts, *i.e.*, long-dated options; 2) pricing multi-asset options, *i.e.*, rainbow and spread options; and 3) portfolio risk management, *i.e.*, volatility targeting and hedging.

This thesis brings together many disciplines including probability theory, quantitative finance, numerical analysis, and mathematical statistics.

1.3 Thesis Structure

This thesis consists of five core chapters that attempt to address the research questions posed in Section 1.2. Chapters 2, 3, 4, and 5 have been published in peer-reviewed journals. Chapter 6 has been accepted for presentation at an international conference. The flow diagram below outlines the structure and key topics of this thesis.

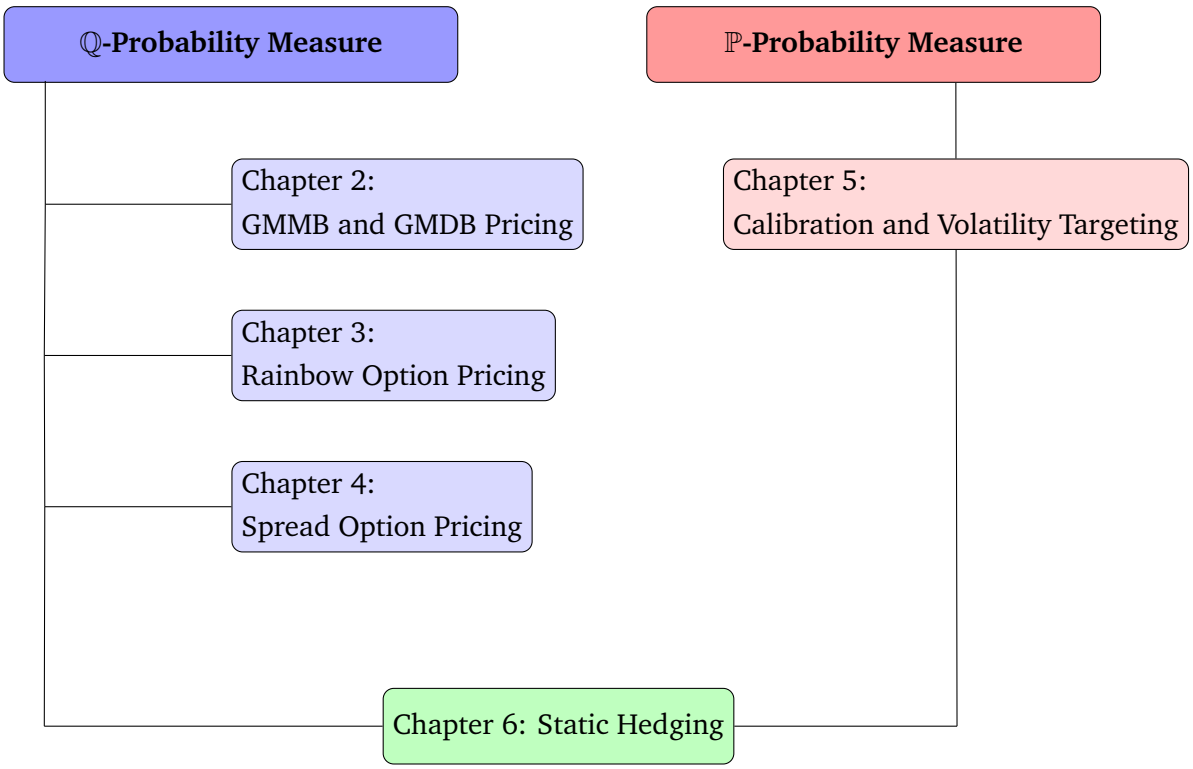

- **Chapter 2: GMMB and GMDB Pricing**

  Chapter 2 considers the application of stochastic volatility models in insurance and is titled "An Economic Scenario Generator for Embedded Derivatives in South Africa" (Levendis & Maré, 2022a). In this chapter, we consider the pricing GMMBs and GMDBs in the South African life insurance industry. The Heston-Hull-White model of Grzelak and Oosterlee (2011) is proposed to price these products, and a stochastic mortality component based on the CIR++ model of Brigo and Mercurio (2001) is added to the model dynamics. We show how to calibrate the Heston-Hull-White model to the FTSE/JSE Top40 implied volatility surface using the Fourier transform technique of Carr and Madan (1999). Furthermore, we calibrate the CIR++ model to survival probability curves for South African males and females. We then test the impact of the various model parameters, including stochastic volatility and stochastic interest rates, on the GMMB and GMDB prices. Finally, we show the performance of a delta-hedging strategy for the GMMB using the Heston-Hull-White model with stochastic mortality.

- **Chapter 3: Rainbow Option Pricing**

  Chapter 3 is titled "Pricing Two-Asset Rainbow Options with the Fast Fourier Transform" (Levendis & Maré, 2023b) and deals with the pricing of rainbow options, specifically "worst-of-2" call options. We consider the two-factor geometric Brownian motion (gBm)

model of Margrabe (1978) and the three-factor stochastic volatility model of Dempster and Hong (2002) for the underlying equity dynamics. The "worst-of-2" call options are priced with the two-dimensional FFT of Hurd and Zhou (2010). The prices and efficiency of the method are compared with the results obtained from a Monte Carlo simulation.

- **Chapter 4: Spread Option Pricing**

  Chapter 4 titled "Efficient Pricing of Spread Options with Stochastic Rates and Stochastic Volatility" (Levendis & Maré, 2022b) focuses on the derivation of the two-asset Heston-Hull-White model - an extension of the single-asset model derived by Grzelak and Oosterlee (2011). The derivation is presented from first principles, and the model is used to price European spread options with the two-dimensional FFT of Hurd and Zhou (2010). The accuracy and efficiency of the FFT are assessed, and the impact of stochastic interest rates on long-dated spread options is tested.

- **Chapter 5: Calibration and Volatility Targeting**

  Chapter 5 titled "Application of Stochastic Volatility Models in the Real-World Measure" (Levendis & Maré, 2023a) considers the calibration of the Heston (1993) and Bates (1996) SVJ models to historical equity returns. Firstly, we aim to replicate the results in Andersen et al. (2002) for the S&P500 using the EMM technique of Gallant and Tauchan (1996). Thereafter, we implement the Heston (1993) and Bates (1996) SVJ models to historical FTSE/JSE Top40 returns to test which model is better at capturing the characteristics of the South African equity market. Finally, we show how a real-world stochastic volatility model can be used to test the performance of a volatility targeting trading strategy.

- **Chapter 6: Static Hedging**

  Chapter 6 focuses on the interplay between the $\mathbb{P}$- and $\mathbb{Q}$-probability measures and is titled "Static Hedging of Vanilla and Exotic Options in a South African Context" (Levendis & Maré, 2023c). We calibrate the SVJJ model of Duffie et al. (2000) to historical FTSE/JSE Top40 returns ($\mathbb{P}$-measure) as well as the FTSE/JSE Top40 implied volatility surface ($\mathbb{Q}$-measure). We then propose a simulation-based framework based on the $\mathbb{P}$- and $\mathbb{Q}$-SVJJ models that can be used to test the performance of a static hedging program, where the written option is either a long-dated vanilla call option or exotic spread call option. We use the optimisation routines of Choie and Novomestky (1989) and Armstrong et al. (2018) to solve for the exchange traded option quantities and cash balance in the hedging portfolio and test the hedging performance at a future date.

Chapter 7 concludes the thesis.

2 An Economic Scenario Generator for Embedded Derivatives in South Africa

> *Volatility is not synonymous of risk but – for those who truly understand it – of wealth.*
>
> **— Francois Rochon**

Keywords: Heston-Hull-White model · stochastic volatility · stochastic interest rates · stochastic mortality · GMMB · GMDB · pricing · hedging

2.1 Introduction

Life insurers often sell policies with embedded guarantees. Pricing these guarantees is challenging due to the financial and insurance risks involved. Two popular products sold by life insurers include the Guaranteed Minimum Maturity Benefit (GMMB) and the Guaranteed Minimum Death Benefit (GMDB). The former provides the policyholder with the greater of the guaranteed amount at maturity or the fund value in which the initial premium was invested, whereas the latter provides the policyholder with the greater of the guaranteed amount or the fund value at the time of death, should death occur before the maturity of the contract. To price these products, a model is required that can capture financial risks, including equity and interest rate volatility, and insurance risks such as mortality. In this chapter[1], we propose such a model.

The Black and Scholes (1973) model is unsuitable for modelling long-dated contingent claims as it does not account for stochastic interest rates or stochastic volatility. Furthermore, it has been shown that the model is inconsistent with the stylised facts of financial returns (see Cont, 2001). Fortunately, the Black and Scholes (1973) model has been extended to incorporate stochastic volatility (see, for example, Heston, 1993 and Bates, 1996).

The Heston (1993) model incorporates stochastic volatility, and, hence, can capture the equity skew observed in financial markets. However, as shown by Kammeyer and Kienitz (2012),

[1]This chapter is based on a paper (Levendis and Maré, 2022a) published in the *South African Actuarial Journal.*

stochastic interest rates are a bigger risk driver for long-dated contingent claims than stochastic volatility. Consequently, a model is required that can incorporate stochastic volatility as well as stochastic interest rates.

The Hull and White (1990) model is a stochastic process that describes the short-term instantaneous interest rate and is popular among practitioners. The model has the benefit in that it can fit the initial term-structure of interest rates and has an analytical solution for zero-coupon bonds. Given the model's analytical tractability, researchers in the field of quantitative finance started to investigate its use with the Heston (1993) stochastic volatility model, leading to what is now known as the Heston-Hull-White model.

A particular challenge faced when pricing GMMB and GMDB products is their long-dated nature. As discussed by Maze (2014), the further a call or put option is from maturity, the more time a stochastic interest rate has to reach its long-run mean and affect the option price. Therefore, it is important for practitioners to include stochastic interest rates when reserving for embedded derivatives. The Heston-Hull-White model will allow for this. Apart from financial risks, insurance risks such as mortality must also be incorporated.

Traditional actuarial modelling assumes deterministic mortality rates based on actuarial life tables. In reality, mortality rates are stochastic and it may be necessary to price contingent claims based on this assumption. In this chapter, we extend the Heston-Hull-White model to include stochastic mortality rates based on either a continuous-time Cox et al. (1985) process or discrete-time AR(1)-ARCH(1)[2] model when pricing products contingent on either survival or death. We call this extension the Heston-Hull-White-Mortality model.

The remainder of this chapter is structured as follows: Section 2.2 reviews relevant literature on guaranteed maturity benefits, the Heston-Hull-White model, and stochastic mortality. Section 2.3 introduces the mathematical theory of the Heston-Hull-White model. Section 2.4 reviews the continuous-time Cox et al. (1985) short rate process and its extension proposed by Brigo and Mercurio (2001), and the discrete time AR(1)-ARCH(1) model that will be used for mortality rates. Section 2.5 shows the formulas that will be used to price the GMMB and GMDB products. Section 2.6 introduces the data that will be used to calibrate the Heston-Hull-White-Mortality model and price the GMMB and GMDB products. Section 2.7 shows our pricing and hedging results, and Section 2.8 concludes the findings.

[2]First-order Autoregressive and Autoregressive Conditional Heteroskedastic model

2.2 Literature Review

In 2007, ING, a large Dutch banking group, presented a problem from the financial industry where the goal was to derive a closed-form or semi-closed form solution for a call option under the hybrid Heston-Hull-White model. To understand the complexity of this challenge, Oosterlee (2007) writes: "*Quite a few mathematicians took up this ING challenge, and during the week three subgroups were formed, each approaching the problem from a different side. A particular challenge here was that some Dutch professors in financial mathematics in earlier attempts were not able to come up with a closed form option pricing solution for this particular model.*" From this challenge, a couple of interesting approaches were suggested to solve the model, see in 't Hout et al. (2007) and Fang and Janssens (2007).

The Heston-Hull-White model has become popular since the so-called ING challenge. Kammeyer and Kienitz (2009) considered an interest rate process that is independent of the equity and volatility processes, and showed the resulting characteristic function and calibration of the model based on the Carr and Madan (1999) fast Fourier transform (FFT) method. Still, the challenge remained to incorporate a non-zero correlation structure between the equity, volatility and interest rate processes.

Grzelak and Oosterlee (2011) made a breakthrough by deriving approximations for specific terms in the Heston-Hull-White model that allow for a non-zero correlation structure between the equity, volatility, and interest rate processes. Their technique, known as *linearisation*, involves approximating the nonlinear terms in the model's covariance matrix with a linear function. The resulting model is then in affine form, which means that the Fourier methods of Duffie et al. (2000) and Carr and Madan (1999) can be applied. Grzelak and Oosterlee (2011) concluded that the approximations yield accurate prices for European options. This was a big contribution to the field of quantitative finance and the model has since been applied to price long-dated contingent claims.

Maze (2014) applied the Heston-Hull-White model to price long-dated European call options and compared the results to the standard Black and Scholes (1973) model. The author showed that call option prices under the Heston-Hull-White model increase compared to the Black and Scholes (1973) model as the time to maturity increases. This is an important finding, especially considering the long-dated nature of embedded derivatives such as GMMB and GMDB products.

Patel (2019) implemented the Heston-Hull-White model to price long-dated European call options and compared the deterministic and stochastic approximations for the Heston-Hull-White model proposed by Grzelak and Oosterlee (2011). The author concludes that the

deterministic approximation is a feasible way to calibrate the Heston-Hull-White model, whereas the stochastic approximation is inefficient.

As mentioned earlier, traditional actuarial modelling assumes deterministic mortality rates. In reality, mortality rates tend to behave randomly. Cairns et al. (2008) discuss various stochastic mortality models including the discrete-time Lee and Carter (1992) model and continuous-time short rate models. Mortality modelling is similar to interest rate modelling, hence, it is relatively simple to extend the theory of interest rates to mortality.

A weakness of many short rate models is that they cannot fit the initial term-structure of interest rates. Examples include the Vasicek (1977) and Cox et al. (1985) models. The Hull and White (1990) model overcomes this issue, but, due to its Gaussian distribution, can produce negative values. Brigo and Mercurio (2001) extended the Cox et al. (1985) model to include a deterministic shift that allows the model to fit the initial term-structure of interest rates whilst maintaining its analytical tractability for zero-coupon bonds. The model is termed the CIR++ model.

Truter (2012) applied the CIR++ model to price defaultable zero-coupon bonds. The author showed that the hazard rate, or instantaneous probability of default, can be modelled as a CIR++ process. From there, an analytical solution exists to calculate survival probabilities (the probability that a counterpart does not default). Simply replacing the hazard rate with the force of mortality (instantaneous rate of mortality), the CIR++ model can be used to calculate survival probabilities in a life insurance context.

Literature on the use of hybrid models to price embedded derivatives or variable annuities is limited. Wang (2011) applied the Heston-Hull-White model to price the Guaranteed Minimum Withdrawal Benefit (GMWB) product and discussed two semi-analytical techniques that can be used to calibrate the model to market data. The author considered the FFT technique of Carr and Madan (1999) and the Fourier-Cosine (COS) technique introduced by Fang and Oosterlee (2008). The author concludes that the Heston-Hull-White model is an appropriate model to use when pricing long-dated claims.

Ignatieva et al. (2016) presented a framework for pricing guaranteed maturity benefits with a regime-switching and stochastic mortality model. The authors considered a continuous-time two-factor affine mortality model and assumed that the unsystematic mortality risk can be diversified away, hence, the instantaneous mortality intensity is the same under the physical and risk-neutral measures.

Balotta et al. (2020) proposed a Lévy-based hybrid model that allows for dependence between interest rates and equity prices. Furthermore, the authors divide insurance risk into two components: mortality risk and surrender risk. The authors conclude that surrender risk

is particularly important in the valuation of variable annuities due to the large number of policyholders who terminate their contracts prematurely.

Veilleux (2016) tested the impact of stochastic volatility, stochastic interest rates, and stochastic mortality on the hedge efficiency of the Guaranteed Lifetime Withdrawal Benefit (GLWB). The author considered various models including a regime-switching log-normal model to capture stochastic volatility, the Hull and White (1990) model for stochastic interest rates, and the Lee and Carter (1992) model for stochastic mortality. The author concludes that stochastic interest rates and stochastic mortality have a significant impact on the hedge efficiency of GLWB benefits.

In South Africa, the Actuarial Society of South Africa's advisory practice note, APN 110: Allowance for Embedded Investment Derivatives[3], gives guidance to actuaries on how to reserve for embedded derivatives including GMMB and GMDB products. The advisory practice note is rather open-ended and leaves the actuary with many modelling choices. An extract from APN 110 reads: "*No specific investment return projection model is prescribed. The actuary may use any market-consistent stochastic investment return projection model that he/she deems appropriate for purposes of quantifying reserves required to meet the potential cost of embedded investment derivatives.*"

The only minimum prescription set by APN 110 is that the model must be market-consistent, meaning that it must match the prices of tradable assets as closely as possible.

Limited research has been conducted in South Africa comparing different modelling approaches for embedded derivatives. Ngugi et al. (2015) applied a Variance-Gamma model to price GMMB and GMDB products wtitten on the FTSE/JSE All Share Index (ALSI) and compared this to a regime-switching model (see Hardy, 2003). The authors conclude that the Variance-Gamma model is more aligned with the stylised facts of financial returns (Cont, 2001). Furthermore, the authors suggest further research on the topic by including stochastic interest rates and stochastic mortality. This suggestion is what motivated this research.

We are not aware of literature pertaining to the Heston-Hull-White model used to price embedded derivatives in the presence of stochastic mortality. Therefore, our contribution to the literature is by pricing GMMB and GMDB products with an extension of the Heston-Hull-White model that we term the Heston-Hull-White-Mortality model. We will apply the model to price and hedge GMMB and GMDB products written on the FTSE/JSE Top40 index in South Africa.

In the next section, we introduce the mathematical theory for the Heston-Hull-White model.

[3] https://www.actuarialsociety.org.za/download/apn-110-allowance-for-embedded-investment-derivatives/

2.3 The Heston-Hull-White Model

Let $\mathbb{Q}$ denote the risk-neutral measure. Under the $\mathbb{Q}$-measure, the Heston-Hull-White model is given by the following system of stochastic differential equations (SDEs):

$$\begin{cases} dS(t) &= r(t)S(t)dt + \sqrt{v(t)}S(t)dW_s(t), & S(0) > 0, \\ dv(t) &= \kappa(\bar{v} - v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), & v(0) > 0, \\ dr(t) &= \lambda_r(\theta(t) - r(t))dt + \eta dW_r(t), & r(0) \in \mathbb{R}, \end{cases}$$

where $S(t)$ is the underlying asset price a time t, $v(t)$ is the instantaneous variance of the underlying asset at time t driven by a Cox et al. (1985) process, and $r(t)$ is the instantaneous short-term interest rate at time t driven by a Hull and White (1990) process.

The parameters of the variance process are defined as follows: κ denotes the mean reversion speed of the variance, $\bar{v}$ is the long-run mean of the variance, and σ_v is the volatility of the volatility.

The parameters of the short-term interest rate process are defined as follows: $\theta(t)$ denotes the time-dependent mean reversion level of the interest rate, λ_r denotes the mean reversion speed of the interest rate, and η is the volatility of the interest rate.

In the system of SDEs, the Brownian motions, $dW_s(t), dW_v(t)$, and $dW_r(t)$ are correlated by $\rho_{s,v}$, $\rho_{s,r}$, and $\rho_{v,r}$ respectively.

Taking $x(t) = \log S(t)$, Grzelak and Oosterlee (2011) consider the log-dynamics for the Heston-Hull-White model as it is often easier to work with this form. Applying Itô's lemma, the log-dynamics for the Heston-Hull-White model are given by the following system of SDEs under the $\mathbb{Q}$-measure:

$$\begin{cases} dx(t) &= (r(t) - \frac{1}{2}v(t))dt + \sqrt{v(t)}dW_s(t), & x(0) = \log S(0) > 0, \\ dv(t) &= \kappa(\bar{v} - v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), & v(0) > 0, \\ dr(t) &= \lambda_r(\theta(t) - r(t))dt + \eta dW_r(t), & r(0) \in \mathbb{R}. \end{cases}$$

Patel (2019) summarised the work of Grzelak and Oosterlee (2011) and showed that the Heston-Hull-White log-dynamics can be expressed in terms of independent Brownian motions. If $\mathbf{B}(t)$ is a standard 3-dimensional Brownian motion, then

$$d\mathbf{X}(t) = \mu(\mathbf{X}(t))dt + \sigma(\mathbf{X}(t))d\mathbf{B}(t), \tag{2.3.1}$$

where

$$\mathbf{X}(t) = \begin{bmatrix} x(t) \\ v(t) \\ r(t) \end{bmatrix},$$

$$\mu(\mathbf{X}(t)) = \begin{bmatrix} r(t) - \frac{1}{2}v(t) \\ \kappa(\bar{v} - v(t)) \\ \lambda_r(\theta(t) - r(t)) \end{bmatrix},$$

and

$$\sigma(\mathbf{X}(t)) = \begin{bmatrix} \sqrt{v(t)} & 0 & 0 \\ \rho_{s,v}\sigma_v\sqrt{v(t)} & \sqrt{1-\rho_{s,v}^2}\sigma_v\sqrt{v(t)} & 0 \\ \rho_{s,r}\eta & \frac{\rho_{v,r}-\rho_{s,r}\rho_{s,v}}{\sqrt{1-\rho_{s,v}^2}}\eta & \sqrt{1-\rho_{s,r}^2-\left(\frac{\rho_{v,r}-\rho_{s,r}\rho_{s,v}}{\sqrt{1-\rho_{s,v}^2}}\right)\eta} \end{bmatrix}.$$

See Patel (2019) for the proof.

The covariance matrix, $\Sigma(\mathbf{X}(t)) = \sigma(\mathbf{X}(t))\sigma(\mathbf{X}(t))^\top$, for the Heston-Hull-White model is given by:

$$\Sigma(\mathbf{X}(t)) = \begin{bmatrix} v(t) & \rho_{s,v}\sigma_v v(t) & \rho_{s,r}\eta\sqrt{v(t)} \\ * & \sigma_v^2 v(t) & \rho_{v,r}\eta\sqrt{v(t)} \\ * & * & \eta^2 \end{bmatrix}. \tag{2.3.2}$$

When the underlying asset and variance processes are correlated to the interest rate process, the Heston-Hull-White model cannot be written in an affine form. The reader is referred to Duffie et al. (2000), Grzelak and Oosterlee (2011), and Patel (2019) for more information on affine models. In short, a model is said to be affine if each element of its drift and covariance can be written as a linear function of the state variables.

From Eq. (2.3.2), it is clear that the covariance matrix contains elements that are nonlinear functions of the state variables, specifically $\sqrt{v(t)}$. In this chapter, we are only interested in the case where the interest rate process is correlated to the spot process, hence, we set $\rho_{v,r} = 0$. To make the model affine, Grzelak and Oosterlee (2011) propose a deterministic approximation by replacing $\sqrt{v(t)}$ with $\mathbb{E}[\sqrt{v(t)}]$. The covariance matrix then becomes:

$$\Sigma(\mathbf{X}(t)) = \begin{bmatrix} v(t) & \rho_{s,v}\sigma_v v(t) & \rho_{s,r}\eta\mathbb{E}[\sqrt{v(t)}] \\ * & \sigma_v^2 v(t) & 0 \\ * & * & \eta^2 \end{bmatrix}. \tag{2.3.3}$$

Grzelak and Oosterlee (2011) approximate $\mathbb{E}[\sqrt{v(t)}]$ with the following function:

$$\Lambda(t) = \sqrt{c(t)(\lambda(t)-1) + c(t)d + \frac{c(t)d}{2(d+\lambda(t))}}, \tag{2.3.4}$$

where

$$c(t) = \frac{1}{4\kappa}\sigma_v^2(1-e^{-\kappa t}), \quad d = \frac{4\kappa\bar{v}}{\sigma_v^2}, \quad \lambda(t) = \frac{4\kappa v(0)e^{-\kappa t}}{\sigma_v^2(1-e^{-\kappa t})}.$$

See Grzelak and Oosterlee (2011), and Patel (2019) for the proof.

Grzelak and Oosterlee (2011) mention that the approximation in Eq. (2.3.4) is still non-trivial and propose a further approximation for $\mathbb{E}[\sqrt{v(t)}]$ with the following function:

$$\mathbb{E}\left[\sqrt{v(t)}\right] \approx a + be^{-ct} := \tilde{\Lambda}(t), \tag{2.3.5}$$

where

$$a = \sqrt{\bar{v} - \frac{\sigma_v^2}{8\kappa}}, \quad b = \sqrt{v(0)} - a, \quad c = -\log(b^{-1}(\Lambda(1) - a)).$$

As we will show in the next section, the approximation in Eq. (2.3.5) plays an important role in the Heston-Hull-White characteristic function.

2.3.1 The Heston-Hull-White Characteristic Function

Grzelak and Oosterlee (2011) assume a constant term-structure of interest rates, $\theta(t) = \theta$, in their derivation of Heston-Hull-White characteristic function. This assumption simplifies the mathematics and we will maintain this assumption throughout the rest of this chapter. From Grzelak and Oosterlee (2011), the characteristic function is given by:

$$\phi_{HHW}(u, \mathbf{X}(t), \tau) = \exp\left(\tilde{A}(u,\tau) + B(u,\tau)x(t) + C(u,\tau)r(t) + D(u,\tau)v(t)\right), \tag{2.3.6}$$

where

$$\begin{aligned}
\tau &= T - t, \\
B(u,\tau) &= iu, \\
C(u,\tau) &= \frac{1}{\lambda_r}(iu-1)\left(1-e^{-\lambda_r\tau}\right), \\
D(u,\tau) &= \frac{1-e^{-D_1\tau}}{\sigma_v^2(1-ge^{-D_1\tau})}\left(\kappa - \sigma_v\rho_{s,v}iu - D_1\right),
\end{aligned}$$

$$A(u,\tau) = \lambda_r \theta I_1(u,\tau) + \kappa\bar{v} I_2(u,\tau) + \frac{1}{2}\eta^2 I_3(u,\tau),$$

$$\tilde{A}(u,\tau) = A(u,\tau) + \rho_{s,r}\eta I_4(u,\tau),$$

with

$$D_1 = \sqrt{(\sigma_v \rho_{s,v} iu - \kappa)^2 - \sigma_v^2 iu(iu-1)},$$

$$g = \frac{\kappa - \sigma_v \rho_{s,v} iu - D_1}{\kappa - \sigma_v \rho_{s,v} iu + D_1},$$

$$I_1(u,\tau) = \frac{1}{\lambda_r}(iu-1)\left(\tau + \frac{1}{\lambda_r}(e^{-\lambda_r \tau} - 1)\right),$$

$$I_2(u,\tau) = \frac{\tau}{\sigma_v^2}(\kappa - \sigma_v \rho_{s,v} iu - D_1) - \frac{2}{\sigma_v^2}\log\left(\frac{1 - ge^{-D_1\tau}}{1-g}\right),$$

$$I_3(u,\tau) = \frac{1}{2\lambda_r^3}(i+u)^2\left(3 + e^{-2\lambda_r\tau} - 4e^{-\lambda_r\tau} - 2\lambda_r\tau\right),$$

$$I_4(u,\tau) \approx -\frac{1}{\lambda_r}\left(iu + u^2\right)\left[\frac{b}{c}(e^{-ct} - e^{-cT} + a\tau + \frac{a}{\lambda_r}(e^{-\lambda_r\tau} - 1)\right.$$

$$\left. + \frac{b}{c - \lambda_r}e^{-cT}(1 - e^{\tau(\lambda_r - c)})\right].$$

The deterministic approximation proposed by Grzelak and Oosterlee (2011) in Eq. (2.3.5) feeds into the valuation of I_4 above, hence the approximation sign. Since the characteristic function for the Heston-Hull-White model is known, Fourier methods can be used to solve the price of a European call option as we show in the next section.

2.3.2 Heston-Hull-White Calibration

Carr and Madan (1999) showed that European call options can be priced via Fourier transform methods if the characteristic function of the model is known analytically. The authors initially considered constant interest rates in their paper. In the case of stochastic interest rates, the Fourier method must be adapted slightly.

Patel (2019) derived the price of a European call option with strike K and maturity T using the Fourier method under stochastic interest rates as:

$$V_{HHW}(k) = \frac{e^{-\alpha k}}{\pi} \int_0^\infty \Re\Big\{e^{-ivk}\psi(v)\Big\}dv, \tag{2.3.7}$$

where $k = \log(K)$, $\alpha > 0$, $i = \sqrt{-1}$, $\Re\{\cdot\}$ denotes the real part of the complex number, and

$$\psi(v) = \frac{\phi_{HHW}(v - (\alpha+1)i)}{(\alpha + 1 + iv)(\alpha + iv)}.$$

The symbol α refers to the dampening coefficient and a suitable value must be chosen in order to compute the option price in Eq. (2.3.7). As mentioned by Patel (2019), choosing α is subjective. The author set $\alpha = 0.75$ and tested the sensitivity of European call option prices to various values of α around 0.75. From the author's testing, it was concluded that $\alpha = 0.75$ is sufficient. We therefore set $\alpha = 0.75$ going forward.

As mentioned, we are not aware of any literature that combines the Heston-Hull-White model with a stochastic mortality model. GMMB and GMDB products are contingent on survival or death; this implies that mortality is an important risk driver when pricing these products. Therefore, we propose an extension of the Heston-Hull-White model by including stochastic mortality rates based on either a continuous-time Cox et al. (1985) process or discrete-time AR(1)-ARCH(1) model.

2.4 Stochastic Mortality

This section is divided into two parts. In the first subsection, we introduce the continuous-time Cox et al. (1985) model for mortality rates. In the second subsection, we present the discrete-time AR(1)-ARCH(1) model of Syuhada and Hakim (2021). The reason for considering both a continuous-time model and a discrete-time model is to allow for calibration depending on the type of mortality data that is available. A continuous-time model is generally calibrated to a survival probability curve derived from historical data or longevity swaps, whereas a discrete-time model is calibrated to a historical time series of mortality rates.

2.4.1 Cox-Ingersoll-Ross Model

Let $\mu(x+t,t)$ denote the force of mortality at time t and age $x+t$. The force of mortality represents the instantaneous probability of death at time t for a life aged $x+t$. Cairns et al. (2008) show that the probability of survival up to time t under $\mathbb{Q}$ is given by:

$$ {}_tp_x = S_x(0,t) = \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^t \mu(x+s,s)ds}\right]. \tag{2.4.1} $$

In order to calculate the expectation in Eq. (2.4.1), a suitable stochastic process must be chosen to describe the dynamics of $\mu(x+t,t)$. Cairns et al. (2008) list certain criteria for mortality rates, including the following:

- Mortality rates should be positive;
- Long-term dynamics of the model should be reasonable;
- The model should be simple to implement using analytical methods or numerical methods.

Based on the criteria listed above, we propose the Cox et al. (1985) process as a plausible model to describe mortality rates. The reason for choosing this process is the following: the model produces positive values (although care must be taken to ensure the Feller condition[4] is met), the model is mean reverting which means that long-term rates will not explode, and the model has an analytical solution for survival probabilities.

For a fixed age x at inception, the original Cox et al. (1985) model is given by the SDE:

$$ d\mu(x+t,t) = \gamma(\omega - \mu(x+t,t))dt + \xi\sqrt{\mu(x+t,t)}dW_\mu(t), \quad \mu(x,0) > 0, \tag{2.4.2} $$

where γ denotes the mean reversion speed of mortality, ω is the long-run mean of mortality rates, and ξ is the volatility of mortality rates.

A shortcoming of this model is that it cannot fit the term-structure of mortality. When the force of mortality is plotted against age, it shows an increasing trend. The original Cox et al. (1985) model does not capture this effect. Fortunately, the model has been extended to incorporate a term-structure.

Brigo and Mercurio (2001) extended the Cox et al. (1985) model to incorporate a term-structure and created the so-called CIR++ model. The model takes the following form where $\varphi_x^{CIR}(t)$ is a deterministic function of time:

$$ \mu(x+t,t) = \varphi_x^{CIR}(t) + X(x+t,t), \tag{2.4.3} $$

[4] The Feller condition is met if $2\gamma\omega > \xi^2$, where γ is the mean reversion speed of mortality, ω is the long-run mean of mortality, and ξ is the volatility of mortality.

where

$$dX(x+t,t) = \gamma(\omega - X(x+t,t))dt + \xi\sqrt{X(x+t,t)}dW_\mu(t), \qquad X(x,0) > 0,$$

$$\varphi_x^{CIR}(t) = f_x^M(0,t) - f_x^{CIR}(0,t).$$

Furthermore, $f_x^M(0,t)$ denotes the market forward mortality curve observed from the initial term-structure of mortality (survival curve), and

$$f_x^{CIR}(0,t) = 2\gamma\omega\frac{e^{th}-1}{2h+(\gamma+h)(e^{th}-1)} + X(x,0)\frac{4h^2e^{th}}{[2h+(\gamma+h)(e^{th}-1)]^2},$$

with $h = \sqrt{\gamma^2+2\xi^2}$.

The survival probability from time t up to time T is then given by:

$$S_x(t,T) = \frac{S_x^M(0,T)L_x(0,t)e^{-M_x(0,t)X(x,0)}}{S_x^M(0,t)L_x(0,T)e^{-M_x(0,T)X(x,0)}}L_x(t,T)e^{-M_x(t,T)[\mu(x+t,t)-\varphi_x^{CIR}(t)]}, \tag{2.4.4}$$

where

$$L_x(t,T) = \left[\frac{2he^{(\gamma+h)(T-t)/2}}{2h+(\gamma+h)(e^{(T-t)h}-1)}\right]^{2\gamma\omega/\xi^2},$$

$$M_x(t,T) = \frac{2(e^{(T-t)h}-1)}{2h+(\gamma+h)(e^{(T-t)}h-1)},$$

and $S_x^M(0,\cdot)$ denotes the market observed survival probability curve at $t=0$ for a life aged x at inception.

A major advantage of term-structure models, and specifically the CIR++ model, is that they fit the input survival probability curve at $t=0$, *i.e*, $S_x(0,\cdot)$.

Throughout the rest of this chapter, we make the same assumption as Ignatieva et al. (2016) in that the life insurer has an adequately large number of policyholders so that unsystematic mortality risk can be diversified away. This implies that the force of mortality under the physical probability measure $\mathbb{P}$ and risk-neutral probability measure $\mathbb{Q}$ are equal:

$$\mu^{\mathbb{P}}(x+t,t) = \mu^{\mathbb{Q}}(x+t,t).$$

Life insurers in South Africa typically use life tables published by the Actuarial Society of South Africa to calculate survival probabilities, since South Africa has no traded market for mortality derivatives. Our assumption that the instantaneous force of mortality is the same under the $\mathbb{P}$ and $\mathbb{Q}$ measures allows us to calibrate the CIR++ model directly to a survival probability curve derived from the life tables.

In the next subsection, we present the discrete-time mortality model of Syuhada and Hakim (2021).

2.4.2 AR(1)-ARCH(1) Model

Syuhada and Hakim (2021) adopted the AR(1)-ARCH(1) model of Giacometti et al. (2012) and Lin et al. (2015) and applied the model to yearly changes in the log mortality rate. Syuhada and Hakim (2021) explain that the conditional mean of the change in log mortality rate is modelled by a first-order Autoregressive AR(1) model, and the conditional variance is modelled by a first-order Autoregressive Conditional Heteroscedastic ARCH(1) stochastic volatility model to capture volatility clustering and the time-varying nature of volatility.

Let $\mu(x,t)$ denote the force of mortality at age x in year t calculated as:

$$\mu(x,t) = \frac{D(x,t)}{E(x,t)}, \tag{2.4.5}$$

where $D(x,t)$ denotes the number of individuals aged x in year t who die before year $t+1$ or age $x+1$, and $E(x,t)$ denotes the number of individuals exposed to death risk at age x in year t.

Let $Y_x(t)$ denote the change in the log mortality rate for each fixed age $x \in \{x_0, x_0+1, x_0+2, ..., X_{max}\}$, that is:

$$Y_x(t) = \log \mu(x,t) - \log \mu(x,t-1) = \log \frac{\mu(x,t)}{\mu(x,t-1)}. \tag{2.4.6}$$

where $t \in \{t_0+1, t_0+2, ..., T_{max}\}$, with X_{max} and T_{max} denoting the maximum age and year in the historical data respectively.

Syuhada and Hakim (2021) model $Y_x(t)$ according to an AR(1)-ARCH(1) model:

$$Y_x(t) = a_x + b_x Y_x(t-1) + \epsilon_x(t), \tag{2.4.7}$$

where

$$\epsilon_x(t) = \sqrt{h_x(t)} z_x(t),$$

$$h_x(t) = \omega_x + \gamma_x \epsilon_x^2(t-1),$$

with $h_x(t)$ denoting the conditional variance, $a_x \in \mathbb{R}$, $b_x \in (-1,1)$, $\omega_x \in (0,\infty)$, $\gamma_x \in [0,1)$, and $\epsilon_x(t)$ is an innovation parameter with $z_x(t) \sim N(0,1)$.

For each fixed age x, the model can be calibrated to a time series of historical mortality rates via maximum likelihood estimation.

Stochastic mortality has become a highly relevant topic due to the COVID-19 pandemic, with South African life insurers reporting absurd mortality rates. See, for example, the report by the Association for Savings and Investment South Africa[5] (ASISA).

This concludes the section on stochastic mortality models. In the next section, we introduce the mathematics of the GMMB and GMDB products.

2.5 Guaranteed Minimum Maturity and Death Benefits

From now on, we will refer to the Heston-Hull-White-Mortality model when combining the Heston-Hull-White model with either the Cox et al. (1985) or AR(1)-ARCH(1) models for mortality. We assume that the force of mortality in the Heston-Hull-White-Mortality model evolves independently from the processes driving $S(t)$, $v(t)$, and $r(t)$. This implies that the Heston-Hull-White model and CIR++/AR(1)-ARCH(1) stochastic mortality model can be calibrated independently. Furthermore, an independent mortality process will simplify the calculations for the GMMB and GMDB products going forward.

We make extensive use of the following relationship when deriving pricing formulae for the GMMB and GMDB products: Given two independent random variables Y and Z, the expected value of their product is equal to the product of their expected values:

$$\mathbb{E}[YZ] = \mathbb{E}[Y]\mathbb{E}[Z]. \tag{2.5.1}$$

We use the same notation as in the textbook by Feng (2018) when presenting the pricing formulae for the GMMB and GMDB products, adjusting the formulae slightly to account for stochastic interest rates and stochastic mortality.

The price of a GMMB at $t = 0$ from a life insurer's perspective is given by the formula:

$$V_{GMMB}(0) = \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^T r(s)ds}(G - F(T))^+ e^{-\int_0^T \mu(x+s,s)ds}\right]$$

[5] https://www.asisa.org.za/media-releases/life-insurers-report-a-surge-in-death-claims-during-covid-19-third-wave/

$$
\begin{aligned}
&- \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} m_e F(s) e^{-\int_0^s \mu(x+u,u)du} \right] ds \\
&= \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T r(s)ds} (G - F(T))^+ \right] \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T \mu(x+s,s)ds} \right] \\
&- m_e \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} F_s \right] \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s \mu(x+u,u)du} \right] ds \\
&= \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T r(s)ds} (G - F(T))^+ \right] {}_Tp_x - m_e \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} F_s \right] {}_sp_x ds, \quad (2.5.2)
\end{aligned}
$$

where

- G is the minimum guaranteed amount defined at inception of the contract;
- $F(t)$ is the value of the fund account at time t in which the initial premium was invested;
- m_e is the annualised fee for the GMMB deducted by the insurer from the fund account expressed as a percentage.

Furthermore, $F(t) = F(0)\frac{S(t)}{S(0)}e^{-mt}$, where S_t denotes the value of the underlying equity index at time t, and m denotes the annualised charge payable by the insurer, expressed as a percentage. We assume that $m_e < m$.

Note from Eq. (2.5.2) that the terms $\mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T r(s)ds} (G - F(T))^+ \right]$ and $\mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T \mu(x+s,s)ds} \right]$ were split using the relation in Eq. (2.5.1) since the force of mortality is independent from all other risk drivers. Furthermore, if the bank account is used as numeraire, the discount factor cannot be removed from the expectation since the interest rate process is stochastic and correlated to the stochastic process driving the fund value. Therefore, the terms $\mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^T r(s)ds} (G - F(T))^+ \right]$ and $m_e \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} F(s) \right] {}_sp_x$ can be solved using the brute force Monte Carlo method. Ignatieva et al. (2016) proposed a Fourier space time-stepping (FST) algorithm to reduce the computational cost of pricing embedded derivatives. See Ignatieva et al. (2016) for details.

The price of a GMDB at $t = 0$ from a life insurer's perspective is given by the formula:

$$
\begin{aligned}
V_{GMDB}(0) &= \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} \mu(x+s,s) e^{-\int_0^s \mu(x+u,u)du} (G(s) - F(s))^+ \right] ds \\
&- \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} m_d e^{-\int_0^s \mu(x+u,u)du} F(s) \right] ds \\
&= \int_0^T \mathbb{E}_{\mathbb{Q}} \left[e^{-\int_0^s r(u)du} (G(s) - F(s))^+ \right] \mathbb{E}_{\mathbb{Q}} \left[\mu(x+s,s) e^{-\int_0^s \mu(x+u,u)du} \right] ds
\end{aligned}
$$

$$
\begin{aligned}
&- m_d \int_0^T \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^s r(s)ds} F_s\right] \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^s \mu(x+u,u)du}\right] ds \\
&= \int_0^T \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^s r(u)du}(G(s) - F(s))^+\right] \mathbb{E}_{\mathbb{Q}}\left[\mu(x+s,s)e^{-\int_0^s \mu(x+u,u)du}\right] ds \\
&- m_d \int_0^T \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^s r(s)ds} F(s)\right] {}_sp_x ds, \qquad (2.5.3)
\end{aligned}
$$

where m_d denotes the annualised fee for the GMDB deducted by the insurer from the fund account. We assume that $m_d < m$.

Recall the probability of survival formula for a life aged x at inception:

$$
{}_tp_x = S_x(0,t) = \mathbb{E}_{\mathbb{Q}}\left[e^{-\int_0^t \mu(x+s,s)ds}\right].
$$

The instantaneous probability of dying at time t and age $x+t$, given survival up to time t, is given by the term $\mathbb{E}_{\mathbb{Q}}\left[\mu(x+t,t)e^{-\int_0^t \mu(x+s,s)ds}\right]$. This term is used extensively in the GMDB pricing formula and can be solved numerically by the brute force Monte Carlo method.

The pricing formula for the GMDB is more complex than the pricing formula for the GMMB. This is because the payoff of the GMDB product can occur at any time in the interval $[0,T]$ depending on when the policyholder dies, whereas the GMMB only pays an amount of $\max(G, F(T))$ if the policyholder survives to time T.

In the next section, we introduce the data used in this study.

2.6 Data

Daily FTSE/JSE Top40 closing prices from September 2005 to November 2020 were downloaded from *za.investing.com*[6]. The historical closing prices are shown in Figure 2.1 below.

[6] *https://za.investing.com/indices/ftse-jse-top-40-historical-data*

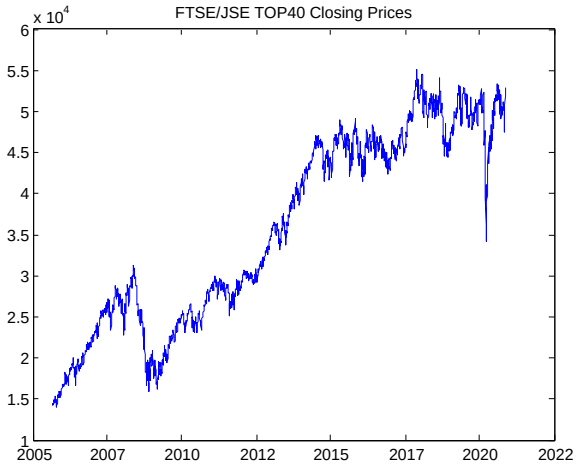

Fig. 2.1: Historical FTSE/JSE Top40 closing prices

For the Hull and White (1990) component of the Heston-Hull-White model, we require a starting value for $r(0)$. We use the South African 3-month T-Bill rate as a proxy for the short-term interest rate. Historical 3-month T-Bill rates were downloaded from the South African Reserve Bank website[7] for the period September 2005 to November 2020 and are shown in Figure 2.2 below:

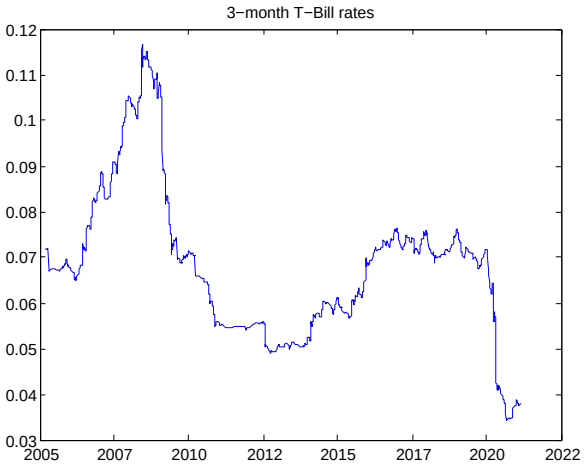

Fig. 2.2: Historical 3-month T-Bill rates in South Africa

[7] *https://www.resbank.co.za/en/home/what-we-do/statistics/key-statistics/selected-historical-rates*

South African mortality rates were sourced from the Actuarial Society of South Africa's investigation: *Report on pensioner mortality 2005-2010*[8], published in February 2017. The investigation was conducted by the Continuous Statistical Investigations (CSI) Committee of the Actuarial Society of South Africa. In the study, the CSI used pensioner data from the start of 2005 to the end of 2010 to construct the force of mortality for each age from 50 to 110, split by male and female.

Figure 2.3 below shows the survival probability, ${}_tp_{50},\ t \in \{1, 2, ..., 60\}$, for a 50-year old South African male and female pensioner:

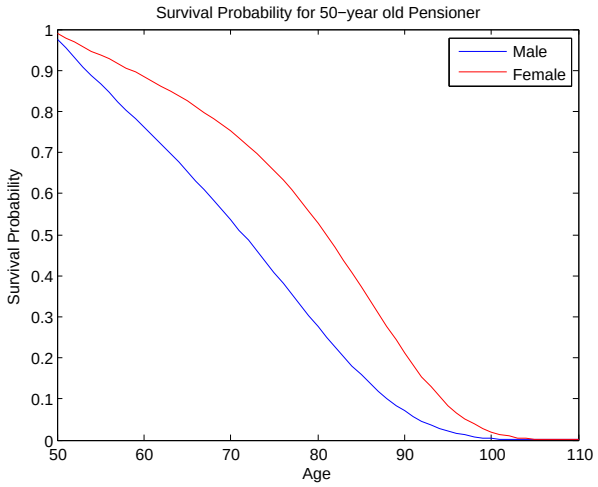

Fig. 2.3: Survival probability for 50-year old male and female pensioner

From Figure 2.3, the survival probability for a South African 50-year old male and female pensioner decreases the further we look into the future. Furthermore, the survival probability for a 50-year old male is considerably lower than a 50-year old female. The CIR++ model is calibrated to the survival probability curves in Figure 2.3.

For the purpose of the AR(1)-ARCH(1) model, we were unable to source historical mortality rates, $\mu(x, t)$, for South African males and females. However, our goal is not to focus on the quality or availability of historical mortality data, but rather to show how to calibrate and apply the AR(1)-ARCH(1) model. Therefore, we use historical mortality rates for the United Kingdom (UK) from 1922 to 2020 sourced from the Human Mortality Database[9] (HMD), the world's leading source of mortality data for developed countries. The historical mortality rates for the UK are split by age and sex.

[8] http://legacy.actuarialsociety.org.za/Societyactivities/CommitteeActivities/ContinuousStatisticalInvestigation(CSI).aspx
[9] https://www.mortality.org/

It is important to note that mortality in the UK is not representative of mortality in South Africa due to various social and economic factors such as income, education, employment, and community safety. However, using UK mortality data does not diminish the point we are trying to make. The only reason for using UK mortality data is that it is comprehensive and freely available from HMD. Any life insurer in South Africa with access to South African mortality data can easily adjust the input data to calibrate the AR(1)-ARCH(1) model.

Figures 2.4 and 2.5 below show the log mortality rates for UK males and females from 1922 to 2020 for ages 0 to 100:

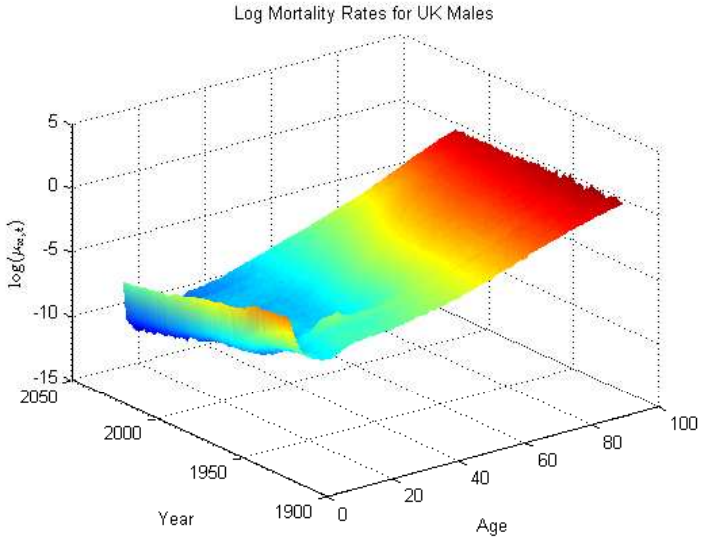

Fig. 2.4: Log mortality rates for UK males

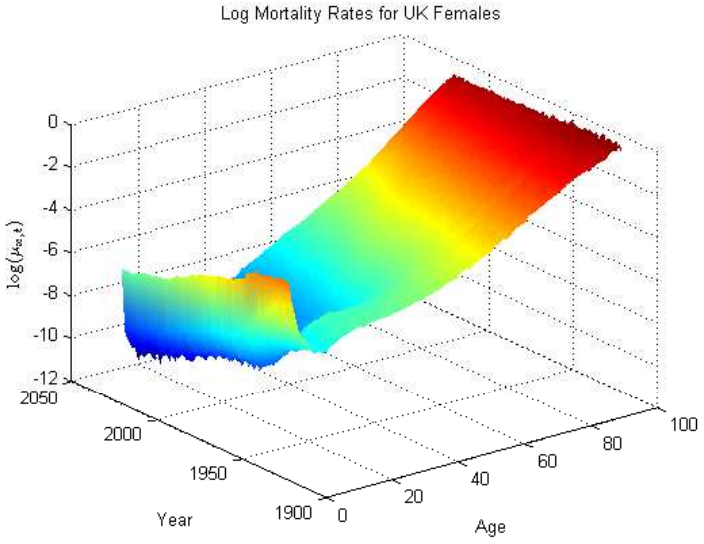

Fig. 2.5: Log mortality rates for UK females

Figures 2.6 and 2.7 below show the yearly changes in log mortality rates for UK males and females:

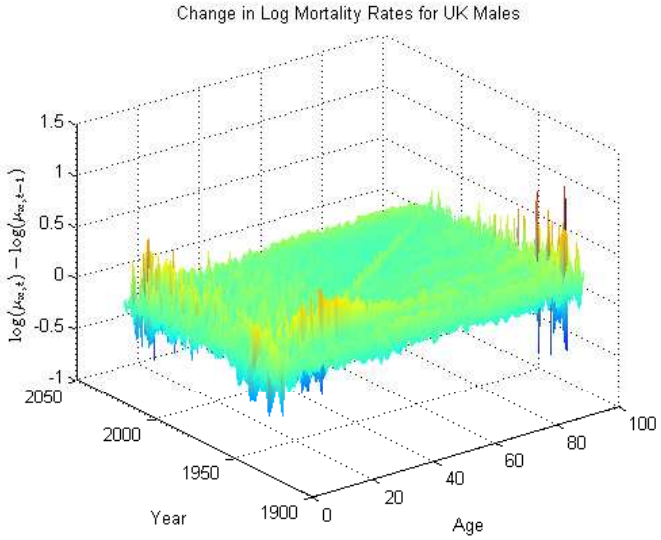

Fig. 2.6: Change in log mortality rates for UK males

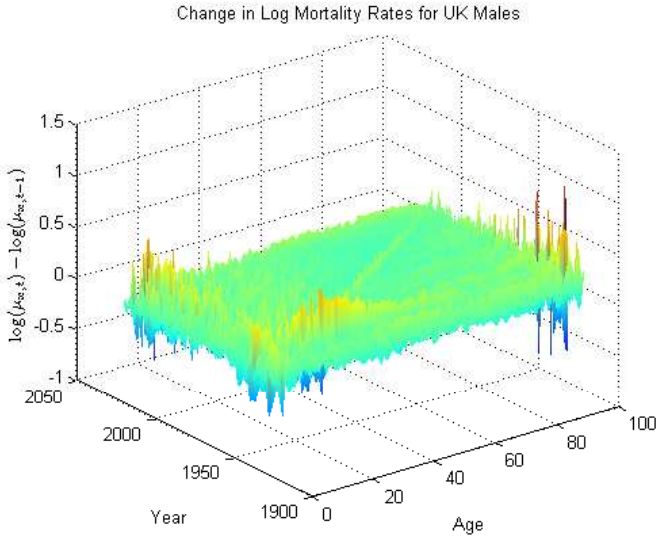

Fig. 2.7: Change in log mortality rates for UK females

The AR(1)-ARCH(1) model of Syuhada and Hakim (2021) can now be applied to the stationary time series of the changes in log mortality rates for each age.

Weekly FTSE/JSE Top40 volatility surfaces from September 2005 to November 2020 were provided by Peresec[10]. Each volatility surface covers a forward moneyness range of 75% to 125% and maturities from 1 to 15 months. For more details on the methodology used to construct the volatility surfaces, see Flint and Maré (2017).

The South African equity options market is illiquid compared to developed markets like Europe and the United States. As mentioned by Flint and Maré (2017), it is uncommon for options to trade beyond an expiry of 15 months in South Africa. This poses a further challenge when pricing GMMB and GMDB products with maturities typically extending far beyond 15 months.

Due to the illiquid equity options market in South Africa, there is no benchmark volatility that can be used for calibration beyond the 15-month mark. This implies that the volatility term-structure produced by the model can yield nonsensical values for long-dated maturities. Flint et al. (2014) tested several models that can produce long-term volatility estimates and showed that different models yield substantial differences for the volatility term-structure.

In this chapter, we allow the Heston-Hull-White-Mortality model to generate the volatility term-structure beyond the 15-month mark. We will test the model by plotting the term-structure of volatility to ensure that the model leads to plausible estimates for long-term volatility. For more detail on long-term volatility estimation, see Flint et al. (2014).

Figure 2.8 below shows the implied volatility surface constructed from FTSE/JSE Top40 index options on 16 November 2020:

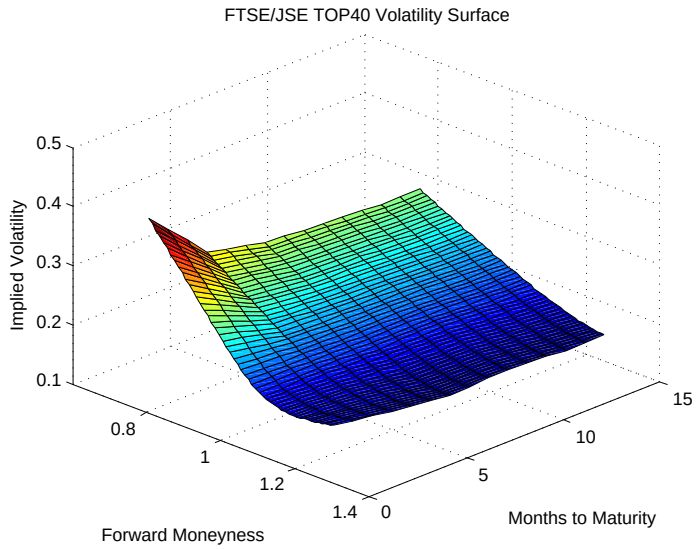

Fig. 2.8: FTSE/JSE Top40 volatility surface on 16 November 2022

[10]https://www.peresec.com/

Note from Figure 2.8 the volatility skew observed across forward moneyness levels and the non-constant volatility across term to maturity. APN 110 requires that the investment return projection model reproduce this surface as closely as possible.

In the next section, we implement the Heston-Hull-White-Mortality model to price GMMB and GMDB products written on the FTSE/JSE Top40 index, and backtest a delta-hedging strategy for the GMMB.

2.7 Numerical Results

This section has two subsections. The first subsection shows the pricing results obtained from the Heston-Hull-White-Mortality model for the GMMB and GMDB products and the sensitivity of the prices to various model parameters. In the second subsection, we perform a backtesting experiment over weekly periods from September 2005 to November 2020 to test the hedging performance of the Heston-Hull-White-Mortality model for the GMMB liability.

2.7.1 Pricing

The first step in pricing the GMMB and GMDB liabilities is to calibrate the Heston-Hull-White parameters to market data. We optimise the set of parameters on the following search space:

$$\begin{aligned}\Omega^{Search} &= D_{\kappa} \times D_{v_0} \times D_{\bar{v}} \times D_{\sigma_v} \times D_{\rho_{s,v}} \times D_{\lambda_r} \times D_{\eta} \times D_{\rho_{s,r}} \\ &= [1,1] \times [0,1] \times [0,1] \times [0,3] \times [-1,1] \times [0,10] \times [0,1] \times [-1,1],\end{aligned}$$

where $D_{\cdot}$ denotes the search domain for each parameter.

Grzelak and Oosterlee (2011) suggest that the Heston-Hull-White model be calibrated in two stages. First, the parameters for the Hull and White (1990) short rate process should be estimated independently from the equity (Heston, 1993) component. Then, keeping the parameters fixed for the short-term interest rate process, the remaining parameters are estimated.

The Hull and White (1990) mean reversion parameter, λ_r, is generally estimated from historical data in practice. Once λ_r has been estimated, the volatility parameter, η, is calibrated to a set of at-the-money swaptions.

The purpose of this chapter is not to focus on the calibration of the Hull and White (1990) model, since this in itself can become rather complex, see, for example, Gurrieri et al. (2009).

Rather, we aim to show the sensitivity of the GMMB and GMDB prices to various short-term interest rate parameters.

Once the parameters for the Hull and White (1990) component of the Heston-Hull-White model have been estimated, the next step is to calibrate the volatility parameters for the Heston (1993) component of the Heston-Hull-White model. We implemented the Fourier method of Carr and Madan (1999) discussed in Section 2.3.2.

The Heston-Hull-White characteristic function in Section 2.3.1 was implemented in Matlab. The next step involved solving the infinite integral in the Fourier transform pricing formula, $\int_0^\infty Re\Big\{e^{-ivk}\psi(v)\Big\}dv$, via numerical integration.

Choosing a suitable upper bound, $v_{\max}$, for the integral, and discretising the integration grid uniformly, the integral was evaluated using the trapezoidal rule of integration. We set $v_{\max} = 50$ and $dv = 0.01$ and used the *trapz*[11] function in Matlab.

For the characteristic function, we fixed the parameters $\{\kappa, \lambda_r, \eta, \rho_{s,r}\}$, and left the parameters $\{v(0), \bar{v}, \sigma_v, \rho_{s,v}\}$ free to estimate, bounded by the constraints in Ω^{Search}. We used the least squares optimisation algorithm *lsqnonlin*[12] in Matlab to minimise the sum of squares between the model and European call option prices. We fixed $\kappa = 1$ since the calibration became unstable when this parameter was left free to estimate. Similar findings were reported by van Dijk et al. (2018).

The details of the calibration for the Heston-Hull-White model to the FTSE/JSE Top40 volatility surface on 16 November 2020 are shown in Table 2.1 below:

Tab. 2.1: Calibration details on 16 November 2020

Number of strikes included	21
Number of maturities included	3
Number of options included	63
θ	0.07
λ_r	0.05
η	0.02
$\rho_{s,r}$	0.3
κ	1
v_0	Free to estimate
$\bar{v}$	Free to estimate
σ_v	Free to estimate
$\rho_{s,v}$	Free to estimate

[11] https://www.mathworks.com/help/matlab/ref/trapz.html
[12] https://www.mathworks.com/help/optim/ug/lsqnonlin.html

The calibration results for the Heston-Hull-White model with $v(0) = 0.0433$, $\bar{v} = 0.05$, $\sigma_v = 0.3817$, and $\rho_{s,v} = -0.9208$ are shown in Figure 2.9 below:

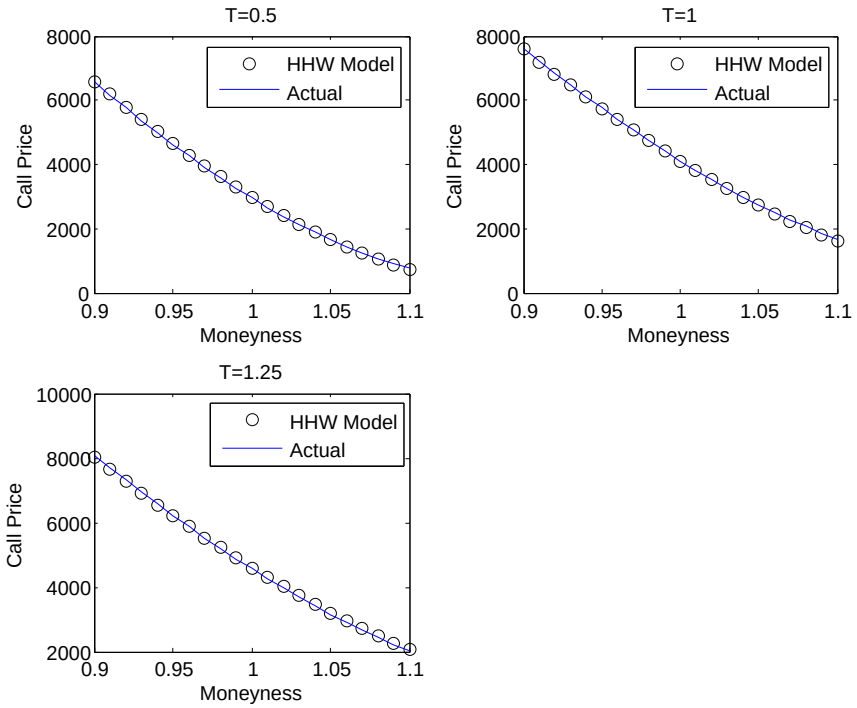

Fig. 2.9: Heston-Hull-White model fit to market data on 16 November 2020

Figure 2.9 shows that the Heston-Hull-White model reproduces the European call option prices well. The model is therefore aligned with the guidelines set by APN 110 in terms of being market-consistent.

As mentioned earlier, the South African equity options market is short-dated. From Figure 2.8, the longest maturity observable is 15 months. Since the GMMB and GMDB liabilities have maturities that extend far beyond this point, we must test whether the calibrated Heston-Hull-White model produces a plausible volatility term-structure for maturities beyond 15 months. Figure 2.10 below shows the volatility term-structure for at-the-money options on 16 November 2020.

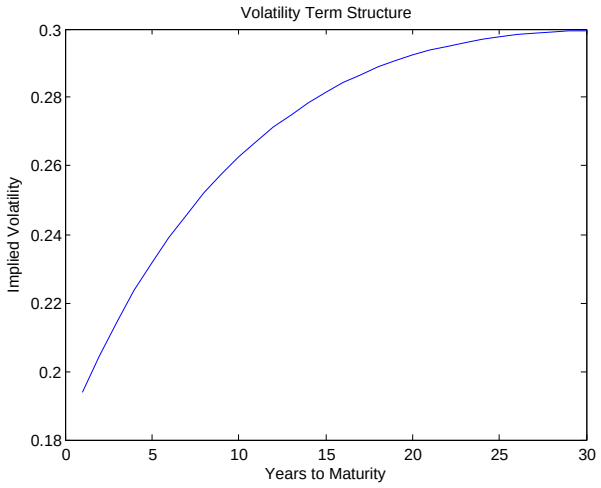

Fig. 2.10: Heston-Hull-White volatility term-structure on 16 November 2020

From Figure 2.10, the 30-year annualised volatility produced by the Heston-Hull-White model on 16 November 2020 is approximately 30%. This is a reasonable estimate and falls within the range of values reported by Flint et al. (2014) in their testing of various models. However, more advanced methods for long-term volatility estimation are available and the reader is referred to Flint et al. (2014) for further detail.

In order to price the GMMB and GMDB liabilities, we use a brute force Monte Carlo simulation to generate sample paths for each state variable. We use an Euler discretisation scheme to discretise the time grid. The discretised Heston-Hull-White model is given by:

$$\begin{cases} S(t+1) &= S(t) + r(t)S(t)\delta t + \sqrt{v(t)}S(t)(W_s(t+1) - W_s(t)), \quad S(0) > 0, \\ v(t+1) &= v(t) + \kappa(\bar{v} - v(t))\delta t + \sigma_v\sqrt{v(t)}(W_v(t+1) - W_v(t)), \quad v(0) > 0, \\ r(t+1) &= r(t) + \lambda_r(\theta - r(t))\delta t + \eta(W_r(t+1) - W_r(t)), \qquad r(0) \in \mathbb{R}, \end{cases}$$

where δt denotes the spacing of the time grid, and $W(t+1) - W(t) \sim \mathcal{N}(0, \delta t)$. To generate correlated Brownian motion increments for the state variables, we use Cholesky factorisation[13].

Using a time spacing of $\delta t = 1/252$, Figure 2.11 below shows a single simulation from the Heston-Hull-White model for each state variable.

[13] https://www.mathworks.com/help/matlab/ref/chol.html

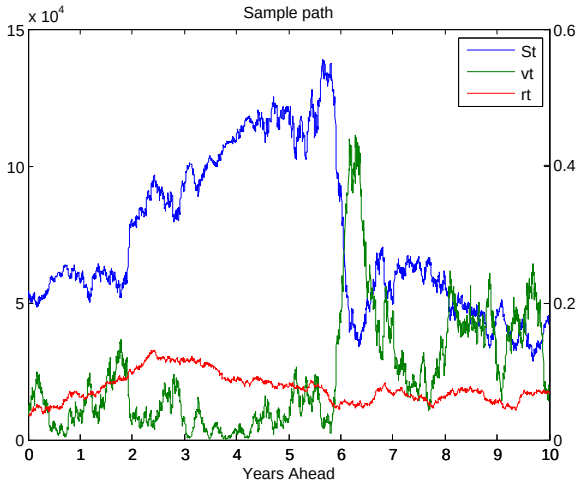

Fig. 2.11: Single sample path for each state variable

In Figure 2.11, the left y-axis refers to the share price, and the right y-axis refers to the volatility and interest rate. Note the inverse relationship between the volatility and share price. From our calibration, we estimated $\rho_{s,v} = -0.9208$, hence, the inverse relationship between the two state variables. This relationship is generally observed in equity markets and Black (1976) refers to this as the leverage effect. Recall that $\rho_{v,r} = 0$ so that the volatility and interest rate processes move independently. In our pricing experiment, we test the sensitivity of the GMMB and GMDB prices to various Heston-Hull-White model parameters.

As mentioned, the CIR++ model is a term-structure model that can fit the input survival probability curve. We consider a 50-year old South African male and female pensioner with parameters shown in Table 2.2 below:

Tab. 2.2: CIR++ parameters for 50-year old male and female pensioners

Parameter	Male	Female
γ	0.90	1.00
ω	0.05	0.03
ξ	0.03	0.01

Figure 2.12 shows the simulation for the force of mortality, $\mu(50+t,t)$, for a 50-year old South African male and female pensioner from the CIR++ model.

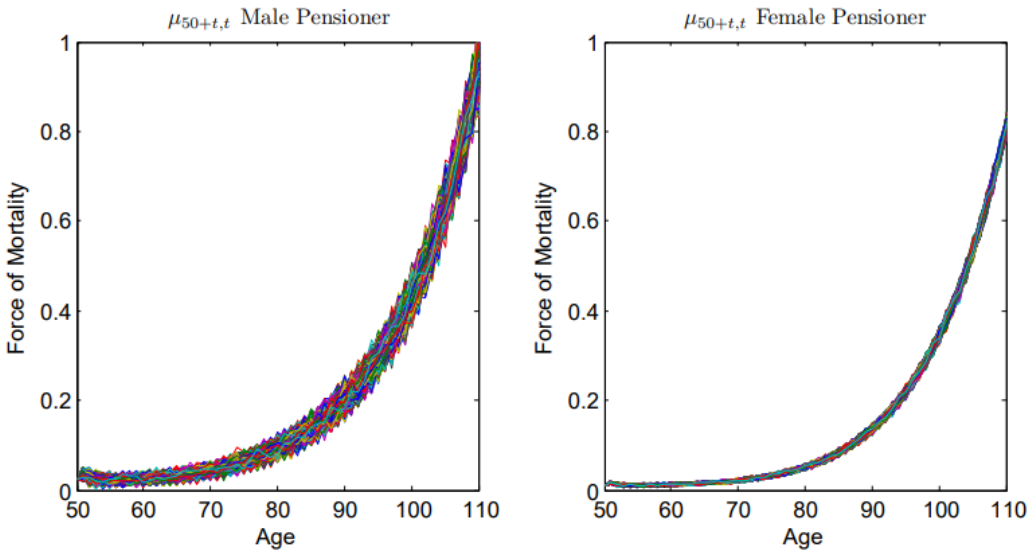

Fig. 2.12: CIR++ simulation for force of mortality

Note that the force of mortality increases with age. Furthermore, the distribution for $\mu(50+t,t)$ is wider for males than females. This is driven by the volatility of mortality, set to $\xi = 0.03$ for males and $\xi = 0.01$ for females.

Figure 2.13 below shows the stochastic survival probability curves calculated from the formula $e^{-\int_0^t \mu(50+s,s)ds}$:

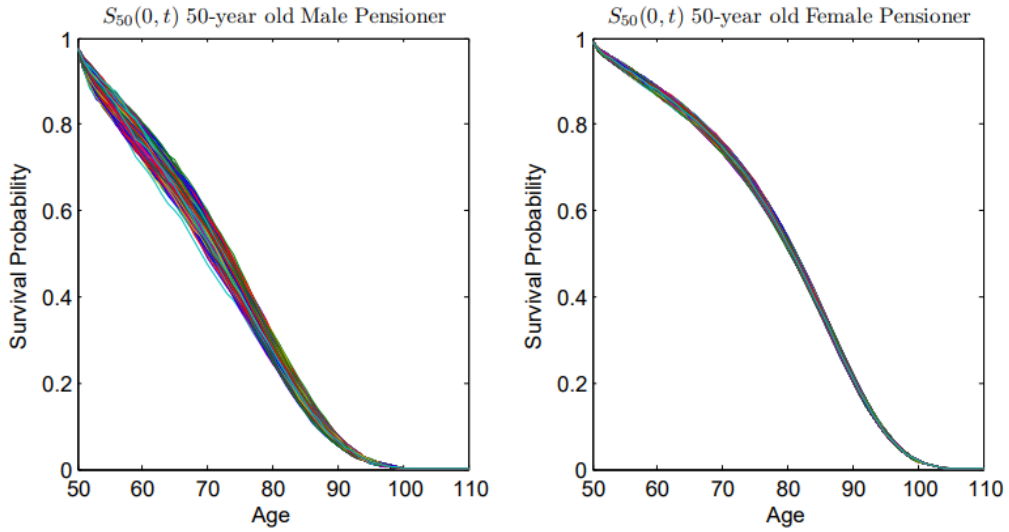

Fig. 2.13: Stochastic survival probability curves

Figure 2.14 below shows the mean survival probability curve, $\mathbb{E}_{\mathbb{Q}}[e^{-\int_0^t \mu(50+s,s)ds}]$, compared to the input survival probability curve for a 50-year old South African male and female pensioner:

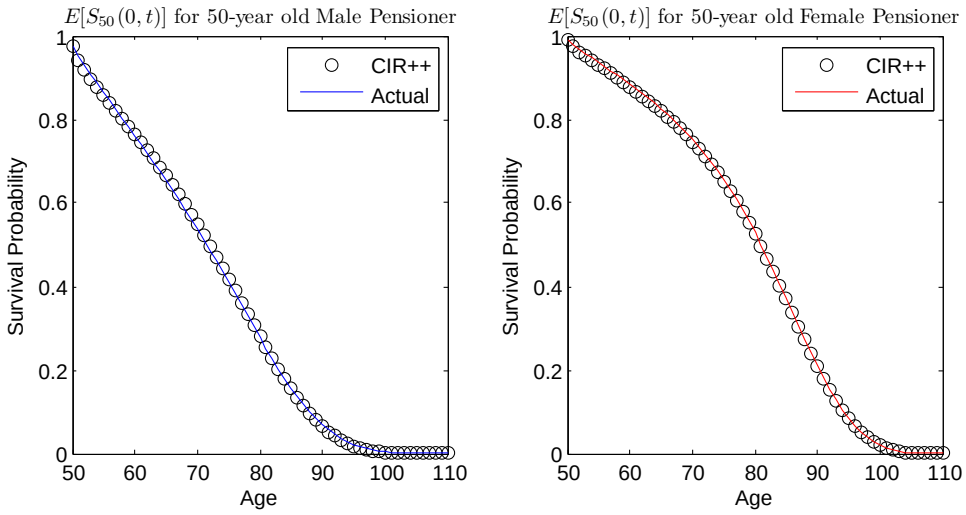

Fig. 2.14: Mean survival probability curve versus input survival probability curve

Figure 2.14 shows that the CIR++ model reproduces the input survival probability curves. Although the CIR++ model has the attractive feature of matching the supplied survival probability curve and incorporating volatility for mortality, we recommend that further research be done to test the impact of non-mean reverting models on mortality.

To illustrate the performance of the AR(1)-ARCH(1) model, we calibrated the model to yearly changes in the log mortality rate for the UK female population in Figure 2.7. For each fixed age, we estimated the parameters over the period 1922 to 2020 using the *ar*[14] and *garch*[15] packages in Matlab. The calibrated parameters are shown in Figure 2.15 below.

[14]https://www.mathworks.com/help/ident/ref/ar.html
[15]https://www.mathworks.com/help/econ/garch.html

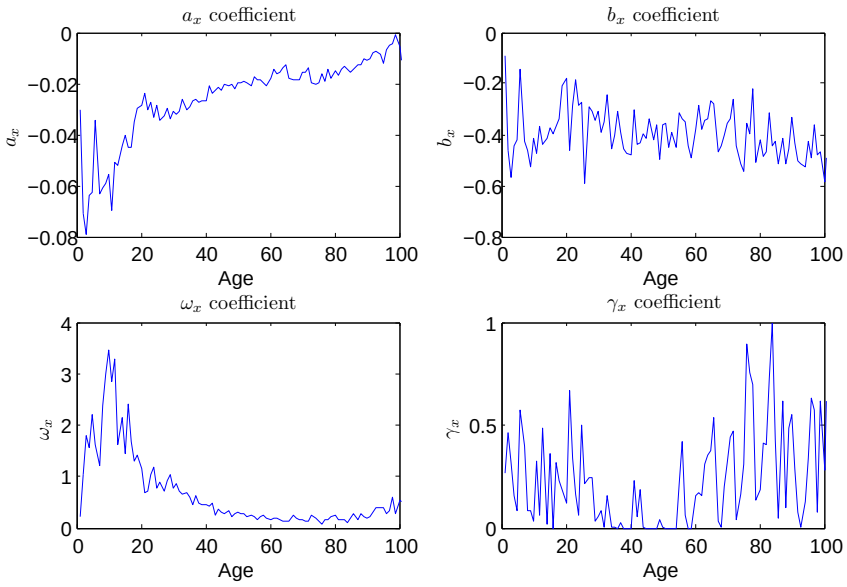

Fig. 2.15: Calibrated AR(1)-ARCH(1) parameters for UK females

Figure 2.16 below shows the predicted force of mortality for each year from 1922 to 2020 for an 85-year old UK female:

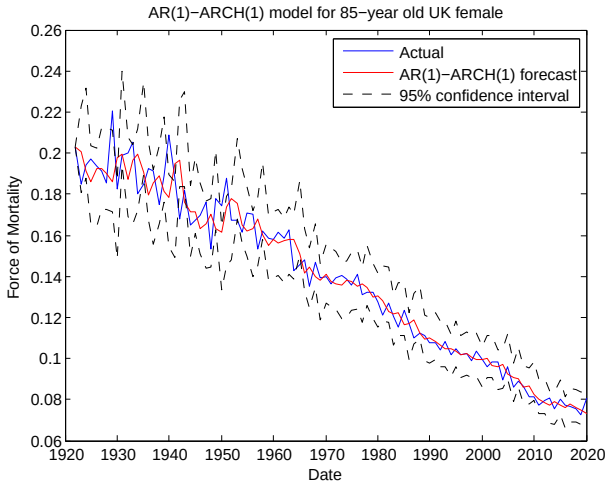

Fig. 2.16: AR(1)-ARCH(1) forecast for 85-year old UK female

The AR(1)-ARCH(1) model predicts the force of mortality well. Furthermore, the actual mortality rates generally lie within the 95% confidence interval.

To test the goodness-of-fit, the standardised residuals (assuming normal innovations for the AR(1)-ARCH(1) model) are compared to the standard normal density function in Figure 2.17 below:

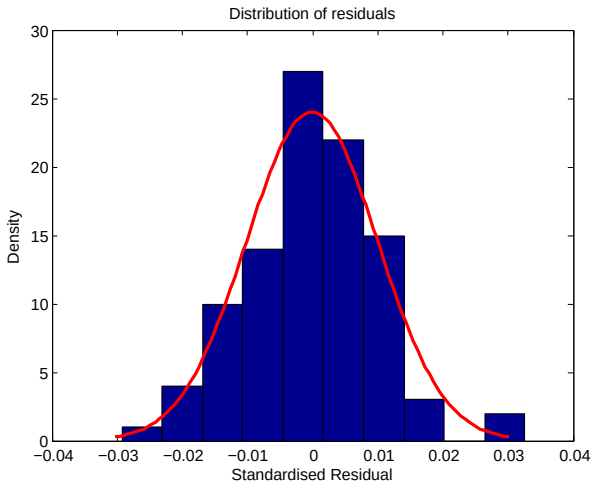

Fig. 2.17: Standardised residuals versus standard normal density

The Jarque-Bera test can be used to test the normality of the data - the null hypothesis is that the standardised residuals are normally distributed. We calculated the p-value for the Jarque-Bera test as 0.08, hence, we do not reject the null hypothesis at a 5% level of significance. We conclude that the AR(1)-ARCH(1) model with normally distributed innovations captures the mean and variance for 85-year old UK females well.

To construct a survival probability curve, the AR(1)-ARCH(1) model is fit to a time series of mortality rates for each fixed age using the age specific parameters in Figure 2.15.

As mentioned, we are not aware of historical mortality data that is publicly available in South Africa. Therefore, we showed how to calibrate the AR(1)-ARCH(1) model to mortality data sourced from the UK. The purpose was to show the mechanics of the AR(1)-ARCH(1) model, and the reader should note that the model can easily be applied to historical South African mortality rates by just changing the input data.

The focus of this chapter is specifically on the South African market, hence, we apply the CIR++ model going forward since survival probability curves for South African male and female pensioners are readily available from the CSI Committee's "Report on pensioner mortality 2005-2010".

Before we present the sensitivity of the GMMB and GMDB prices to various Heston-Hull-White-Mortality model parameters, we first create a base set of parameters. We use the parameters

from Tables 2.1 and 2.2 as our base. Table 2.3 below shows the base parameters for a 50-year old male and female pensioner in South Africa:

Tab. 2.3: Base parameters on 16/11/2020

Hull-White component	θ	0.07
	$r(0)$	0.04
	λ_r	0.05
	η	0.02
	$\rho_{s,r}$	0.3
Heston component	κ	1.00
	$v(0)$	0.04
	$\bar{v}$	0.05
	σ_v	0.38
	$\rho_{s,v}$	-0.92
CIR++ component	$\mu_{Male}(50,0)$	0.02
	$\mu_{Female}(50,0)$	0.01
	γ_{Male}	0.90
	γ_{Female}	1.00
	ω_{Male}	0.05
	ω_{Female}	0.03
	ξ_{Male}	0.03
	ξ_{Female}	0.01

In the results below, we ran 10,000 simulations and considered a single premium of R100,000 invested in the GMMB and GMDB respectively. It is simple to extend the analysis to take into account recurring premiums, see Feng (2018). We report our results as a ratio of the GMMB and GMDB price to the initial premium invested at $t = 0$:

$$R_{GMMB} = \frac{V_{GMMB}(0)}{Initial\ Premium},$$

$$R_{GMDB} = \frac{V_{GMDB}(0)}{Initial\ Premium}.$$

First, we consider three scenarios for the GMMB price from a life insurer's perspective: 1) deterministic interest rates by setting $\eta = 0$, 2) an independent interest rate process by setting $\rho_{s,r} = 0$, and 3) constant equity volatility by setting $\sigma_v = 0$. We keep all other parameters in Table 2.3 fixed. Figure 2.18 below shows R_{GMMB} for all three scenarios with a guaranteed rate of $g = 6\%$ per annum, where $G = G(0) \times (1 + g)^T$, with $G(0) = 100,000$.

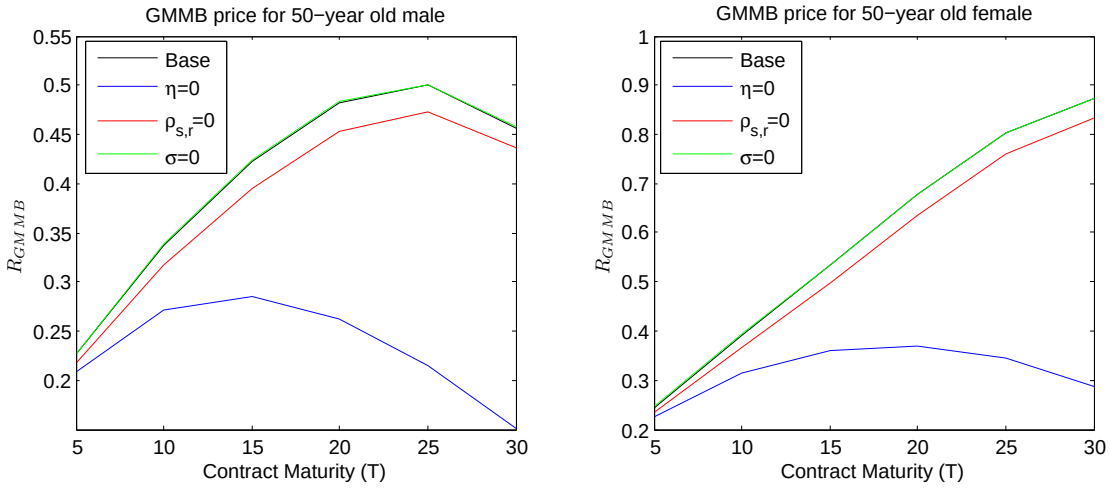

Fig. 2.18: GMMB price for a 50-year old South African pensioner

From Figure 2.18, we observe the following:

- From a life insurer's perspective, a 6% guarantee per annum relates to a life contingent put option that is written in-the-money, hence, there may be a significant liability;
- The GMMB price for a 50-year old female is higher than the price for a 50-year old male. This is driven solely by the fact that the survival probability for South African females is higher than males;
- Deterministic interest rates, $\eta = 0$, yield GMMB prices that are significantly lower than the base model. The difference is amplified as the term to maturity increases;
- Independent equity and interest rates, $\rho_{s,r} = 0$, yield GMMB prices that are lower than base model, but higher than the deterministic interest rate model;
- Constant equity volatility, $\sigma_v = 0$, yields GMMB prices that are very close to the base model.

Our findings suggest that interest rate risk is the dominating factor when pricing GMMB products. Assuming deterministic rates may lead to GMMB prices that are significantly understated, and, hence, pose a large risk to the life insurer. Assuming independent interest rate and equity processes may also not be sufficient. An interesting observation is that stochastic volatility makes nearly no impact on the price of the GMMB.

Next, we present three scenarios for the GMDB price from a life insurer's perspective: 1) deterministic interest rates by setting $\eta = 0$, 2) an independent interest rate process by setting $\rho_{s,r} = 0$, and 3) constant equity volatility by setting $\sigma_v = 0$. As with the GMMB product, we

keep all other parameters in Table 2.3 fixed. Figure 2.19 below shows R_{GMDB} for all three scenarios with a guaranteed roll-up rate of $g = 6\%$ per annum, where $G(t) = G(0)e^{gt}$, with $G(0) = 100,000$.

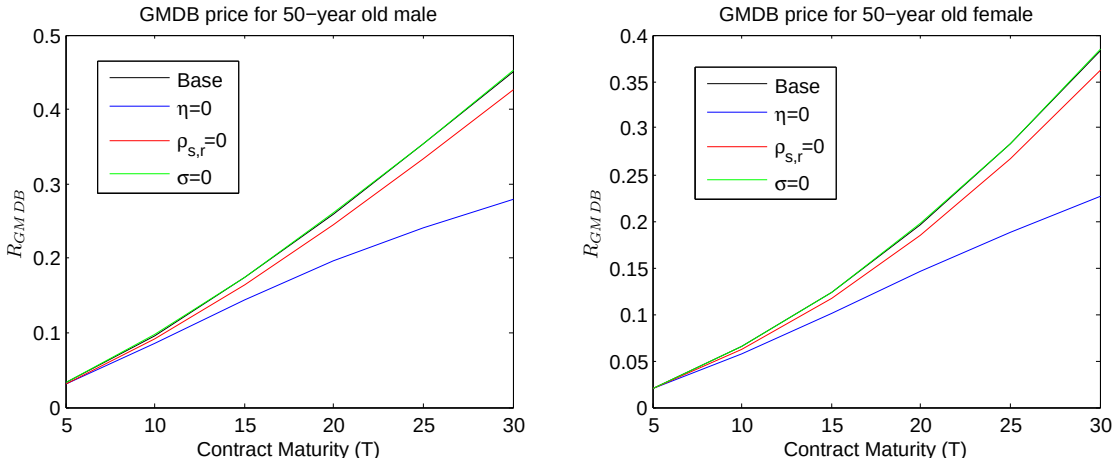

Fig. 2.19: GMDB price for a 50-year old South African pensioner

From Figure 2.19, we observe the following:

- The GMDB price increases as the term to maturity increases, whereas the GMMB price starts to decrease as the contract maturity becomes larger;
- The GMDB price for a 50-year old male is higher than the price for a 50-year old female. This is because the force of mortality for South African males is higher than females;
- Deterministic interest rates, $\eta = 0$, yield GMDB prices that are considerably lower than the base model. The difference is amplified as the term to maturity increases;
- Independent equity and interest rates, $\rho_{s,r} = 0$, yield lower prices than the base model, but higher prices than the deterministic interest rate model;
- Constant volatility, $\sigma_v = 0$, yields GMDB prices that are very close to the base model;

Once again, our findings suggest that interest rates are the dominating factor when pricing GMDB products.

Stochastic mortality has application in risk management problems and can be applied to Value-at-Risk (VaR) (Syuhada & Hakim, 2021) and longevity stress testing (Browne et al., 2009), for example.

In the next section, we apply the Heston-Hull-White-Mortality model to hedge the GMMB.

2.7.2 Hedging

Standard Black and Scholes (1973) theory suggests that a contingent claim can be replicated with a combination of shares and cash. This is referred to as delta-hedging. Since the GMMB and GMDB payoffs resemble the payoff of a put option (although life contingent), Black and Scholes (1973) theory can be applied to hedge these products.

Let $\Delta(t)$ denote the number of shares that an insurer must hold in a hedging portfolio to hedge against movements in the underlying equity index. Furthermore, let $B(t)$ denote the money market account from which money can be borrowed or invested. No upfront premium is payable for the GMMB. Therefore, the value of the money market account at $t = 0$ is given by:

$$B(0) = -\Delta(0)S(0), \tag{2.7.1}$$

where $\Delta(0) \in [-1, 0]$ in the case of a put option.

For $t > 0$, the money market account grows at the risk-free rate, r. We assume that the risk-free rate is the 3-month T-Bill rate in South Africa. The change in the money market account over a discrete time interval is given by:

$$B(t) = B(t-1)e^{rdt} - (\Delta(t) - \Delta(t-1))S(t), \quad B(0) = -\Delta(0)S(0). \tag{2.7.2}$$

Note that we consider a discrete hedging portfolio since it is impractical to hedge continuously. At every time step, the hedging portfolio must be rebalanced to reflect the sensitivity of the derivative to the underlying asset price as new information becomes available.

The hedging portfolio at time t consists of a combination of shares and cash and is given by:

$$\Pi(t) = \Delta(t)S(t) + B(t). \tag{2.7.3}$$

In order to calculate $\Delta(t)$ in the Heston-Hull-White-Mortality model, we approximate $\Delta(t)$ with a forward finite difference scheme as follows:

$$\Delta(t) \approx \frac{V_{GMMB}(t, S(t) + 0.01, v(t), r(t), \mu(t)) - V_{GMMB}(t, S(t), v(t), r(t), \mu(t))}{0.01}. \tag{2.7.4}$$

Before we proceed with our backtesting experiment, we calibrate the Heston-Hull-White-Mortality model to weekly FTSE/JSE Top40 implied volatility surfaces from September 2005 to November 2020. The calibrated parameters are shown in Figure 2.20 below.

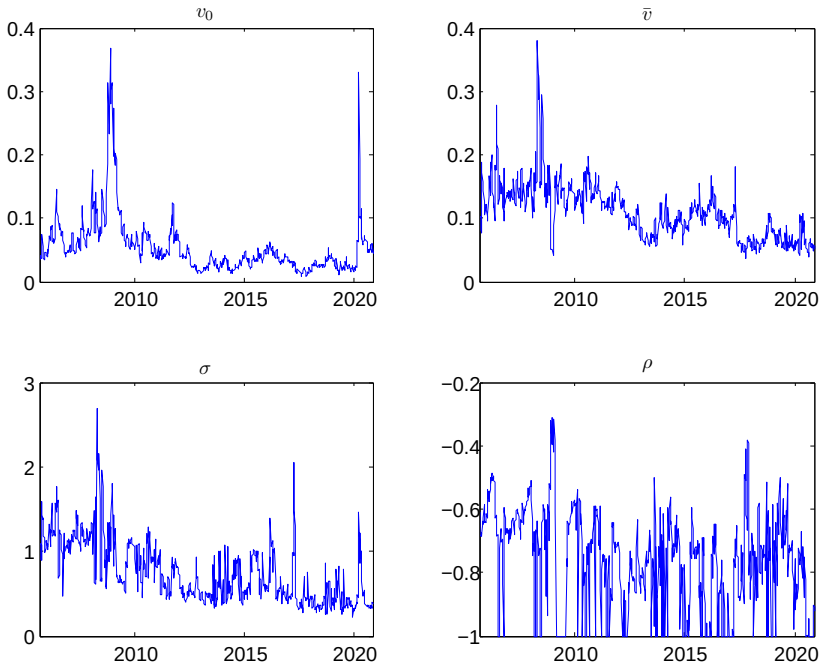

Fig. 2.20: Calibrated parameters over time

Note the clear spikes in volatility during the global financial crisis of 2007/2008 and the COVID-19 pandemic. In the calibration, we set $\kappa = 1$ since an unrestricted κ produced unstable results.

For this experiment, we rebalance the hedging portfolio once every week. We consider a 6% per annum GMMB product sold to a 50-year old male pensioner in September 2005, expiring in November 2020. We further assume that the investor is alive at each valuation date for this backtesting experiment.

Figure 2.21 below shows the value of the GMMB liability over time calculated from the Heston-Hull-White-Mortality model.

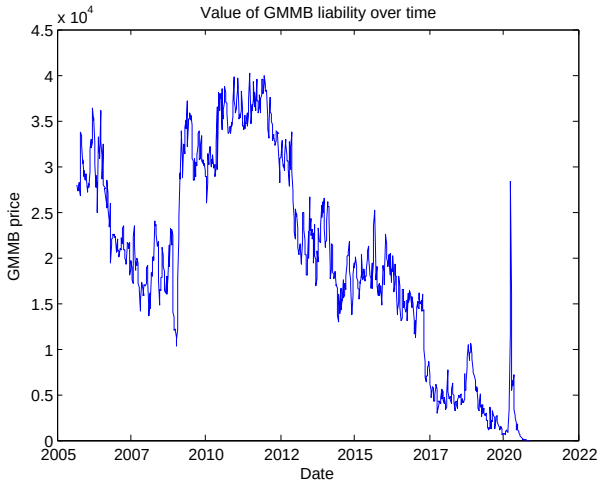

Fig. 2.21: GMMB liability over time

Note the extreme volatility in the GMMB price over time. The volatility is driven by movements in the underlying FTSE/JSE Top40 index, interest rates, and mortality. Furthermore, the GMMB liability displays significant jumps in the 2007/2008 global financial crisis and COVID-19 periods. Hedging is a key tool that can be used by life insurers to manage the risk of selling embedded derivatives.

Figure 2.22 below shows the value for Δ over time calculated from Eq. (2.7.4):

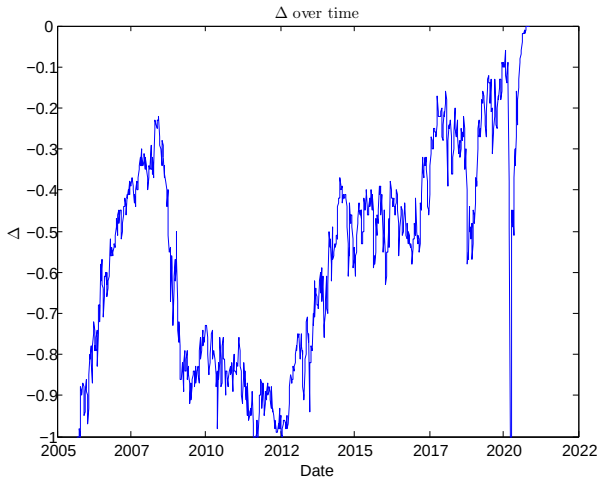

Fig. 2.22: GMMB Δ over time

In periods of stress and poor equity performance, Δ approaches -1. This means that equity exposure in the hedging portfolio increases. Note that Δ is close to -1 in the time of the global financial crisis and the COVID-19 pandemic. Alternatively, when equity performs well, Δ approaches 0 (equity exposure decreases). Hedging the GMMB requires active management by consistently updating the allocation in a short equity position and cash.

From Feng (2018), we define $V^*_{GMMB,t}$ as:

$$V^*_{GMMB}(t) := V_{GMMB}(t) - V_{GMMB}(0)e^{rt}. \tag{2.7.5}$$

$V^*_{GMMB}(t)$ represents the deviation from the risk-neutral expectation of the GMMB calculated at inception of the contract. Figure 2.23 below shows the performance of the delta-hedging strategy in Eq. (2.7.3) compared to Eq. (2.7.5):

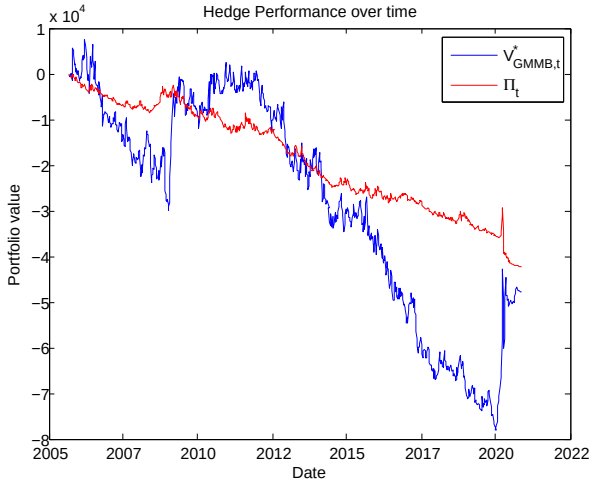

Fig. 2.23: Hedge performance over time

Delta-hedging captures the general trend of the movements in the GMMB liability, but fails to capture significant volatility and jumps. This makes sense since the market is incomplete under the Heston-Hull-White-Mortality model. Volatility, interest rates, and mortality are all factors that are not directly tradable in the South African market. To capture the effect of volatility, interest rates, and mortality on the GMMB liability, more products that derive their value from these risk factors need to be added to the hedging portfolio.

Although not perfect, delta-hedging does seem to provide some benefit for the GMMB liability. The hedging strategy captures the downward trend through time.

Hedging is a complex problem and highly relevant in the life insurance industry. Therefore, further research in the field of embedded derivative hedging is of significant importance.

2.8 Conclusion

In this chapter, we showed the importance of incorporating stochastic interest rates when reserving for GMMB and GMDB products since interest rate risk is the dominating factor when pricing these claims. Ignoring interest rate risk can lead to prices that are significantly understated from a life insurer's perspective. Stochastic volatility had no significant impact on the GMMB and GMDB liabilities.

Furthermore, we proposed the continuous-time CIR++ and discrete-time AR(1)-ARCH(1) models for mortality. The stochastic mortality models performed very well. The CIR++ model is able to reproduce the input survival probability curve, whereas the AR(1)-ARCH(1) model provides a good fit to historical data. Both models can be applied to risk management problems such as VaR and longevity stress testing, and we suggest further research on this topic.

Insurance liabilities change through time as the underlying risk factors change. To hedge these risks, we implemented a simple delta-hedging strategy based on the Heston-Hull-White-Mortality model and showed that a delta-hedging strategy captures the general trend of the movements in the liability value. However, under the Heston-Hull-White-Mortality model, the market is incomplete. Further research should be conducted to test whether a volatility (delta-vega) or interest rate (delta-rho) hedging strategy can produce better hedging results than a simple delta-hedging strategy. This is a significant challenge faced in the life insurance industry and finding a good hedge may lead to less volatility in the profit and loss (P&L) in the insurer's income statement.

The proposed Heston-Hull-White-Mortality model aligns with the requirements as set by APN 110 and is a good starting point for life insurers when deciding on a model for embedded derivatives.

3

Pricing Two-Asset Rainbow Options with the Fast Fourier Transform

> *I certainly view volatility as my friend. Volatility is on sale because 99% of the institutions out there are doing their best to avoid it.*
>
> **— Michael Burry**

Keywords: Characteristic function · fast Fourier transform · rainbow option · three-factor stochastic volatility model · two-factor geometric Brownian motion

3.1 Introduction

A rainbow option refers to an option that depends on more than one underlying risky asset, where each asset is seen as a colour of the rainbow (Ouwehand & West, 2006). These options come in various forms including the "best-of-d" and "worst-of-d" call options on d underlying assets. Pricing rainbow options is often challenging due to the absence of a closed-form solution; hence, numerical methods must be employed. This is the focus of this chapter[1].

The market for rainbow options is illiquid and these options are typically structured on demand. However, this could change in the future. Roberts (2018) lists various applications for rainbow options in the industry. Firstly, rainbow options can be used to gain exposure to the market at a lower cost whilst reducing risk (Klyueva, 2014). Furthermore, "best-of-d" call options can be used to hedge currency risk if a company has the option to settle their liabilities in various foreign currencies (Guillaume, 2008). Due to their potential, numerous methods have been proposed in the literature to price rainbow options.

The first contribution made in the literature on the pricing of two-asset rainbow options can be traced back to Stulz (1982). The author derived formulas for European call and put options on the minimum or maximum of two risky assets that involve the calculation of the bivariate

[1]This chapter is based on a paper (Levendis and Maré, 2023b) published in the *South African Statistical Journal*.

normal density function. Ouwehand and West (2006) derived formulas for "best-of-3" and "worst-of-3" call options using a change of numeraire technique. Eberlein et al. (2010) presented a general framework for the Fourier transform method of Carr and Madan (1999), which is used when certain conditions hold like the existence of the dampened characteristic function. The framework allows for the pricing of "best-of-d" and "worst-of-d" call and put options.

Fourier transformation is a mathematical method that decomposes a function into the sum of simpler trigonometric functions. When uniformly spaced samples of a continuous function are input to a Fourier transform, the transformation is called the discrete-time Fourier transform (DTFT). The input data are discrete and the output DTFT is a continuous function. If samples of equal length are taken from the DTFT output, then the transformation is called the discrete Fourier transform (DFT). Computing the DFT directly requires a total of N^2 operations to be performed, where N is the number of input data points. Therefore, the total computation time is $O(N^2)$.

The fast Fourier transform (FFT), pioneered by Cooley and Tukey (1965), is an efficient algorithm for computing the DFT. The Cooley-Tukey FFT algorithm requires N to be a power of 2, *i.e.*, $N = 2^m$, where $m \geq 0$. This requirement leads to significant time saving compared to the direct computation approach, with the total running time for Cooley-Tukey FFT algorithm being $O(N \log(N))$.

Carr and Madan (1999) developed a Fourier transform for the price of a European call option in terms of the characteristic function of the log of the stock price at the option maturity. The authors applied the FFT algorithm to compute European call option prices given the Fourier transform and concluded that the FFT yields significant improvement in terms of computation speed.

Hurd and Zhou (2010) extended the Fourier transform of Carr and Madan (1999) to the two-dimensional case of spread options and concluded that the FFT produces accurate and efficient spread option prices.

Building on the work of Eberlein et al. (2010), Roberts (2018) applied the two-dimensional FFT method of Hurd and Zhou (2010) to price "worst-of-2" call options based on the two-factor geometric Brownian motion (gBm) model. Roberts (2018) claims to be the first author to have applied the two-dimensional FFT to price two-asset rainbow options.

In order to compute option prices using the two-dimensional FFT, the double integral that appears in the Fourier transform must be approximated by truncating the domain $\mathbb{R}^2$ to a suitable lower and upper bound. The lower and upper bounds are user defined and termed the "truncation width".

Roberts (2018) shows that the FFT price for a "worst-of-2" call option can converge to the price obtained from the Stulz (1982) formula using as few as 512 terms. However, Roberts (2018) mentions that there is no single truncation width that works consistently well for all terms in the FFT and one first needs a price estimate before an appropriate value for the truncation width can be chosen. Moreover, further investigation is needed to test the accuracy of the FFT method for further in-the-money rainbow options. Roberts (2018) showed that an in-the-money "worst-of-2" call option did not converge to three decimals whereas all other strike prices did. Lastly, Roberts (2018) suggests that more complex underlying models be considered to price rainbow options other than the two-factor gBm.

In this chapter, we first attempt to replicate the FFT results in Roberts (2018) to test whether we arrive at the same conclusions for "worst-of-2" call options. We then extend the work of Roberts (2018) per the author's suggestion by applying the three-factor stochastic volatility model in addition to the two-factor gBm model for the underlying assets. We then apply the two-dimensional FFT method of Hurd and Zhou (2010) to price "worst-of-2" call options based on these dynamics. Our contribution to the literature is, therefore, the pricing of two-asset rainbow options with stochastic volatility using the FFT. To our knowledge, we are the first authors to apply the FFT method of Hurd and Zhou (2010) to "worst-of-2" call options based on dynamics other than the two-factor gBm.

The field of quantitative finance draws on many concepts from mathematical statistics including moment generating functions, characteristic functions, stochastic processes, and the multivariate normal distribution, to name a few. This chapter highlights the interplay between the fields of statistics and quantitative finance and that statistics plays a vital role in further development of quantitative finance.

The remainder of this chapter is structured as follows: Section 3.2 introduces the underlying dynamics that will be considered to price two-asset rainbow options. Section 3.3 shows the characteristic functions for the two-asset gBm and three-factor stochastic volatility models. Section 3.4 introduces the Hurd and Zhou (2010) two-dimensional FFT algorithm. Section 3.5 shows the numerical results, and Section 3.6 concludes the chapter.

3.2 Stock Price Dynamics for Two-Asset Options

This section introduces two stochastic processes for options that depend on two underlying assets. The choice of model generally depends on the behaviour of the market. One generally aims to choose a model that describes the market dynamics of interest rates, stock prices, and volatility as closely as possible. The first, and most basic, model we consider is two-factor gBm.

3.2.1 Two-Factor Geometric Brownian Motion

The two-factor gBm model was first introduced by Margrabe (1978) to price exchange options - the option to exchange one risky asset for another. The two-factor gBm model takes the following form:

$$\begin{cases} dS_1(t) &= (r-\delta_1)S_1(t)dt + \sigma_1 S_1(t)dW_{S_1}(t), \\ dS_2(t) &= (r-\delta_2)S_2(t)dt + \sigma_2 S_2(t)dW_{S_2}(t), \end{cases}$$

where $S_1(t)$, $S_2(t)$ denote the stock prices at time t; $dS_1(t)$, $dS_2(2)$ are the increments of the respective stock prices from time t to time $t+dt$, with dt an infinitesimal quantity; σ_1, σ_2 are the annualised volatility estimates for the two stocks; δ_1, δ_2 are the respective dividend yields; r is the constant risk-free rate, and $dW_{S_1}(t)dW_{S_2}(t) = \rho_{S_1,S_2}dt$, with ρ_{S_1,S_2} the correlation coefficient between the two stock prices.

The two-factor gBm model assumes that the two stock prices each follow a log-normal distribution with constant volatility. Clearly this assumption is very limiting since the model assumes symmetric returns for the log of the stock price based on a normal distribution. Empirical evidence has shown that returns for the log of the stock price tend to be negatively skewed (see Cont, 2001); hence, a model that can account for asymmetry is preferred to describe the behaviour of the equity market. Furthermore, the two-factor gBm model is unable to produce the fat-tailed returns often observed in equity markets (see Cont, 2001).

Next, we introduce the three-factor stochastic volatility model that addresses some of the shortcomings in the two-factor gBm model.

3.2.2 Three-Factor Stochastic Volatility

The three-factor stochastic volatility model was first introduced by Dempster and Hong (2002) to price spread options. The model assumes that both assets are driven by the same Cox et al. (1985) variance process and takes the following form:

$$\begin{cases} dS_1(t) &= (r-\delta_1)S_1(t)dt + \sigma_1\sqrt{v(t)}S_1(t)dW_{S_1}(t), \\ dS_2(t) &= (r-\delta_2)S_2(t)dt + \sigma_2\sqrt{v(t)}S_2(t)dW_{S_2}(t), \\ dv(t) &= \kappa\Big(\bar{v}-v(t)\Big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \end{cases}$$

where $v(t)$ denotes the variance of the stock prices at time t; $dv(t)$ is the increment of the variance from time t to time $t+dt$; κ is the mean reversion speed of the variance; $\bar{v}$ is the

long-run mean of the variance, and σ_v is the volatility of volatility. The Brownian motions are correlated as follows:

$$dW_{S_1}(t)dW_{S_2}(t) = \rho_{S_1,S_2}dt,$$
$$dW_{S_1}(t)dW_v(t) = \rho_{S_1,v}dt,$$
$$dW_{S_2}(t)dW_v(t) = \rho_{S_2,v}dt.$$

The three-factor stochastic volatility model has the ability to produce asymmetric and fat-tailed returns and is, therefore, more flexible than the two-factor gBm model.

In the next section, we introduce the work of Eberlein et al. (2010) along with the characteristic functions for the two-factor gBm and three-factor stochastic volatility models. The next section forms an integral part of the two-dimensional FFT method of Hurd and Zhou (2010), which will be discussed in Section 4.

3.3 Rainbow Options, Fourier Transform and Characteristic Functions

Fourier transform methods have led to significant computational gains since their introduction to option pricing by Carr and Madan (1999). The Fourier transform framework was further generalised to multi-asset options by Eberlein et al. (2010) where the authors listed conditions under which Fourier transform formulas are valid. In this section, we first introduce the payoff functions for "best-of-2" and "worst-of-2" call options. Next, we discuss the general Fourier transform framework of Eberlein et al. (2010). Lastly, we show the characteristic functions for the two-factor gBm and three-factor stochastic volatility models.

3.3.1 Best-of-2 and Worst-of-2 Call Options

Let $V_{max}(S_1(t), S_2(t))$ and $V_{min}(S_1(t), S_2(t))$ denote the values for a "best-of-2" and "worst-of-2" call option depending on two assets $\{S_1(t), S_2(t)\}$ at time t with strike price K and maturity T. At maturity T, the payoff formula for the call options is given by:

$$V_{max}(S_1(T), S_2(T)) = \max\Big(\max\Big(S_1(T), S_2(T)\Big) - K, 0\Big), \tag{3.3.1}$$

$$V_{min}(S_1(T), S_2(T)) = \max\Big(\min\Big(S_1(T), S_2(T)\Big) - K, 0\Big). \tag{3.3.2}$$

From Eberlein et al. (2010), the Fourier transform corresponding to the payoff functions in Eq. (3.3.1) and (3.3.2) can be derived and will be discussed in the next subsection.

3.3.2 The Result of Eberlein, Glau, and Papapantoleon

Eberlein et al. (2010) presented a general framework for which Fourier transform formulas are valid. For the two-dimensional case, the authors consider any payoff function $f : \mathbb{R}^2 \to \mathbb{R}_+$, for example Eq. (3.3.1) or (3.3.2), and the dampened payoff function:

$$g(\boldsymbol{x}) := e^{-\boldsymbol{\alpha}^\top \boldsymbol{x}} f(\boldsymbol{x}) \quad \text{for} \quad \boldsymbol{x} \in \mathbb{R}^2,$$

where $\boldsymbol{\alpha} \in \mathbb{R}^2$ is the dampening coefficient.

Let $\hat{g}$ denote the Fourier transform of the function g; $L^1_{bc}(\mathbb{R}^2)$ be the space of bounded, continuous functions in $L^1(\mathbb{R}^2)$, where $L^1(\mathbb{R}^2)$ is the space of all integrable functions on $\mathbb{R}^2$. Moreover, let $\boldsymbol{X}(0) = (x_1(0), x_2(0))^\top = (\log S_1(0), \log S_2(0))^\top$ and let $M_{\boldsymbol{X}_T}$ denote the moment generating function for the random variable $\boldsymbol{X}(T) = (x_1(T), x_2(T))^\top = (\log S_1(T), \log S_2(T))^\top$ at the option maturity T.

Eberlein et al. (2010) make the following assumptions:

1. Assume that $g \in L^1_{bc}(\mathbb{R}^2)$ and $\hat{g} \in L^1(\mathbb{R}^2)$;
2. Assume that $M_{\boldsymbol{X}_T}(\boldsymbol{\alpha})$ exists;

Under these assumptions, the authors present the following Fourier transform formula for two-asset options:

Theorem 1 *If the asset price processes are modelled as two-factor gBm or three-factor stochastic volatility processes and assumptions 1 and 2 hold, then the payoff function for a two-asset option V at $t = 0$ can be written as:*

$$V(\boldsymbol{X}(0)) = \frac{e^{\boldsymbol{\alpha}^\top \boldsymbol{X}(0)}}{(2\pi)^2} \int_{\mathbb{R}^2} e^{i\boldsymbol{u}^\top \boldsymbol{X}(0)} M_{\boldsymbol{X}_T}(\boldsymbol{\alpha} + i\boldsymbol{u}) \hat{f}(i\boldsymbol{\alpha} - \boldsymbol{u}) d\boldsymbol{u},$$

where $i = \sqrt{-1}$, $\boldsymbol{u} = [u_1, u_2] \in \mathbb{R}^2$, and $\hat{f}(\cdot)$ denotes the Fourier transform of the payoff function.

See Eberlein et al. (2010) for the proof.

Using Theorem 1, Eberlein et al. (2010) show that the Fourier valuation formula for a "worst-of-2" call option at $t = 0$ with strike K and maturity T is given by:

$$V_{min}(x_1(0), x_2(0)) = \frac{e^{-rT}}{4\pi^2} \int_{\mathbb{R}^2} e^{x_1(0)(\alpha_1+iu_1)} e^{x_2(0)(\alpha_2+iu_2)} M_{\boldsymbol{X}_T}(\alpha_1 + iu_1, \alpha_2 + iu_2) \times \frac{K^{1-\alpha_1-\alpha_2-iu_1-iu_2}}{(\alpha_1 + iu_1)(\alpha_2 + iu_2)(\alpha_1 + \alpha_2 - 1 + iu_1 + iu_2)} du. \tag{3.3.3}$$

Assuming the moment generating function exists, the relationship between the moment generating function $M_{\boldsymbol{X}_T}$ and the characteristic function $\phi_{\boldsymbol{X}_T}$ is given by:

$$M_{\boldsymbol{X}_T}(\boldsymbol{u}) = \phi_{\boldsymbol{X}_T}(-i\boldsymbol{u}),$$

hence, the moment generating function in Eq. (3.3.3) can be replaced by $\phi_{\boldsymbol{X}_T}$.

In the next section, we introduce the characteristic functions for the two-factor gBm and three-factor stochastic volatility models.

3.3.3 Two-Factor gBm Characteristic Function

The characteristic function represents the joint distribution of $\boldsymbol{X}(T)$ at the option maturity T. From Dempster and Hong (2002), the characteristic function for the two-asset gBm model is given by:

$$\phi_{gBm}(u_1, u_2) = \exp\Big(iu_1x_1(0) + iu_2x_2(0) + \zeta T + \sum_{j=1,2} u_j(r - \delta_j)T\Big),$$

where

$$\zeta := -\frac{1}{2}\Big[(\sigma_1^2u_1^2 + \sigma_2^2u_2^2 + 2\rho_{x_1,x_2}\sigma_1\sigma_2u_1u_2) + i(\sigma_1^2u_1 + \sigma_2^2u_2)\Big].$$

The characteristic function can be used as input to Eq. (3.3.3) to price "worst-of-2" call options under the two-factor gBm model.

As mentioned, the two-factor gBm model is not consistent with the behaviour of the equity market. Therefore, we consider the three-factor stochastic volatility model next.

3.3.4 Three-Factor Stochastic Volatility Characteristic Function

Dempster and Hong (2002) derived an expression for the characteristic function of the three-factor stochastic volatility model where both assets are driven by a single Cox et al. (1985) variance process. The expression for the characteristic function is given by:

$$\phi_{SV}(u_1,u_2) = \exp\left(iu_1x_1(0) + iu_2x_2(0) + \left(\frac{2\zeta(1-e^{-\beta T})}{2\beta-(\beta-\gamma)(1-e^{-\beta T})}\right)v(0) + \sum_{j=1,2} u_j(r-\delta_j)T - \frac{\kappa\bar{v}}{\sigma_v^2}\Gamma\right),$$

where

$$\begin{aligned}
\Gamma &:= \left[2\log\left(\frac{2\beta-(\beta-\gamma)(1-e^{-\beta T})}{2\beta}\right) + (\beta-\gamma)T\right],\\
\zeta &:= -\frac{1}{2}\Big[(\sigma_1^2u_1^2 + \sigma_2^2u_2^2 + 2\rho_{x_1,x_2}\sigma_1\sigma_2u_1u_2) + i(\sigma_1^2u_1 + \sigma_2^2u_2)\Big],\\
\gamma &:= \kappa - i(\rho_{x_1,v}\sigma_1u_1 + \rho_{x_2,v}\sigma_2u_2)\sigma_v,\\
\beta &:= \sqrt{\gamma^2 - 2\sigma_v^2\zeta}.
\end{aligned}$$

In the next section, we introduce the two-dimensional FFT method of Hurd and Zhou (2010).

3.4 The Two-Dimensional FFT Method

Hurd and Zhou (2010) initially developed the two-dimensional FFT method to price spread options. Using Eq. (3.3.3) and applying the same logic to "worst-of-2" call options with $K=1$ and maturity T, the Fourier representation for the "worst-of-2" call option with payoff function $V_{min}(x_1(T),x_2(T)) = \max\left(\min(e^{x_1(T)},e^{x_2(T)})-1,0\right)$ is given by:

$$V_{min}(x_1(0),x_2(0)) = \frac{e^{-rT}}{4\pi^2}\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} e^{x_1(0)(\alpha_1+iu_1)}e^{x_2(0)(\alpha_2+iu_2)}\phi_{\{gBm,SV\}}(u_1-i\alpha_1,u_2-i\alpha_2) \times \frac{1^{1-\alpha_1-\alpha_2-iu_1-iu_2}}{(\alpha_1+iu_1)(\alpha_2+iu_2)(\alpha_1+\alpha_2-1+iu_1+iu_2)}d\boldsymbol{u}, \tag{3.4.1}$$

where $\alpha_1,\alpha_2>0$, $\phi_{\{gBm,SV\}}$ is the characteristic function under either the two-factor gBm or three-factor stochastic volatility model, and $\epsilon_1,\epsilon_2<0$ with $\epsilon_1+\epsilon_2<-1$. The parameters ϵ_1,ϵ_2

can be chosen freely within their given constraints. Roberts (2018) showed that "worst-of-2" call option prices are insensitive to the choice of ϵ_1 and ϵ_2. Furthermore, let

$$\hat{V}_{min}(u_1, u_2) := \frac{1^{1-\alpha_1-\alpha_2-iu_1-iu_2}}{(\alpha_1 + iu_1)(\alpha_2 + iu_2)(\alpha_1 + \alpha_2 - 1 + iu_1 + iu_2)}.$$

The integral in Eq. (3.4.1) can be approximated as follows. Let

$$\Gamma = \{\boldsymbol{u}(\boldsymbol{k}) = (u_1(k_1), u_2(k_2)) \mid \boldsymbol{k} = (k_1, k_2) \in \{0, 1, ..., N-1\}^2\},\ u_i(k_i) = -\bar{u} + k_i\xi,$$

where $N = 2^m$ with $m \geq 0$, ξ is the lattice spacing, and $\bar{u} = \frac{N\xi}{2}$.

Furthermore, let the reciprocal lattice be given by:

$$\Gamma^* = \{\boldsymbol{x}(\boldsymbol{l}) = (x_1(l_1), x_2(l_2)) \mid \boldsymbol{l} = (l_1, l_2) \in \{0, 1, ..., N-1\}^2\},\ x_i(l_i) = -\bar{x} + l_i\xi^*,$$

where $\xi^* = \frac{\pi}{\bar{u}}$ is the reciprocal lattice spacing and $\bar{x} = \frac{N\xi^*}{2}$.

The integral in Eq. (3.4.1) is approximated by the following double sum for each pair $(x_1(l_1), x_2(l_2))$ in Γ^*:

$$V_{min}(x_1(l_1), x_2(l_2)) \approx \frac{e^{-rT}}{4\pi^2} \sum_{k_1=0}^{N-1} \sum_{k_2=0}^{N-1} e^{x_1(l_1)(\alpha_1 + i(u_1(k_1) + i\epsilon_1))} e^{x_2(l_2)(\alpha_2 + i(u_1(k_1) + i\epsilon_1))}$$

$$\times \phi_{\{gBm,SV\}}(u_1(k_1) + i\epsilon_1 - i\alpha_1, u_2(k_2) + i\epsilon_2 - i\alpha_2)\hat{V}_{min}(u_1(k_1) + i\epsilon_1, u_2(k_2) + i\epsilon_2)$$

$$= (-1)^{l_1+l_2} e^{-rT} \left(\frac{\xi N}{2\pi}\right)^2 e^{(\alpha_1 - \epsilon_1)x_1(l_1) + (\alpha_2 - \epsilon_2)x_1(l_2)} [\text{ifft2}(H)](l_1, l_2),$$

where

$$H(k_1, k_2) = (-1)^{k_1+k_2} \phi_{\{gBm,SV\}}(u_1(k_1) + i\epsilon_1 - i\alpha_1, u_2(k_2) + i\epsilon_2 - i\alpha_2)$$
$$\times \hat{V}_{min}(u_1(k_1) + i\epsilon_1, u_2(k_2) + i\epsilon_2),$$

and $\text{ifft2}(H)$ represents the two-dimensional FFT applied to matrix H.

Recall that Eq. (3.4.1) is for the specific case where $K = 1$. Roberts (2018) explains that Eq. (3.4.1) can be generalised for any $K > 0$ by scaling the initial share prices:

$$\boldsymbol{x}(0) = \left(\log\left(\frac{S_1(0)}{K}\right), \log\left(\frac{S_2(0)}{K}\right)\right).$$

To determine the option price, a grid search is done to find $\boldsymbol{x}(0)$ in Γ^*. If $\boldsymbol{x}(0)$ does not perfectly fall on the lattice Γ^*, an interpolation scheme must be used to find $\boldsymbol{x}(0)$.

The implementation of the two-dimensional FFT method is outlined in Alfeus and Schlögl (2018). Algorithm 1 below follows directly from their paper.

Algorithm 1 Two-dimensional FFT for rainbow options

---

1. **Input :** N, a power of two; $\bar{u}$, truncation width; $\boldsymbol{\epsilon}$, damping factor.
2. Set $\boldsymbol{x}(0) = \left(\log\left(\frac{S_1(0)}{K}\right), \log\left(\frac{S_2(0)}{K}\right)\right) \in (x_1(l_1), x_2(l_2))$.
3. **for all** $\boldsymbol{k}, \boldsymbol{l} \in \{1, 2, ..., N-1\}^2$ **do**

$$H(k_1, k_2) = (-1)^{k_1+k_2}\phi_{\{gBm,SV\}}(u_1(k_1) + i\epsilon_1 - \alpha_1, u_2(k_2) + i\epsilon_2 - \alpha_2)$$
$$\times \hat{V}_{min}(u_1(k_1) + i\epsilon_1, u_2(k_1) + i\epsilon_1);$$

$$C(l_1, l_2) = (-1)^{l_1+l_2}\left(\frac{\xi N}{2\pi}\right)^2 e^{(\alpha_1-\epsilon_1)x_1(l_1)+(\alpha_2-\epsilon_2)x_2(l_2)};$$

4. **end**
5. $V_{min}(x_1(l_1), x_2(l_2)) = \Re(C \times \text{ifft2}(H))$ where $\Re(\cdot)$ denotes the real part of the complex number.
6. $P \leftarrow K \times V_{min}(\boldsymbol{x}(0))$ using an interpolation scheme to find $\boldsymbol{x}(0)$ in Γ^*.

Output : P.

In the next section, we present our numerical results for "worst-of-2" call options based on the FFT method of Hurd and Zhou (2010) and various dynamics for the underlying assets.

3.5 Numerical Results

In this section, we first compare our pricing results of the FFT and two-factor gBm model with the results in Roberts (2018). We also test various values for the truncation width in the FFT algorithm to see whether an optimal value exists. Thereafter, we test the accuracy of the FFT under the three-factor stochastic volatility by comparing the FFT prices to the prices obtained from a Monte Carlo simulation.

3.5.1 FFT and Two-Factor gBm

Table 3.1 below shows our results for "worst-of-2" call options using the FFT and two-factor gBm model. The prices are compared to the prices obtained from the Stulz (1982) formula, as

shown in Roberts (2018). Furthermore, we use the exact same values for N and $\bar{u}$ as Roberts (2018) to compare the accuracy of our implementation.

Tab. 3.1: FFT prices under two-factor gBm with $S_1(0) = 100,\ S_2(0) = 96,\ \delta_1 = 0.05,\ \delta_2 = 0.05,\ r = 0.1,\ \sigma_1 = 0.1,\ \sigma_2 = 0.2,\ \rho_{x_1,x_2} = 0.5,\ \epsilon_1 = -3,\ \epsilon_2 = -1,\ \alpha_1 = 0.75,\ \alpha_2 = 0.75,\ T = 1.$

K	Stulz	$N = 64$ $\bar{u} = 30$	$N = 128$ $\bar{u} = 50$	$N = 256$ $\bar{u} = 90$	$N = 512$ $\bar{u} = 120$	$N = 1024$ $\bar{u} = 130$
90	8.274176	8.276002	8.273253	8.274178	8.274173	8.274158
92	7.118883	7.122826	7.177980	7.118794	7.118871	7.118862
94	6.055238	6.041106	6.055158	6.055197	6.055214	6.055220
96	5.087925	5.072448	5.087190	5.087892	5.087903	5.087914
98	4.220092	4.210246	4.219182	4.220043	4.220087	4.220090
100	3.452949	3.451947	3.452913	3.452951	3.452948	3.452948
102	2.785485	2.791147	2.785838	2.785476	2.785481	2.785482
104	2.214392	2.222128	2.214398	2.214413	2.214401	2.214396

Table 3.2 below shows the absolute difference between our FFT prices and the Stulz (1982) prices across strike:

Tab. 3.2: Absolute difference between FFT and Stulz prices

K	$N = 64$ $\bar{u} = 30$	$N = 128$ $\bar{u} = 50$	$N = 256$ $\bar{u} = 90$	$N = 512$ $\bar{u} = 120$	$N = 1024$ $\bar{u} = 130$
90	0.001826	0.000923	0.000002	0.000003	0.000018
92	0.006057	0.000903	0.000089	0.000012	0.000021
94	0.014132	0.000080	0.000041	0.000024	0.000018
96	0.015477	0.000735	0.000033	0.000022	0.000011
98	0.009846	0.000910	0.000049	0.000005	0.000002
100	0.001002	0.000036	0.000002	0.000001	0.000001
102	0.005662	0.000353	0.000009	0.000004	0.000003
104	0.007736	0.000006	0.000021	0.000009	0.000004

Our results do not support the findings in Roberts (2018). Firstly, our prices are noticeably more accurate than the prices in Roberts (2018). Based on the author's implementation, the FFT did not converge to three decimal places for $K = 90$. However, we find that the FFT converges for $K = 90$ up to five decimals using 256 or 512 terms. Roberts (2018) recommends that further investigation be done to conclude the accuracy of the FFT applied to further in-the-money options. Based on our implementation, there is nothing that suggests the FFT method is less accurate for further in-the-money options.

Another interesting observation is that our FFT results converge faster than the results in Roberts (2018). The author mentioned that all FFT prices apart from $K = 90$ converged to three decimal places from 512 terms. Based on our implementation, we find that the FFT converges to at least three decimals using 256 terms for all strikes. From 512 terms, the FFT converges to at least four decimals for all strikes.

Lastly, Roberts (2018) stated that it is preferable to use the Fourier-cosine series expansion (COS) method of Ruijter and Oosterlee (2012) to price rainbow options since it is faster and more robust than the FFT method. In terms of accuracy, our FFT implementation is very much aligned with the COS pricing results in Roberts (2018) (and in many cases even more accurate). This leads us to believe that there is an error in the FFT implementation in Roberts (2018). Upon further investigation, Roberts (2018) mentioned that the dampening factors in the FFT method should be restricted to $\alpha_2 < 0$ and $\alpha_1 + \alpha_2 > 1$, where, in fact, the dampening factors should be restricted to $\alpha_1, \alpha_2 > 0$ and $\alpha_1 + \alpha_2 > 1$ as per Eberlein et al. (2010). The question remains whether it is truly better to use the COS method rather than the FFT method for "worst-of-2" call options.

3.5.2 Testing the Truncation Width

The FFT algorithm of Hurd and Zhou (2010) requires an appropriate choice for the truncation width $\bar{u}$. Figure 3.1 below shows the FFT price for a "worst-of-2" call option with $K = 98$ and $T = 1$ as a function of $\bar{u}$:

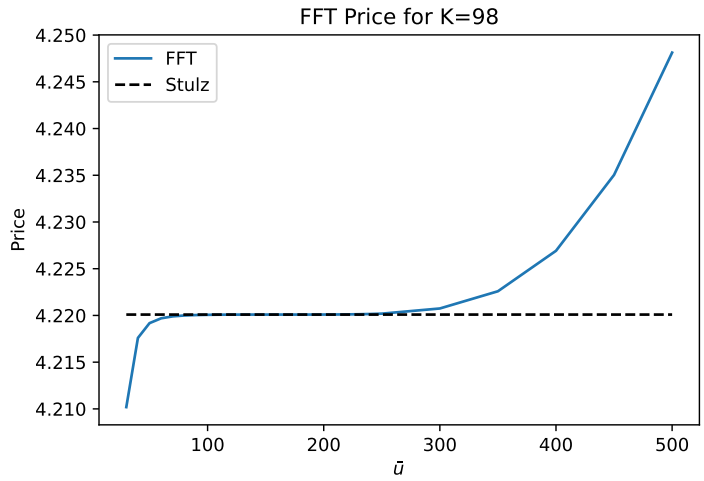

Fig. 3.1: Two-factor gBm FFT Price as a function of truncation width $\bar{u}$ with $N = 512$

From Figure 3.1, it is clear that the choice of $\bar{u}$ can have a significant impact on the FFT price. If $\bar{u}$ is chosen too small, the FFT price will be understated. Alternatively, choosing a value that is too large will overestimate the price.

Based on our implementation, the FFT converges to at least four decimal places for $\bar{u} \in [90, 200]$ for $K = 98$ and $N = 512$ compared to the Stulz (1982) price. Roberts (2018) mentioned that it is not possible to choose an optimal value for $\bar{u}$. However, we have shown that a value of $N = 512$ and $\bar{u} \in [90, 200]$ will be sufficient to achieve convergence up to four decimals under the two-factor gBm model. In Table 3.1, for $N = 512$, we used a value of $\bar{u} = 120$, which falls in the interval $[90, 200]$. All FFT prices converged to at least four decimal places, which further supports our finding.

An area where the COS method does seem to outperform the FFT is the rate of convergence. Based on the results in Roberts (2018), the COS method converged to three decimal places using as few as 64 terms. Based on our implementation of the FFT, convergence to three decimals was only achieved from 256 terms. Although the COS method does seem to be faster, there is not much that differentiates the pricing results between the COS and FFT methods. Therefore, if speed is an important factor, the COS method is preferable to the FFT method which aligns with the results in Roberts (2018). However, it takes only 3 seconds to price a single option when using the FFT method with 256 terms.

In summary, it is possible to choose an optimal value for $\bar{u}$ to achieve a certain level of convergence to the Stulz (1982) price. Next, we implement the three-factor stochastic volatility model of Dempster and Hong (2002).

3.5.3 FFT and Three-Factor Stochastic Volatility

In this subsection, we compare the FFT prices for "worst-of-2" call options under the three-factor stochastic volatility model of Dempster and Hong (2002) with a Monte Carlo simulation and show the convergence speed of the FFT by varying the number of terms N.

To test the accuracy of the FFT, Table 3.3 below compares the FFT pricing results for "worst-of-2" call options based on the three-factor stochastic volatility model with a Monte Carlo simulation of 20,000,000 samples for various K, where the Monte Carlo simulation is seen as the benchmark price. We show the results up to 6 decimal places.

Tab. 3.3: MC and FFT prices under three-factor stochastic volatility with $S_1(0) = 100$, $S_2(0) = 96$, $\delta_1 = 0.05$, $\delta_2 = 0.05$, $r = 0.1$, $\sigma_1 = 1$, $\sigma_2 = 0.5$, $v(0) = 0.04$, $\kappa = 1$, $\bar{v} = 0.04$, $\sigma_v = 0.05$, $\rho_{x_1,x_2} = 0.5$, $\rho_{x_1,v} = -0.5$, $\rho_{x_2,v} = 0.25$, $\epsilon_1 = -3$, $\epsilon_2 = -1$, $\alpha_1 = 0.75$, $\alpha_2 = 0.75$, $T = 1$, $N = 512$, $\bar{u} = 100$.

K	MC Price	FFT Price	Absolute Difference
90	7.642450	7.642304	0.000145
92	6.436579	6.436327	0.000252
94	5.340944	5.340803	0.000141
96	4.363270	4.363219	0.000052
98	3.507569	3.507650	0.000081
100	2.773671	2.773815	0.000141
102	2.157198	2.157236	0.000038
104	1.650093	1.650149	0.000056

The results in Table 3.3 indicate that the FFT and Monte Carlo prices are aligned to at least three decimal places for all strikes considered which confirms the accuracy of the FFT.

To illustrate the FFT convergence, Figure 3.2 below plots the logarithm of the absolute difference for a "worst-of-2" call option with $K = 98$ and $T = 1$ as a function of N, where $N = 8196$ is the benchmark price. The results are based on the same parameters shown in Table 3.3:

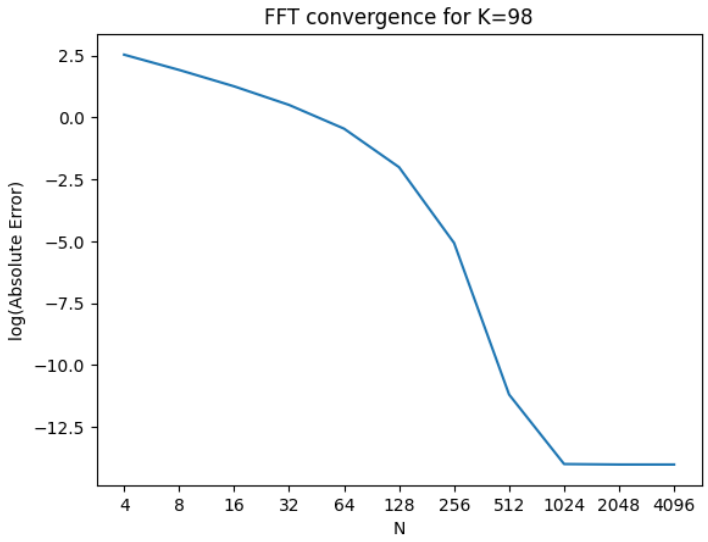

Fig. 3.2: FFT convergence under the three-factor stochastic volatility model

Figure 3.2 shows that the FFT converges at an exponential rate. Monte Carlo simulation, on the other hand, is known to converge at a rate of $O(\sqrt{N})$ - to achieve a tenfold increase in accuracy,

a hundredfold increase in the number of simulations is required (see Glasserman, 2003). FFT is far superior to Monte Carlo simulation in terms of efficiency.

In the next section, we conclude the chapter.

3.6 Conclusion

In this chapter, we implemented the two-factor gBm model and three-factor stochastic volatility model of Dempster and Hong (2002) to price "worst-of-2" call options. Using the two-dimensional FFT method of Hurd and Zhou (2010), we showed that it is possible to achieve a certain level of convergence to the Stulz (1982) and Monte Carlo prices under the two-factor gBm and three-factor stochastic volatility models respectively. We also showed that an optimal value for the truncation width can be chosen which contradicts a previous finding in the literature.

The FFT gives practitioners a powerful way of calculating rainbow option prices without the need to perform a costly Monte Carlo simulation. As rainbow option trading becomes more popular, the FFT method is likely to shine.

4 Efficient Pricing of Spread Options with Stochastic Rates and Stochastic Volatility

> *Volatility is a symptom that people have no clue of the underlying value.*
>
> **— Jeremy Grantham**

Keywords: Spread option · two-asset Heston-Hull-White model · discounted characteristic function · fast Fourier transform · stochastic interest rates

4.1 Introduction

Let S_1, S_2 denote the prices of two different equities, T the maturity date, K the strike price, and r the risk-free interest rate. Let $\mathbb{Q}$ be the risk-neutral measure associated with the bank account as numeraire, $B(t) = e^{\int_0^t r(s)ds}$. The value of a European equity spread call option at $t = 0$ is then given by:

$$V_{Spread}(0) = \mathbb{E}_{\mathbb{Q}}\Big[e^{-\int_0^T r(s)ds} \max\Big(S_1(T) - S_2(T) - K, 0\Big)\Big]. \tag{4.1.1}$$

Eq. (4.1.1) has no closed-form solution when the underlying assets are driven by stochastic interest rates and stochastic volatility. Monte Carlo simulation is one possible way of solving the problem, but it is too slow for practical use. This chapter[1] proposes an efficient numerical method to solve Eq. (4.1.1) when the state variables are driven by stochastic volatility and stochastic interest rates.

A breakthrough was made by Dempster and Hong (2002) where the authors extended the Carr and Madan (1999) Fourier transform method to two factors. The authors derived lower and upper bound expressions for the price of a European spread call option and showed how to

[1] This chapter is based on a paper (Levendis and Maré, 2022b) published in the *Journal of Risk and Financial Management*.

compute these expressions by means of the two-dimensional FFT when the underlying asset processes are driven by gBm or stochastic volatility.

Hurd and Zhou (2010) proposed a method for pricing spread options based on a square integrable Fourier representation of the payoff function. The authors claim that the method can be applied to any model for which the characteristic function of the joint asset process is known in closed form. The authors considered gBm, three-factor stochastic volatility, and exponential Lévy models for the underlying asset processes.

Little has been said about the impact of stochastic interest rates on spread options. We do, however, know that stochastic interest rate risk dominates that of stochastic volatility for long-dated European call and put options (see Kammeyer & Kienitz, 2012). Furthermore, there is empirical evidence that suggests changes in interest rates and changes in stock prices are negatively correlated in general (Alam & Uddin, 2009). None of the models considered by Dempster and Hong (2002) or Hurd and Zhou (2010) account for stochastic interest rates or the correlation between interest rates and stock prices.

Grzelak and Oosterlee (2011) made a remarkable contribution to the literature, which derived approximations for the characteristic function of the Heston (1993) model with stochastic interest rates driven by a Hull and White (1990) process. This model is called the Heston–Hull–White model. Through these approximations, the model allows for a full matrix of correlations between stock, volatility, and interest rate processes. The authors showed that the model can be calibrated to European options efficiently via Fourier techniques.

In this chapter, we extend the Heston–Hull–White model of Grzelak and Oosterlee (2011) to two assets and the FFT method of Hurd and Zhou (2010) to stochastic interest rates. We then compare the efficiency of the extended Hurd and Zhou (2010) method to a Monte Carlo simulation and assess the impact of stochastic interest rates on spread option prices.

The rest of this chapter is organised as follows: The two-asset Heston–Hull–White model is introduced in Section 4.2. In Section 4.3, we discuss the result of Grzelak and Oosterlee (2011) that will be used to derive the discounted characteristic function. The discounted characteristic function for the two-asset Heston-Hull–White model is derived in Section 4.4. The extension of the Hurd and Zhou (2010) FFT algorithm to stochastic interest rates is detailed in Section 4.5. Our numerical results are shown in Section 4.6, and Section 4.7 concludes the chapter.

4.2 The Two-Asset Heston–Hull–White Model

Combining the three-factor stochastic volatility model of Dempster and Hong (2002) and the Heston–Hull–White model of Grzelak and Oosterlee (2011) yields:

$$\begin{cases} dS_1(t) = (r(t) - \delta_1)S_1(t)dt + \sigma_1\sqrt{v(t)}S_1(t)dW_{x_1}(t), \\ dS_2(t) = (r(t) - \delta_2)S_2(t)dt + \sigma_2\sqrt{v(t)}S_2(t)dW_{x_2}(t), \\ dv(t) = \kappa\Big(\bar{v} - v(t)\Big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dr(t) = \lambda\Big(\theta(t) - r(t)\Big)dt + \eta dW_r(t), \end{cases} \tag{4.2.1}$$

where σ_1, σ_2 denote the volatility of S_1, S_2, respectively; δ_1, δ_2 are the dividend yields for S_1, S_2; κ and λ denote the mean reversion speed of the variance and short rate processes; σ_v and η denote the volatility of the volatility and short rate processes, respectively, and $\bar{v}$ and $\theta(t)$ denote the mean level of the variance and short rate, respectively.

The underlying process are correlated as follows:

$$\begin{aligned} dW_{x_1}(t)dW_{x_2}(t) &= \rho_{x_1,x_2}dt, \quad & dW_{x_1}(t)dW_v(t) &= \rho_{x_1,v}dt, \\ dW_{x_1}(t)dW_r(t) &= \rho_{x_1,r}dt, \quad & dW_{x_2}(t)dW_v(t) &= \rho_{x_2,v}dt, \\ dW_{x_2}(t)dW_r(t) &= \rho_{x_2,r}dt, \quad & dW_v(t)dW_r(t) &= \rho_{v,r}dt. \end{aligned}$$

Taking the log-transform $x_1(t) = \log S_1(t)$, $x_2(t) = \log S_2(t)$ and applying Itô's lemma, the system of stochastic differential equations (SDEs) in Eq. (4.2.1) can be rewritten as:

$$\begin{cases} dx_1(t) = \Big(r(t) - \delta_1 - \frac{1}{2}\sigma_1^2 v(t)\Big)dt + \sigma_1\sqrt{v(t)}dW_{x_1}(t), \\ dx_2(t) = \Big(r(t) - \delta_2 - \frac{1}{2}\sigma_2^2 v(t)\Big)dt + \sigma_2\sqrt{v(t)}dW_{x_2}(t), \\ dv(t) = \kappa\Big(\bar{v} - v(t)\Big)dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dr(t) = \lambda\Big(\theta(t) - r(t)\Big)dt + \eta dW_r(t). \end{cases} \tag{4.2.2}$$

The system of SDEs in Eq. (4.2.2) can be expressed as:

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

where

$$\mathbf{X}(t) = \begin{bmatrix} x_1(t) \\ x_2(t) \\ v(t) \\ r(t) \end{bmatrix},$$

$$\mu(\mathbf{X}(t)) = \begin{bmatrix} r(t) - \delta_1 - \frac{1}{2}\sigma_1^2 v(t) \\ r(t) - \delta_2 - \frac{1}{2}\sigma_2^2 v(t) \\ \kappa\Big(\bar{v} - v(t)\Big) \\ \lambda\Big(\theta(t) - r(t)\Big) \end{bmatrix},$$

$$\mathbf{B}(t) = \begin{bmatrix} B_{x_1}(t) \\ B_{x_2}(t) \\ B_v(t) \\ B_r(t) \end{bmatrix},$$

$$\begin{aligned} \mathbf{\Sigma}(\mathbf{X}(t)) &= \sigma(\mathbf{X}(t))\sigma(\mathbf{X}(t))^T \\ &= \begin{bmatrix} \sigma_1^2 v(t) & \rho_{x_1,x_2}\sigma_1\sigma_2 v(t) & \rho_{x_1,v}\sigma_1\sigma_v v(t) & \rho_{x_1,r}\sigma_1\eta\sqrt{v(t)} \\ * & \sigma_2^2 v(t) & \rho_{x_2,v}\sigma_2\sigma_v v(t) & \rho_{x_2,r}\sigma_2\eta\sqrt{v(t)} \\ * & * & \sigma_v^2 v(t) & \rho_{v,r}\sigma_v\eta\sqrt{v(t)} \\ * & * & * & \eta^2 \end{bmatrix}, \end{aligned}$$

with $\mathbf{B}(t)$ a vector of independent Brownian motions and $\sigma(\mathbf{X}(t))$ the Cholesky decomposition of the symmetric covariance matrix $\mathbf{\Sigma}(\mathbf{X}(t))$.

The influential paper by Duffie et al. (2000) states that each element in the drift and covariance matrices must be a linear function of the state variables in $\mathbf{X}(t)$ in order for a model to be in affine form. If this is the case, then the discounted characteristic function for the state vector $\mathbf{X}(t)$ can be written as:

$$\begin{aligned} \phi(\mathbf{u}, \mathbf{X}(t), t, T) &= \mathbb{E}_{\mathbb{Q}}\Big[e^{-\int_t^T r(s)ds + i\mathbf{u}^T\mathbf{X}(T)} \mid \mathcal{F}(t)\Big] \\ &= e^{A(\mathbf{u},\tau)+B(\mathbf{u},\tau)x_1(t)+C(\mathbf{u},\tau)x_2(t)+D(\mathbf{u},\tau)v(t)+E(\mathbf{u},\tau)r(t)}, \end{aligned} \tag{4.2.3}$$

where $\tau := T - t$, and $\mathbf{u} = [u_1, u_2, 0, 0]^T$.

From the covariance matrix $\mathbf{\Sigma}(\mathbf{X}(t))$, it is clear that there are elements that are nonlinear functions of the state variables, in particular $\sqrt{v(t)}$.

Grzelak and Oosterlee (2011) proposed replacing the term $\sqrt{v(t)}$ with $\mathbb{E}[\sqrt{v(t)}]$ so that the Heston–Hull–White model could be expressed in affine form as in Eq. (4.2.3). Their result follows in the next section.

4.3 The Result of Grzelak and Oosterlee

In order to write the Heston–Hull–White model in affine form, Grzelak and Oosterlee (2011) proposed the following approximation for $\sqrt{v(t)}$:

Lemma 1 *Approximation for $\mathbb{E}[\sqrt{v(t)}]$. Given that $v(t)$ follows a Cox et al. (1985) process, $\mathbb{E}[\sqrt{v(t)}]$ can be approximated by:*

$$\mathbb{E}\Big[\sqrt{v(t)}\Big] \approx \sqrt{c(t)(\lambda(t)-1)+c(t)d+\frac{c(t)d}{2(d+\lambda(t))}} =: \Lambda(t),$$

where $c(t)=\frac{1}{4\kappa}\sigma_v^2(1-e^{-\kappa t})$, $d=\frac{4\kappa\bar{v}}{\sigma_v^2}$, and $\lambda(t)=\frac{4\kappa v(0)e^{-\kappa t}}{\sigma_v^2(1-e^{-\kappa t})}$.

See Grzelak and Oosterlee (2011) for details.

The authors mention that the approximation for $\mathbb{E}[\sqrt{v(t)}]$ is still nontrivial and may lead to challenges when deriving the characteristic function for the Heston–Hull–White model. Therefore, a further simplified approximation for $\mathbb{E}[\sqrt{v(t)}]$ was proposed as shown in Lemma 2 below:

Lemma 2 *Further approximation for $\mathbb{E}[\sqrt{v(t)}]$. $\mathbb{E}[\sqrt{v(t)}]$ can be further approximated by a function of the form:*

$$\mathbb{E}\Big[\sqrt{v(t)}\Big] \approx a+be^{-ct} =: \tilde{\Lambda}(t),$$

where $a=\sqrt{\bar{v}-\frac{\sigma_v^2}{8\kappa}}$, $b=\sqrt{v(0)}-a$, and $c=-\log(b^{-1}(\Lambda(1)-a))$.

See Grzelak and Oosterlee (2011) for details.

Note that $\tilde{\Lambda}(t)$ in Lemma 2 is undefined for $\bar{v}<\sigma_v^2/8\kappa$; hence, parameter constraints must be imposed in order to use this result.

In the next section, we derive the discounted characteristic function for the two-asset Heston–Hull–White model using the approximation of Grzelak and Oosterlee (2011).

4.4 The Two-Asset Heston–Hull–White Characteristic Function

Replacing $\sqrt{v(t)}$ with $\mathbb{E}\left[\sqrt{v(t)}\right]$ in the covariance matrix $\mathbf{\Sigma}(\mathbf{X}(t))$ yields the approximated covariance matrix for the two-asset Heston–Hull–White model:

$$\tilde{\mathbf{\Sigma}}(\mathbf{X}(t)) = \begin{bmatrix} \sigma_1^2 v(t) & \rho_{x_1,x_2}\sigma_1\sigma_2 v(t) & \rho_{x_1,v}\sigma_1\sigma_v v(t) & \rho_{x_1,r}\sigma_1\eta\mathbb{E}\left[\sqrt{v(t)}\right] \\ * & \sigma_2^2 v(t) & \rho_{x_2,v}\sigma_2\sigma_v v(t) & \rho_{x_2,r}\sigma_2\eta\mathbb{E}\left[\sqrt{v(t)}\right] \\ * & * & \sigma_v^2 v(t) & \rho_{v,r}\sigma_v\eta\mathbb{E}\left[\sqrt{v(t)}\right] \\ * & * & * & \eta^2 \end{bmatrix}.$$

For the derivation of the discounted characteristic function, we drop the function arguments for convenience. Hence, $x_1 := x_1(t)$, $x_2 := x_2(t)$ $r := r(t)$, $v := v(t)$, and $\phi := \phi(\mathbf{u}, \mathbf{X}(t), t, T)$. Furthermore, to simplify the calculations, we consider a constant term-structure of interest rates $\theta(t) = \theta$. The method can be generalised to include a term-structure of interest rates, see Grzelak and Oosterlee (2011) for details.

Using the drift vector $\mu(\mathbf{X}(t))$ and the covariance matrix $\tilde{\mathbf{\Sigma}}(\mathbf{X}(t))$ and applying the multidimensional Itô lemma to $\phi(\mathbf{u}, \mathbf{X}(t), t, T)$ yields the partial differential equation (PDE):

$$\begin{aligned} 0 =& \frac{\partial\phi}{\partial t} + (r - \delta_1 - \frac{1}{2}\sigma_1^2 v)\frac{\partial\phi}{\partial x_1} + (r - \delta_2 - \frac{1}{2}\sigma_2^2 v)\frac{\partial\phi}{\partial x_2} + \kappa(\bar{v} - v)\frac{\partial\phi}{\partial v} + \lambda(\theta - r)\frac{\partial\phi}{\partial r} \\ &+ \frac{1}{2}\sigma_1^2 v\frac{\partial^2\phi}{\partial x_1^2} + \frac{1}{2}\sigma_2^2 v\frac{\partial^2\phi}{\partial x_2^2} + \frac{1}{2}\sigma_v^2 v\frac{\partial^2\phi}{\partial v^2} + \frac{1}{2}\eta^2\frac{\partial^2\phi}{\partial r^2} + \rho_{x_1,x_2}\sigma_1\sigma_2 v\frac{\partial^2\phi}{\partial x_1\partial x_2} \\ &+ \rho_{x_1,v}\sigma_1\sigma_v v\frac{\partial^2\phi}{\partial x_1\partial v} + \rho_{x_1,r}\sigma_1\eta\mathbb{E}\left[\sqrt{v}\right]\frac{\partial^2\phi}{\partial x_1\partial r} + \rho_{x_2,v}\sigma_2\sigma_v v\frac{\partial^2\phi}{\partial x_2\partial v} \\ &+ \rho_{x_2,r}\sigma_2\eta\mathbb{E}\left[\sqrt{v}\right]\frac{\partial^2\phi}{\partial x_2\partial r} + \rho_{v,r}\sigma_v\eta\mathbb{E}\left[\sqrt{v}\right]\frac{\partial^2\phi}{\partial v\partial r} - r\phi, \end{aligned} \tag{4.4.1}$$

subject to the terminal condition $\phi(\mathbf{u}, \mathbf{X}(T), T, T) = e^{i(u_1 x_1(T) + u_2 x_2(T))}$.

The PDE in Eq. (4.4.1) is in affine form as a consequence of the linearisation technique proposed by Grzelak and Oosterlee (2011). Therefore, its solution is of the form:

$$\phi(\mathbf{u}, \mathbf{X}(t), t, T) = e^{A(\mathbf{u},t,T) + B(\mathbf{u},t,T)x_1(t) + C(\mathbf{u},t,T)x_2(t) + D(\mathbf{u},t,T)v(t) + E(\mathbf{u},t,T)r(t)}.$$

Calculating the partial derivatives in Eq. (4.4.1) with $A := A(\mathbf{u}, t, T)$, $B := B(\mathbf{u}, t, T)$, $C := C(\mathbf{u}, t, T)$, $D := D(\mathbf{u}, t, T)$, and $E := E(\mathbf{u}, t, T)$ yields:

$$
\begin{aligned}
&\frac{\partial \phi}{\partial t} = \phi\left(\frac{\partial A}{\partial t} + x_1\frac{\partial B}{\partial x_1} + x_2\frac{\partial C}{\partial t} + v\frac{\partial D}{\partial v} + r\frac{\partial E}{\partial r}\right),\\
&\frac{\partial \phi}{\partial x_1} = B\phi, \qquad \frac{\partial \phi}{\partial x_1^2} = B^2\phi, \qquad \frac{\partial \phi}{\partial x_1 \partial x_2} = BC\phi,\\
&\frac{\partial \phi}{\partial x_2} = C\phi, \qquad \frac{\partial \phi}{\partial x_2^2} = C^2\phi, \qquad \frac{\partial \phi}{\partial x_1 \partial v} = BD\phi, \qquad \frac{\partial \phi}{\partial x_2 \partial r} = CE\phi,\\
&\frac{\partial \phi}{\partial v} = D\phi, \qquad \frac{\partial \phi}{\partial v^2} = D^2\phi, \qquad \frac{\partial \phi}{\partial x_1 \partial r} = BE\phi, \qquad \frac{\partial \phi}{\partial v \partial r} = DE\phi,\\
&\frac{\partial \phi}{\partial r} = E\phi, \qquad \frac{\partial \phi}{\partial r^2} = E^2\phi, \qquad \frac{\partial \phi}{\partial x_2 \partial v} = CD\phi.
\end{aligned}
$$

Substituting the partial derivatives into Eq. (4.4.1) gives the following PDE:

$$
\begin{aligned}
0 =& \frac{\partial A}{\partial t} + x_1\frac{\partial B}{\partial t} + x_2\frac{\partial C}{\partial t} + v\frac{\partial D}{\partial t} + r\frac{\partial E}{\partial t} + (r - \delta_1 - \frac{1}{2}\sigma_1^2 v)B + (r - \delta_2 - \frac{1}{2}\sigma_2^2 v)C\\
&+ \kappa(\bar{v} - v)D + \lambda(\theta - r)E + \frac{1}{2}\sigma_1^2 v B^2 + \frac{1}{2}\sigma_2^2 v C^2 + \frac{1}{2}\sigma_v^2 v D^2 + \frac{1}{2}\eta^2 E^2\\
&+ \rho_{x_1,x_2}\sigma_1\sigma_2 v BC + \rho_{x_1,v}\sigma_1\sigma_v v BD + \rho_{x_1,r}\sigma_1\eta\mathbb{E}\left[\sqrt{v}\right]BE\\
&+ \rho_{x_2,v}\sigma_2\sigma_v v CD + \rho_{x_2,r}\sigma_2\eta\mathbb{E}\left[\sqrt{v}\right]CE + \rho_{v,r}\sigma_v\eta\mathbb{E}\left[\sqrt{v}\right]DE - r.
\end{aligned}
$$

Collecting the terms for x_1, x_2, v, and r and performing a change of variable $\tau = T - t$, the following set of ordinary differential equations (ODEs) must be solved:

$$
\begin{cases}
\frac{\partial B}{\partial \tau} = & 0,\\
\frac{\partial C}{\partial \tau} = & 0,\\
\frac{\partial D}{\partial \tau} = & -\frac{1}{2}\sigma_1^2 B - \frac{1}{2}\sigma_2^2 C - \kappa D + \frac{1}{2}\sigma_1^2 B^2 + \frac{1}{2}\sigma_2^2 C^2 + \frac{1}{2}\sigma_v^2 D^2 + \rho_{x_1,x_2}\sigma_1\sigma_2 BC\\
& + \rho_{x_1,v}\sigma_1\sigma_v BD + \rho_{x_2,v}\sigma_2\sigma_v CD,\\
\frac{\partial E}{\partial \tau} = & B + C - \lambda E - 1,\\
\frac{\partial A}{\partial \tau} = & \kappa\bar{v}D + \lambda\theta E + \frac{1}{2}\eta^2 E^2 + \rho_{x_1,r}\sigma_1\eta\mathbb{E}[\sqrt{v}]BE + \rho_{x_2,r}\sigma_2\eta\mathbb{E}[\sqrt{v}]CE\\
& + \rho_{v,r}\sigma_v\eta\mathbb{E}[\sqrt{v}]DE - \delta_1 B - \delta_2 C,
\end{cases}
$$

with initial conditions $B(\mathbf{u}, \tau) = iu_1$, $C(\mathbf{u}, \tau) = iu_2$, $D(\mathbf{u}, \tau) = 0$, $E(\mathbf{u}, \tau) = 0$, and $A(\mathbf{u}, \tau) = 0$.

The solutions to the ODEs are given by:

$$
\begin{aligned}
B(\mathbf{u},\tau) &= iu_1,\\
C(\mathbf{u},\tau) &= iu_2,\\
D(\mathbf{u},\tau) &= \frac{-Q-D_1}{2R(1-Ge^{-D_1\tau})}(1-e^{-D_1\tau}),\\
E(\mathbf{u},\tau) &= (iu_1+iu_2-1)\lambda^{-1}(1-e^{-\lambda\tau}),\\
A(\mathbf{u},\tau) &= \kappa\bar{v}I_1(\mathbf{u},\tau)+\lambda\theta I_2(\mathbf{u},\tau)+\frac{1}{2}\eta^2 I_3(\mathbf{u},\tau)+\rho_{x_1,r}\sigma_1\eta I_4(\mathbf{u},\tau)\\
&\quad+\rho_{x_2,r}\sigma_2\eta I_5(\mathbf{u},\tau)+\rho_{v,r}\sigma_v\eta I_6(\mathbf{u},\tau)-\delta_1 iu_1\tau-\delta_2 iu_2\tau,
\end{aligned}
$$

where

$$
\begin{aligned}
D_1 &= \sqrt{Q^2-4PR},\\
G &= \frac{-Q-D_1}{-Q+D_1}\\
P &= -\frac{1}{2}\Big[\sigma_1^2u_1^2+\sigma_2^2u_2^2+2\rho_{x_1,x_2}\sigma_1\sigma_2u_1u_2+i(\sigma_1^2u_1+\sigma_2^2u_2)\Big]\\
Q &= \rho_{x_1,v}\sigma_1\sigma_v iu_1+\rho_{x_2,v}\sigma_2\sigma_v iu_2-\kappa\\
R &= \frac{1}{2}\sigma_v^2.
\end{aligned}
$$

The solutions to the integrals $I_1(\mathbf{u},\tau)$, $I_2(\mathbf{u},\tau)$, $I_3(\mathbf{u},\tau)$, $I_4(\mathbf{u},\tau)$, $I_5(\mathbf{u},\tau)$, and $I_6(\mathbf{u},\tau)$ are given by:

$$
\begin{aligned}
I_1(\mathbf{u},\tau) &= \frac{1}{2R}\left[(-Q-D_1)\tau-2\log\Big(\frac{1-Ge^{-D_1\tau}}{1-G}\Big)\right],\\
I_2(\mathbf{u},\tau) &= \frac{1}{\lambda}(iu_1+iu_2-1)(\tau+\frac{1}{\lambda}(e^{-\lambda\tau}-1)),\\
I_3(\mathbf{u},\tau) &= \frac{1}{2\lambda^3}(i+u_1+u_2)^2[3+e^{-2\lambda\tau}-4e^{-\lambda\tau}-2\lambda\tau],\\
I_4(\mathbf{u},\tau) &= -\frac{1}{\lambda}(iu_1+u_1u_2+u_1^2)\left[\frac{b}{c}(e^{-ct}-e^{-cT})+a\tau+\frac{a}{\lambda}(e^{-\lambda\tau}-1)\right.\\
&\quad\left.+\frac{b}{c-\lambda}e^{-cT}\Big(1-e^{-\tau(\lambda-c)}\Big)\right],\\
I_5(\mathbf{u},\tau) &= -\frac{1}{\lambda}(iu_2+u_1u_2+u_2^2)\left[\frac{b}{c}(e^{-ct}-e^{-cT})+a\tau+\frac{a}{\lambda}(e^{-\lambda\tau}-1)\right.\\
&\quad\left.+\frac{b}{c-\lambda}e^{-cT}\Big(1-e^{-\tau(\lambda-c)}\Big)\right],
\end{aligned}
$$

$$I_6(\mathbf{u},\tau) = \int_0^\tau \mathbb{E}\Big[\sqrt{v(T-s)}\Big] D(\mathbf{u},s)E(\mathbf{u},s),$$

where $\mathbb{E}\Big[\sqrt{v(T-s)}\Big] \approx a + b^{-c(T-s)}$, with a, b, and c defined in Lemma 2.

This concludes the derivation of the discounted characteristic function for the two-asset Heston–Hull–White model. In the next section, we extend the FFT method of Hurd and Zhou (2010) to cater for stochastic interest rates.

4.5 The Result of Hurd and Zhou (Extended)

Hurd and Zhou (2010) stated the following theorem for the square integrable Fourier representation of the basic spread option payoff function $F(x_1,x_2) = \max(e^{x_1} - e^{x_2} - 1, 0)$:

Theorem 1 *For any real numbers $\boldsymbol{\epsilon} = (\epsilon_1,\epsilon_2)$ with $\epsilon_2 > 0$, $\epsilon_1 + \epsilon_2 < -1$, and $\mathbf{x} = (x_1,x_2)$,*

$$F(\mathbf{x}) = \frac{1}{4\pi^2}\int\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} e^{i\mathbf{u}\mathbf{x}^\top}\hat{F}(u_1,u_2)du_1du_2,$$

$$\hat{F}(u_1,u_2) = \frac{\Gamma(i(u_1+u_2)-1)\Gamma(-iu_2)}{\Gamma(iu_1+1)},$$

where $\Gamma(z)$ is the complex gamma function defined for $\Re(z) > 0$ by the integral $\Gamma(z) = \int_0^\infty e^{-t}t^{z-1}dt$ and $\mathbf{u}\mathbf{x}^\top = u_1x_1 + u_2x_2$, where $\mathbf{x}^\top$ is the unconjugated transpose of $\mathbf{x}$.

See Hurd and Zhou (2010) for the proof.

Lemma 3 below is adapted from Hurd and Zhou (2010) to account for stochastic interest rates:

Lemma 3 *Let $N(T) = e^{-\int_0^T r(s)ds}$ and $\mathbf{x}(t) = [\log S_1(t), \log S_2(t)]$. For any $t > 0$, the increment $\mathbf{x}(t) - \mathbf{x}(0)$ is independent of $\mathbf{x}(0)$, which implies:*

$$\mathbb{E}_\mathbb{Q}\Big[N(T)e^{i\mathbf{u}\mathbf{x}(T)^\top}\Big] = e^{i\mathbf{u}\mathbf{x}(0)^\top}\phi(\mathbf{u};T),$$

with $\phi(\mathbf{u};T) := \mathbb{E}_\mathbb{Q}\Big[N(T)e^{i\mathbf{u}(\mathbf{x}(T)-\mathbf{x}(0))^\top}\Big]$.

See Hurd and Zhou (2010) for details.

Using Theorem 1 and Lemma 3, the price of a European spread call option with stochastic interest rates can be written as a two-dimensional Fourier transform in the variable $\mathbf{x}(0)$. The derivation is shown below. Consider the spread option payoff function:

$$V_{Spread}(0) = \mathbb{E}_{\mathbb{Q}}\Big[N(T)\max\Big(e^{x_1(T)} - e^{x_2(T)} - 1, 0\Big)\Big].$$

Changing from the risk-neutral measure $\mathbb{Q}$ to the T-Forward measure using the zero-coupon bond $P(0,T)$ as numeraire, we obtain:

$$\begin{aligned} V_{Spread}(0) &= P(0,T)\mathbb{E}_T\Big[\frac{1}{4\pi^2}\int\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} e^{i\mathbf{u}\mathbf{x}(T)^\top}\hat{F}(u_1,u_2)du_1du_2\Big] \\ &= \frac{1}{4\pi^2}P(0,T)\int\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} \mathbb{E}_T\Big[e^{i\mathbf{u}\mathbf{x}(T)^\top}\Big]\hat{F}(u_1,u_2)du_1du_2. \end{aligned}$$

Using the result $P(0,T)\mathbb{E}_T\Big[e^{i\mathbf{u}\mathbf{x}(T)^\top}\Big] = \mathbb{E}_{\mathbb{Q}}\Big[N(T)e^{i\mathbf{u}\mathbf{x}(T)^\top}\Big]$, the price of the spread option becomes:

$$\begin{aligned} V_{Spread}(0) = \frac{1}{4\pi^2}P(0,T)\Big\{\int\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} e^{i\mathbf{u}\mathbf{x}(0)^\top}\frac{1}{P(0,T)}\mathbb{E}_{\mathbb{Q}}\Big[N(T)e^{i\mathbf{u}(\mathbf{x}(T)-\mathbf{x}(0))^\top}\Big] \\ \times \hat{F}(u_1,u_2)du_1du_2\Big\} \\ = \frac{1}{4\pi^2}\int\int_{\mathbb{R}^2+i\boldsymbol{\epsilon}} e^{i\mathbf{u}\mathbf{x}(0)^\top}\phi(\mathbf{u};T)\hat{F}(u_1,u_2)du_1du_2. \end{aligned} \tag{4.5.1}$$

Eq. (4.5.1) is for the specific case where $K = 1$. Roberts (2018) shows that the result can be extended to $K > 0$ by scaling the two initial stock prices as follows:

$$V_{Spread}(S_1(0), S_2(0), K, T) = K \times V_{Spread}\left(\frac{S_1(0)}{K}, \frac{S_2(0)}{K}, 1, T\right). \tag{4.5.2}$$

The implementation of Eq. (4.5.2) using the FFT technique is outlined in Alfeus and Schlögl (2018). Algorithm 2 below follows directly from their paper.

Algorithm 2 Two-dimensional FFT for spread options

1. **Input :** N, a power of two; $\bar{u}$, truncation width; $\boldsymbol{\epsilon}$, damping factor.
2. Set $\boldsymbol{x}(0) = \left(\log\left(\frac{S_1(0)}{K}\right), \log\left(\frac{S_2(0)}{K}\right)\right) \in (x_1(l_1), x_2(l_2))$.
3. **for all** $\boldsymbol{k}, \boldsymbol{l} \in \{1, 2, ..., N-1\}^2$ **do**

$$H(k_1, k_2) = (-1)^{k_1+k_2}\phi(\mathbf{u}(\mathbf{k}) + i\boldsymbol{\epsilon})\hat{F}(\mathbf{u}(\mathbf{k}) + i\boldsymbol{\epsilon});$$

$$C(l_1, l_2) = (-1)^{l_1+l_2}\left(\frac{\xi N}{2\pi}\right)^2 e^{(\alpha_1-\epsilon_1)x_1(l_1)+(\alpha_2-\epsilon_2)x_2(l_2)};$$

4. **end**
5. $V_{Spread}(x_1(l_1), x_2(l_2)) = \Re(C \times \text{ifft2}(H))$ where $\Re(\cdot)$ denotes the real part of the complex number.
6. $P \leftarrow K \times V_{Spread}(\boldsymbol{x}(0))$ using an interpolation scheme to find $\boldsymbol{x}(0)$ in Γ^*.

Output : P.

In Algorithm 2, the double integral in Eq. (4.5.1) is approximated by a double sum over the lattice:

$$\Gamma = \{\boldsymbol{u}(\boldsymbol{k}) = (u_1(k_1), u_2(k_2)) \mid \boldsymbol{k} = (k_1, k_2) \in \{0, 1, ..., N-1\}^2\},\ u_i(k_i) = -\bar{u} + k_i\xi,$$

where $N = 2^m$ with $m \geq 0$, ξ is the lattice spacing, and $\bar{u} = \frac{N\xi}{2}$

Furthermore, $\mathbf{x}(0) = \left[\log\left(\frac{S_1(0)}{K}\right), \log\left(\frac{S_2(0)}{K}\right)\right]$ is chosen to lie on the reciprocal lattice:

$$\Gamma^* = \{\boldsymbol{x}(\boldsymbol{l}) = (x_1(l_1), x_2(l_2)) \mid \boldsymbol{l} = (l_1, l_2) \in \{0, 1, ..., N-1\}^2\},\ x_i(l_i) = -\bar{x} + l_i\xi^*,$$

where $\xi^* = \frac{\pi}{\bar{u}}$ is the reciprocal lattice spacing and $\bar{x} = \frac{N\xi^*}{2}$.

This concludes the extension of the Hurd and Zhou (2010) FFT algorithm to stochastic interest rates. In the next section, we test the accuracy of the two-asset Heston–Hull–White model and the impact of the stochastic interest rates on spread option prices.

4.6 Numerical Results

This section is divided into three parts. First, we compare our implementation of the Hurd and Zhou (2010) FFT algorithm with the results shown in their paper. Secondly, we compare the convergence of the FFT to the Monte Carlo simulation. Lastly, we assess the impact of stochastic rates on the spread option prices.

4.6.1 Implementation Testing

To test the accuracy of our model, we implemented the two-asset Heston-Hull-White model and two-dimensional FFT of Hurd and Zhou (2010) in Python.

The two-asset Heston–Hull–White model reduces to the three-factor stochastic volatility model of Dempster and Hong (2002) when $\eta = 0$. Table 4.1 below compares our implementation with the results published in Hurd and Zhou (2010):

Tab. 4.1: $S_1(0) = 100, S_2(0) = 96, \delta_1 = 0.05, \delta_2 = 0.05, v(0) = 0.04, r(0) = 0.1, \sigma_1 = 1.0, \sigma_2 = 0.5, \kappa = 1, \bar{v} = 0.04, \sigma_v = 0.05, \lambda = 1, \theta = 0.1, \eta = 0, \rho_{x_1,x_2} = 0.5, \rho_{x_1,v} = -0.5, \rho_{x_1,r} = 0, \rho_{x_2,v} = 0.25, \rho_{x_2,r} = 0, \rho_{v,r} = 0, N = 256, \bar{u} = 40, \epsilon_1 = -3, \epsilon_2 = 1, T = 1.$

Strike	Hurd and Zhou Price	Model Price	Absolute Difference
2.0	7.548502	7.549344	0.000842
2.2	7.453536	7.454381	0.000845
2.4	7.359381	7.360137	0.000756
2.6	7.266037	7.266787	0.000749
2.8	7.173501	7.174295	0.000794
3	7.081775	7.082660	0.000885
3.2	6.990857	6.991678	0.000821
3.4	6.900745	6.901351	0.000606
3.6	6.811440	6.812176	0.000736
3.8	6.722939	6.723817	0.000878
4.0	6.635242	6.635881	0.000639

The results confirm that our implementation of the FFT algorithm was accurate.

Next, we test the convergence of the FFT and Monte Carlo simulation for spread options.

4.6.2 Convergence

Table 4.2 below shows the FFT price and execution time for a European spread call option with varying K and N.

Tab. 4.2: Convergence of FFT using the parameters in Table 4.1

N	FFT $K = 2$	FFT $K = 3$	FFT $K = 4$	Time (seconds)
4	4.354906	6.532359	8.709812	0.012491
8	1.488913	2.233370	2.977827	0.013598
16	0.697647	1.046470	1.395293	0.041658
32	0.450374	0.675562	0.900749	0.059231
64	0.936496	1.404743	1.872991	0.206991
128	7.553730	7.087006	6.640188	0.787242
256	7.549344	7.082660	6.635881	3.233390
512	7.549344	7.082660	6.635881	12.481379
1024	7.549344	7.082660	6.635881	50.788263
2048	7.549344	7.082660	6.635881	203.205588
4096	7.549344	7.082660	6.635881	817.879709

The FFT algorithm converges to the solution in approximately 3.23 seconds with $N = 256$ steps.

Figure 4.1 below shows the convergence and execution time for the Monte Carlo simulation:

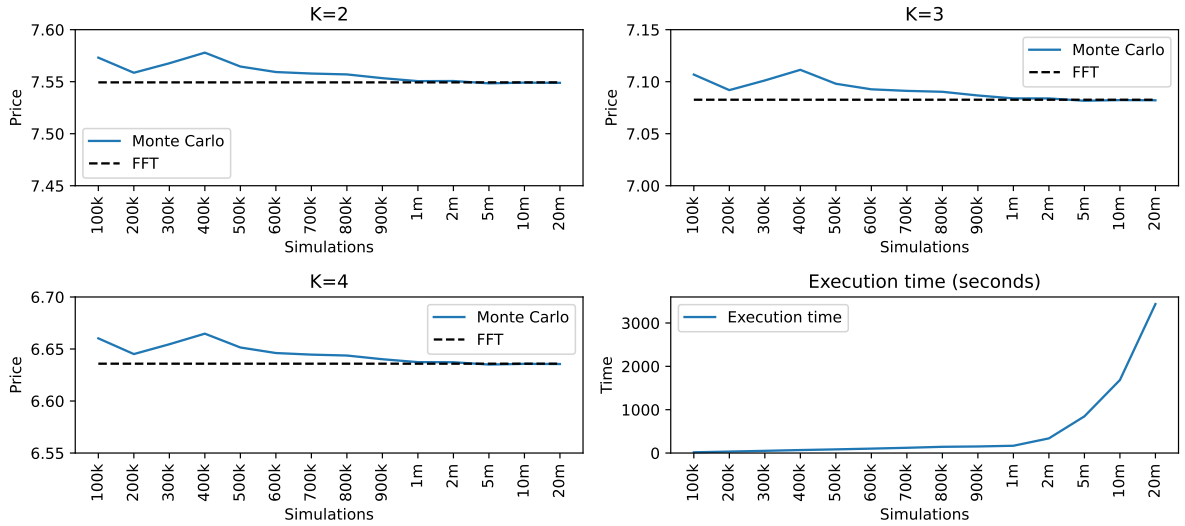

Fig. 4.1: Monte Carlo convergence using the parameters in Table 4.1

The Monte Carlo simulation converged to the FFT price with 1,000,000 simulations in approximately 168 seconds. The FFT significantly outperformed the Monte Carlo in terms of efficiency being up to 50 times faster.

We conclude this chapter by assessing the impact of stochastic interest rates on the spread option prices.

4.6.3 Impact of Stochastic Interest Rates

We consider the following two cases to test the impact of stochastic interest rates: Case 1 where interest rates and equity prices are positively correlated; and Case 2 where interest rates and equity prices are negatively correlated.

Figure 4.2 below shows the results for Case 1 using the same parameters as in Table 4.1 except for η, $\rho_{x_1,r}$, and $\rho_{x_2,r}$:

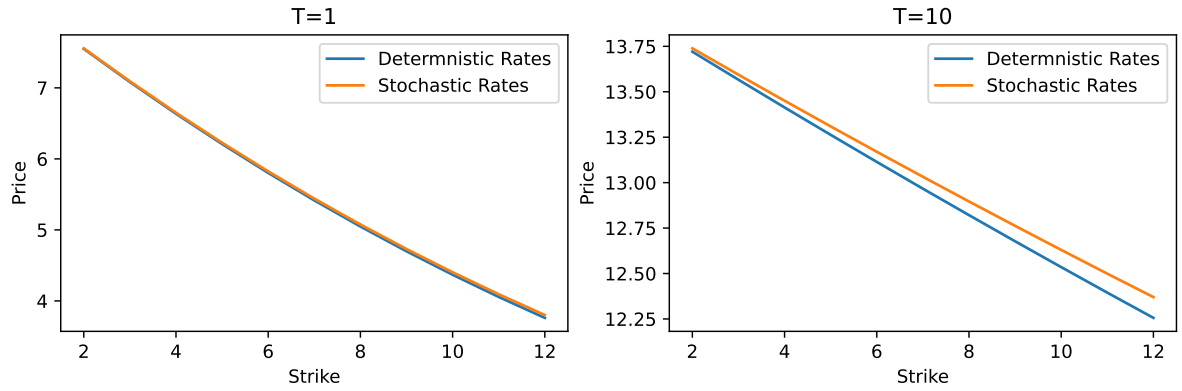

Fig. 4.2: Impact of stochastic interest rates with $\eta = 0.05$, $\rho_{x_1,r} = 0.75$, and $\rho_{x_2,r} = 0.6$.

Similarly, Figure 4.3 below shows the results for Case 2.

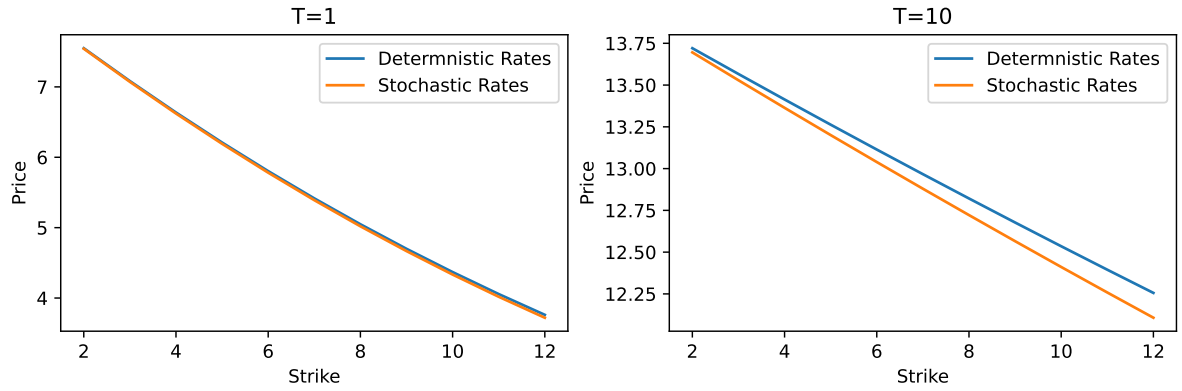

Fig. 4.3: Impact of stochastic interest rates with $\eta = 0.05$, $\rho_{x_1,r} = -0.75$, and $\rho_{x_2,r} = -0.6$.

For short-dated European spread call options, the impact of stochastic interest rates was insignificant. However, for long-dated European spread call options, the price difference between deterministic interest rates and stochastic interest rates widened for further out-the-money options. Moreover, when interest rates and stock prices were positively correlated, the spread option price under stochastic interest rates was higher than under deterministic interest rates. The opposite held when interest rates and stock prices were negatively correlated.

4.7 Conclusions

In this chapter, we extended the Heston-Hull–White model of Grzelak and Oosterlee (2011) to two underlying assets and the FFT algorithm of Hurd and Zhou (2010) to account for stochastic interest rates.

Based on our implementation of the two-asset Heston–Hull–White model, we observed that the FFT algorithm was approximately 50 times faster than the Monte Carlo simulation. We also observed that the price of a long-dated European spread call option was sensitive to stochastic interest rates and equity-rate correlation. This difference became more significant for out-of-the-money options.

We hope that practitioners will find use in our extensions of the Heston–Hull–White model of Grzelak and Oosterlee (2011) and the FFT algorithm of Hurd and Zhou (2010), which will lead them to test the impact of stochastic interest rates on their spread option portfolios.

5 Application of Stochastic Volatility Models in the Real-World Measure

> *Tolerating short-term volatility is how we get to enjoy long-term gains.*
>
> **— David Poppe**

Keywords: Heston model · Bates stochastic volatility jump model · calibration · real-world measure · efficient method of moments · volatility targeting

5.1 Introduction

Equity prices tend to fall sharply in times of economic crisis. Think back to 19 October 1987, *Black Monday*, when the Dow plummeted more than 20% in a single day, for example. Whether it be war, a global pandemic, or a great recession, there is no way of knowing when a market crash will occur. This chapter[1] focuses on jump risk in financial markets.

The seminal Black and Scholes (1973) model used for the pricing of contingent claims, for example, is based on the assumption of log-normal asset returns – the driving force of asset returns in this model assumes a geometric Brownian motion (gBm). Such models are flawed in this respect since they do not account for severe market shocks or the stochastic nature of asset-return volatility.

Strides have been made in the asset-pricing literature to account for the non-normality of equity returns (see, Cont, 2001 for a discussion on the statistical properties of asset returns). The Heston (1993) stochastic volatility model, for example, is frequently used in the pricing of contingent claims based on equities. It provides a reasonably good calibrated fit to the long-term implied volatility skew observed in the equity derivatives market (see, Gatheral, 2006). However, for short maturities, the model fails to produce the steep slope of the skew

[1] This chapter is based on two papers: (Levendis and Maré, 2023a) published in *ORiON* and (Venter et al., 2022) published in *Cogent Economics & Finance*.

(see, for example, El Euch et al., 2019). This indicates that the model is unable to account for short-term shocks (so-called jumps) that may be caused by adverse market events. Statistical tests have been proposed in Aït-Sahalia and Jacod (2009) and Aït-Sahalia et al. (2012) to verify whether jumps are present in financial time series.

A well-established model that can provide a good fit to the short-term skew is the Bates (1996) stochastic volatility jump (SVJ) model. The Bates (1996) SVJ model is an extension of the classical Heston (1993) model that adds random jumps based on a Poisson process. Poklewski-Koziell (2012) performed a detailed analysis of the Heston (1993) and Bates (1996) SVJ models and showed that the Bates (1996) SVJ model produces a good fit to the S&P500 implied volatility surface compared to the Heston (1993) model. Note that Poklewski-Koziell (2012) calibrated the Heston (1993) and Bates (1996) SVJ models to option prices (this calibration is said to be in the risk-neutral pricing measure).

Models with jump dynamics are better at characterising markets with significant implied volatility skews than models without jumps (see, Gatheral, 2006). However, past literature on jump diffusion models tends to focus almost exclusively on calibration to option prices (*i.e.*, the risk-neutral measure). In the absence of a liquidly traded options market, calibration using standard least-squares techniques to minimise the sum of squared differences between market and model prices is infeasible. This is a common problem that plagues illiquid option markets like South Africa. Another challenge is that the estimated density resulting from the calibration to option prices can differ substantially from the estimated density of the historical log returns (see, Grobler and Visagie, 2019).

Calibration of continuous-time models to discrete historical observations has been explored by Andersen et al. (2002); they applied the efficient method of moments (EMM) technique of Gallant and Tauchan (1996) to calibrate a class of SVJ models with Poisson jumps of time-varying intensity to daily S&P500 returns. The EMM is a simulation-based technique that estimates the continuous-time model parameters from the expectation of the derivative of the log-likelihood function, where the log-likelihood function takes the form of a simpler discrete-time auxiliary semi-non-parametric model. Andersen et al. (2002) concluded that stochastic volatility and jumps are both important factors required to characterise daily S&P500 returns.

To our knowledge, there has been no attempt made to calibrate stochastic volatility models, including the Heston (1993) and Bates (1996) SVJ models, to historical equity returns in South Africa (*i.e.*, calibration to the so-called real-world measure). The literature tends to focus solely on calibration to option prices (risk-neutral measure), which, in the context of South African single stocks, for example, is typically impractical (as no generally liquid traded market exists). We, therefore, view this as an excellent opportunity to contribute to the asset-pricing literature with specific focus on calibration of stochastic volatility models in illiquid markets.

The real-world measure is often neglected in favour of the risk-neutral measure due to the pricing of contingent claims. However, the real-world measure is extremely useful and important in risk management and asset/liability applications (see, for example, van Dijk et al., 2018); simulation-based analysis of trading and investment strategies (see, Olivieri et al., 2022); analysis of so-called xVA (counterparty credit, margin, and capital costs); investment based pricing and evaluation of asset price behaviour; and product development, to name a few.

The goal of this chapter is to calibrate the Heston (1993) and Bates (1996) SVJ models to historical FTSE/JSE Top40 returns by making use of the EMM and to test which model is better at characterising the evolution of equity-based risk and return in the South African market. Furthermore, we show how a real-world stochastic volatility model can be used in the context of portfolio management by testing a simple volatility targeting strategy.

The remainder of this chapter is structured as follows: Section 5.2 presents the dynamics for the Heston (1993) and Bates (1996) SVJ models as well as the EMM technique of Gallant and Tauchan (1996). Section 5.3 shows the calibration results of the EMM applied to the S&P500 and FTSE/JSE Top40. Section 5.4 focuses on a practical application of real-world stochastic volatility models by testing a volatility targeting strategy, and Section 5.5 concludes the chapter.

5.2 Stochastic Volatility Models

This section is divided into three parts. The first subsection is dedicated to the Heston (1993) stochastic volatility model. The second subsection focuses on the Bates (1996) SVJ model, and the third subsection discusses the EMM methodology of Gallant and Tauchan (1996).

5.2.1 The Heston Stochastic Volatility Model

Under the real-world probability measure, $\mathbb{P}$, the Heston (1993) model is given by the system of stochastic differential equations (SDEs):

$$\begin{aligned} dS(t) &= \mu S(t)dt + \sqrt{v(t)}S(t)dW_x(t), \\ dv(t) &= (\alpha - \beta v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt, \end{aligned} \tag{5.2.1}$$

where μ denotes the expected rate of return, β is the mean reversion speed of the variance, $\frac{\alpha}{\beta}$ is the long-run mean of the variance, σ_v is the volatility of volatility, and $\rho_{x,v}$ is the correlation between the stock and variance processes.

Note that Eq. (5.2.1) is an alternative representation of the Heston (1993) model. In previous chapters, the mean reversion speed was denoted by κ, hence, $\beta = \kappa$. Furthermore, the long-run mean of the variance was denoted by $\bar{v}$, hence, $\frac{\alpha}{\beta} = \bar{v}$. The representation in Eq. (5.2.1) was chosen to align with the representation in Andersen et al. (2002), which is used as a base for comparison of calibration results.

Taking the log-transform $x(t) = \log S(t)$ and applying Itô's lemma to Eq. (5.2.1), we obtain:

$$\begin{aligned} dx(t) &= (\mu - \frac{1}{2}v(t))dt + \sqrt{v(t)}dW_x(t), \\ dv(t) &= (\alpha - \beta v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt. \end{aligned} \tag{5.2.2}$$

We simulated 100,000 paths over 100 daily time steps from Eq. (5.2.2) using an Euler discretisation scheme to illustrate the impact of certain Heston (1993) model parameters on the returns distribution at the end of 100 days. Figure 5.1 below shows the impact of σ_v on returns:

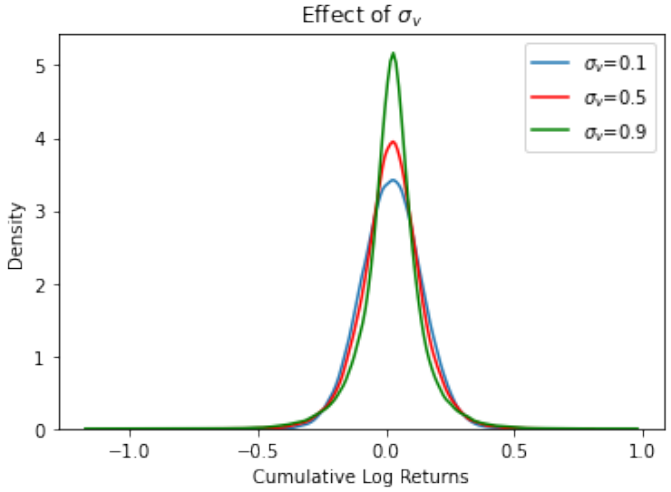

Fig. 5.1: Effect of σ_v on returns

Note that the kurtosis of the returns distribution increases as σ_v increases.

Next, we show the impact of the correlation between the stock and variance processes, denoted by $\rho_{x,v}$. Figure 5.2 below illustrates the impact of $\rho_{x,v}$ on the returns distribution:

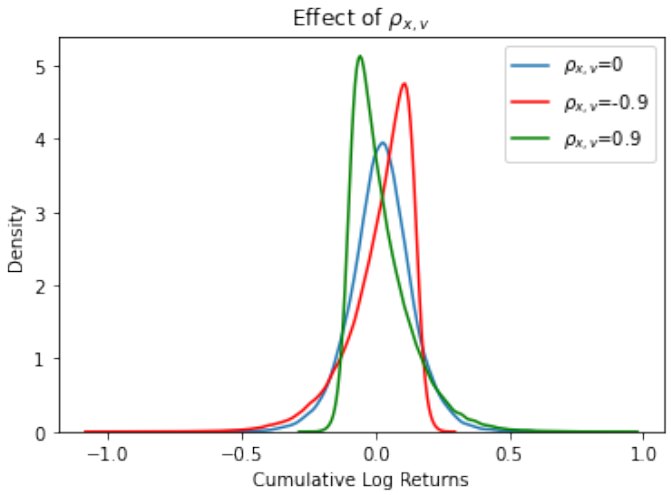

Fig. 5.2: Effect of $\rho_{x,v}$ on returns

The parameter $\rho_{x,v}$ controls the skewness of the returns distribution - a negative value for $\rho_{x,v}$ results in a negatively skewed distribution and vice versa.

Even though the Heston (1993) model displays some desirable properties, it still underestimates the kurtosis of returns observed in practice (see, for example, González-Urteaga, 2012). Therefore, we consider the Bates (1996) SVJ model next.

5.2.2 The Bates Stochastic Volatility Jump Model

Under the $\mathbb{P}$-measure, the Bates (1996) SVJ model is presented by the system of SDEs:

$$\begin{aligned} dS(t) &= (\mu - \lambda\mu_J)S(t)dt + \sqrt{v(t)}S(t)dW_x(t) + JS(t)dN(t), \\ dv(t) &= (\alpha - \beta v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt, \end{aligned} \tag{5.2.3}$$

where $N(t)$ is a Poisson process with jump intensity λ. Furthermore, J denotes the percentage jump size of the underlying where

$$\log(1+J) \sim \mathcal{N}\Big(\log(1+\mu_S) - 0.5\sigma_S^2, \sigma_S^2\Big),$$

$$\mu_J = \exp\left\{\mu_S + \frac{\sigma_S^2}{2}\right\} - 1,$$

with μ_S and σ_S the mean and volatility of the jump size.

Taking the log-transform $x(t) = \log S(t)$ and applying Itô's lemma to Eq. (5.2.3), we get:

$$\begin{aligned} dx(t) &= (\mu - \lambda\mu_J - \frac{1}{2}v(t))dt + \sqrt{v(t)}dW_x(t) + \log(1+J)dN(t), \\ dv(t) &= (\alpha - \beta v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt. \end{aligned} \tag{5.2.4}$$

Similar to what was done using the Heston (1993) model, we simulated 100,000 paths over 100 daily time steps from Eq. (5.2.4). Figure 5.3 below shows the impact of λ on returns:

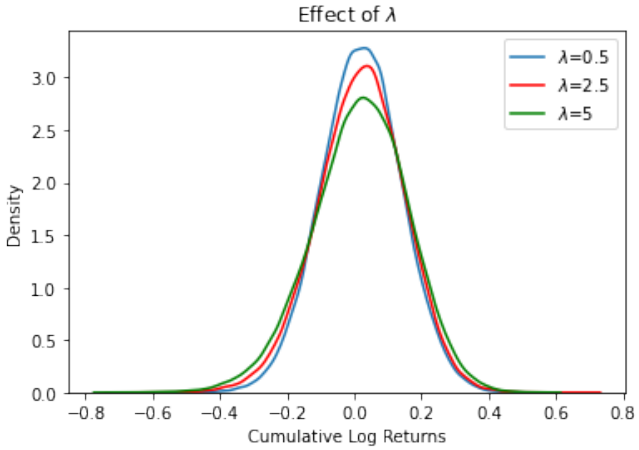

Fig. 5.3: Effect of λ on returns

The kurtosis increases as λ increases. Figure 5.4 below illustrates the impact of μ_S on returns.

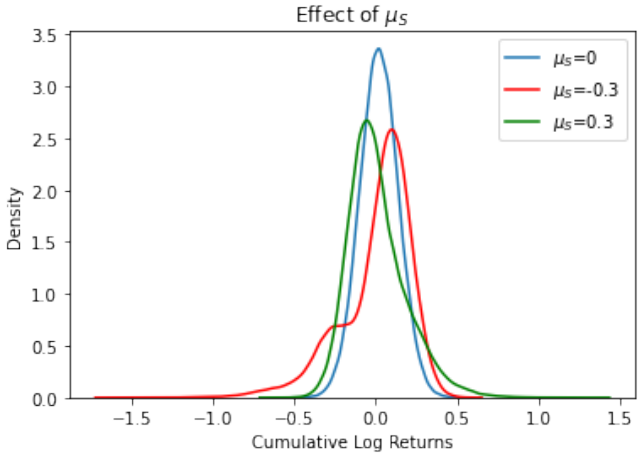

Fig. 5.4: Effect of μ_S on returns

Similar to the parameter $\rho_{x,v}$, the mean jump size μ_S controls the skewness of the distribution – a negative value for μ_S produces a negatively skewed distribution and vice versa. Lastly, Figure 5.5 below shows the impact of σ_S on returns:

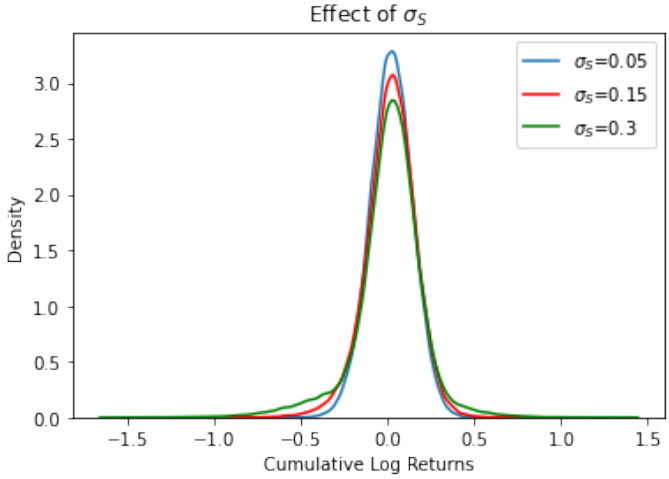

Fig. 5.5: Effect of σ_S on returns

Note that the kurtosis of the distribution increases as σ_S increases.

It should be clear that the Bates (1996) SVJ model provides more flexibility than the Heston (1993) model to capture additional features of returns.

In the next subsection, we present the EMM technique of Gallant and Tauchan (1996) that will be used to calibrate the Heston (1993) and Bates (1996) SVJ models to historical equity returns.

5.2.3 Efficient Method of Moments

The EMM procedure is outlined in Andersen et al. (1999). Suppose there is a historical time series $Y_t = \{y_1, ...y_t\}$ and the aim is to estimate the vector of stochastic volatility model parameters $\boldsymbol{\theta}$ from this series. There is no analytical expression for the likelihood function of the Heston (1993) or Bates (1996) SVJ models. Therefore, the first step in the EMM procedure is to choose an auxiliary model (score generator) with transition density function $f(y_t|Y_{t-1}, \boldsymbol{\eta})$, parameterised by the pseudo parameter vector $\boldsymbol{\eta}$. To this end, we choose the semi-nonparametric density of Gallant and Nychka (1987), where a leading parametric term is chosen to capture the majority of the dependency in the conditional mean and variance. The transition density function is then extended by adding Hermite polynomials that capture any remaining non-Gaussian features in the time series. The semi-nonparametric density is given by:

$$f_K(y_t|Y_{t-1}; \boldsymbol{\eta}) = \left(v + (1-v) \frac{[P_K(z_t)]^2}{\int_{\mathbb{R}} [P_K(u)]^2 \varphi(u) du} \right) \frac{\varphi(z_t)}{\sqrt{h_t}},$$

where $v = 0.01$ to avoid instability during the estimation procedure when $P_K(z_t) = 0$, $\varphi(\cdot)$ is the standard normal density function, and

$$z_t = \frac{y_t - \mu_t}{\sqrt{h_t}},$$

with

$$\mu_t = 0,$$

$$h_t = \omega + \gamma_0 y_{t-1}^2 + \gamma_1 h_{t-1} \sim GARCH(1,1),$$

where μ_t and h_t denote the conditional mean and variance, respectively. Note that μ_t need not be zero, and h_t may be specified by alternative discrete-time models such as ARCH or EGARCH (see, Andersen et al., 1999, for details).

The Hermite polynomials are given by:

$$P_K(z_t) = \sum_{i=0}^{K_z} a_i z_t^i,$$

where $a_0 = 1$ and K_z denotes the order of the Hermite polynomial. Note that the EMM procedure requires the dimension of the pseudo parameter vector $\boldsymbol{\eta}$ to be greater than or equal to the dimension of the stochastic volatility parameter vector $\boldsymbol{\theta}$.

Once an auxiliary model has been chosen, the pseudo parameter vector $\boldsymbol{\eta}$ is estimated using maximum likelihood. The maximum likelihood estimator $\hat{\boldsymbol{\eta}}_T$ satisfies the first-order conditions

$$\frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\boldsymbol{\eta}}\log f(y_t|Y_{t-1},\hat{\boldsymbol{\eta}}_T) = \frac{1}{T}\sum_{t=1}^{T}s_f(Y_t,\hat{\boldsymbol{\eta}}_T) = 0,$$

where $s_f(Y_t,\hat{\boldsymbol{\eta}}_T) = \frac{\partial}{\partial\boldsymbol{\eta}}\log f(y_t|Y_{t-1},\hat{\boldsymbol{\eta}}_T)$ denotes the score function of the auxiliary model.

The second step in the EMM procedure is to simulate a series $\hat{y}_n(\boldsymbol{\theta})$, $n = 1,...,N$, from the stochastic volatility model for a given $\boldsymbol{\theta}$ and evaluate the sample moments at the fixed maximum likelihood estimate $\hat{\boldsymbol{\eta}}_T$, *i.e.*,

$$m_N(\boldsymbol{\theta},\hat{\boldsymbol{\eta}}_T) = \frac{1}{N}\sum_{n=1}^{N}\frac{\partial}{\partial\boldsymbol{\eta}}\log f(\hat{y}_n(\boldsymbol{\theta})|\hat{Y}_{n-1}(\boldsymbol{\theta}),\hat{\boldsymbol{\eta}}_T).$$

Andersen et al. (1999) mention that $m_N(\boldsymbol{\theta},\hat{\boldsymbol{\eta}}_T) \to m(\boldsymbol{\theta},\hat{\boldsymbol{\eta}}_T)$ as $N \to \infty$; hence, the simulated time series should be large enough so that the Monte Carlo error can be ignored.

Finally, the vector of stochastic volatility parameters $\boldsymbol{\theta}$ is estimated by minimising the objective function

$$\hat{\theta} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}}[m_N(\boldsymbol{\theta},\hat{\boldsymbol{\eta}}_T)'\hat{I}_T^{-1}m_N(\boldsymbol{\theta},\hat{\boldsymbol{\eta}}_T)], \tag{5.2.5}$$

where $\hat{I}_T$ is a consistent estimator of the asymptotic covariance matrix I of the sample pseudo score vector. The estimator $\hat{I}_T$ is calculated as the outer product of scores, *i.e.*,

$$\hat{I}_T = \frac{1}{T}\sum_{t=1}^{T}\frac{\partial}{\partial\boldsymbol{\eta}}\log f(y_t|Y_{t-1},\hat{\boldsymbol{\eta}}_T)\frac{\partial}{\partial\boldsymbol{\eta}}\log f(y_t|Y_{t-1},\hat{\boldsymbol{\eta}}_T)'.$$

A major advantage of the EMM is that T multiplied by the minimised value in Eq. (5.2.5) follows a χ^2 distribution with $n_{\boldsymbol{\eta}} - n_{\boldsymbol{\theta}}$ degrees of freedom, where $n_{\boldsymbol{\eta}}$ and $n_{\boldsymbol{\theta}}$ denote the number of parameters in the semi-nonparametric and stochastic volatility models. Therefore, a goodness-of-fit test can be performed by comparing the final estimate of the objective function multiplied by the number of observations used in the calibration with the relevant percentile from the χ^2 distribution. Andersen et al. (1999) explain that the null hypothesis in this case is that the model has been correctly specified. Hence, if the minimised value in Eq. (5.2.5) multiplied by the number of observations is less than the critical value from the χ^2 distribution, the hypothesis is not rejected.

In the next section, we calibrate the Heston (1993) and Bates (1996) SVJ models to historical S&P500 and FTSE/JSE Top40 returns using the EMM.

5.3 Empirical Results

This section is divided into two parts. The first subsection focuses on the S&P500 and compares our EMM implementation of the Heston (1993) and Bates (1996) SVJ models to the results in Andersen et al. (2002). The reason for this is to validate the accuracy of our implementation. We further extend the analysis of Andersen et al. (2002) by calibrating the Heston (1993) model over different periods to test the stability of the model parameters. The second subsection uses the EMM to calibrate the Heston (1993) and Bates (1996) SVJ models to the FTSE/JSE Top40 to test which model is better at capturing risk and return in the South African equity market.

5.3.1 S&P500

The S&P500 is considered by many to be the best indicator of global equity market performance. It consists of 500 leading companies that are publicly traded in the United States. Andersen et al. (2002) calibrated the Heston (1993) and Bates (1996) SVJ models using the EMM over the period 2 January 1953 to 31 December 1996 to S&P500 returns. Our goal was to replicate their results over the same time period to confirm the accuracy of our EMM implementation.

Figure 5.6 shows the S&P500 closing prices from 2 January 1953 to 31 December 1996:

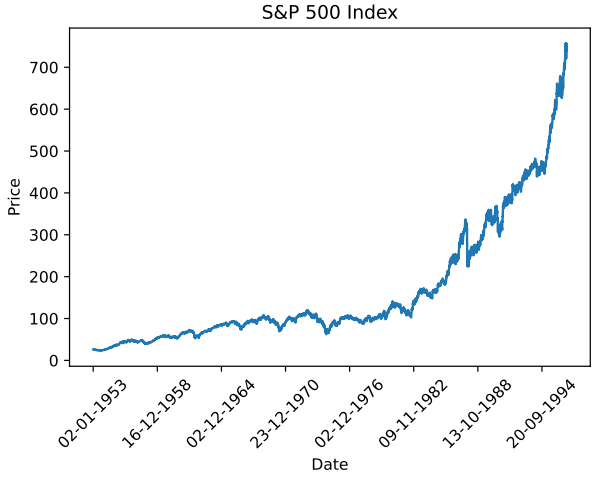

Fig. 5.6: S&P500 daily historical closing prices

Figure 5.7 below shows the daily log returns for the S&P500 over the same period:

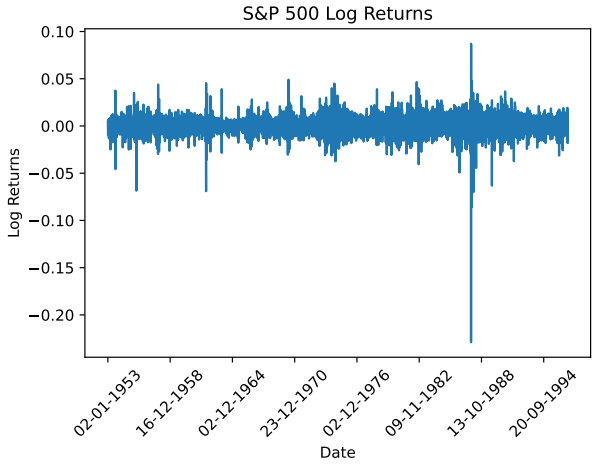

Fig. 5.7: S&P500 log returns

The most significant event over the period 2 January 1953 to 31 December 1996 was *Black Monday*, 19 October 1987, when the S&P500 fell by more than 20%. Prior to *Black Monday*, equity market behaviour was relatively stable. The concept of a *volatility skew* was unknown and the assumption of log-normal returns seemed reasonable.

Figure 5.8 below shows the daily S&P500 log returns with the normal distribution superimposed over the returns:

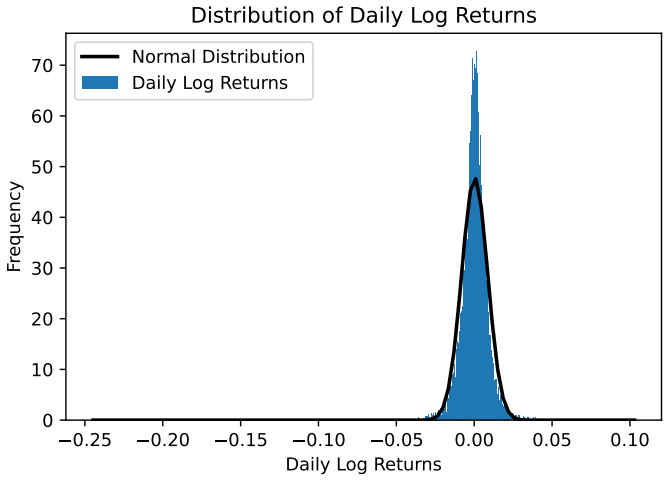

Fig. 5.8: S&P500 and normal density

It is clear from Figure 5.8 that the normal distribution cannot capture the high peak and fat tails of the S&P500 returns. Assuming a log-normal distribution of returns is, therefore, not consistent with historical equity behaviour and significantly underestimates the size and frequency of equity price drops.

Our first goal was to reproduce the calibration results in Andersen et al. (2002) for the Heston (1993) and Bates (1996) SVJ models over the same period from 2 January 1953 to 31 December 1996 to validate the accuracy of our EMM implementation. The results for the Heston (1993) stochastic volatility model are shown below.

Heston Stochastic Volatility Model

Table 5.1 below shows a comparison of the Heston (1993) model parameters based on our EMM implementation to the parameters in Andersen et al. (2002):

Tab. 5.1: Comparison of annualised Heston parameters for S&P500

Parameter	**Andersen *et al.***	**Our implementation**
μ	0.0756	0.0768
α	0.0438	0.0485
β	3.2508	4.2142
σ_v	0.1850	0.1768
$\rho_{x,v}$	-0.5877	-0.4323
$T\hat{\theta}$	31.9400	15.6099

The calibrated parameters achieved using our method and those of Andersen et al. (2002) align well. We do not expect to match their parameters exactly since we are likely using a different optimisation routine and have access to better software compared to what was available 20 years ago. The goodness-of-fit statistic, $T\hat{\theta}$, is also shown in Table 5.1 for completeness.

Interestingly, the goodness-of-fit statistic is significantly smaller in our implementation than in Andersen et al. (2002). This indicates that our implementation of the Heston (1993) model yields a better fit to the S&P500 returns than the Heston (1993) model in Andersen et al. (2002). Figures 5.9 and 5.10 below show a visual comparison of the densities generated by the model parameters in Table 5.1.

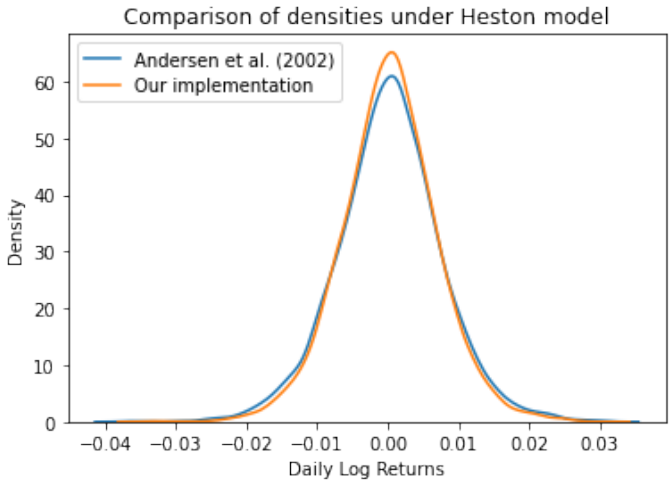

Fig. 5.9: Comparison of densities under the Heston stochastic volatility model for S&P500

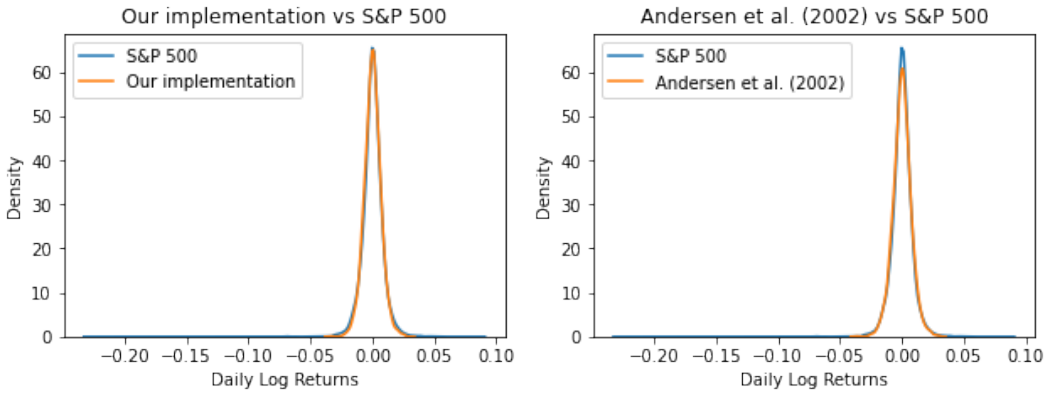

Fig. 5.10: Comparison of Heston and S&P500 densities

Figure 5.9 indicates that the two sets of model parameters yield very similar results. Figure 5.10 shows that our implementation of the Heston (1993) model captures the peak of the S&P500 returns slightly better than Andersen et al. (2002).

Table 5.2 below shows the first four statistical moments as well as the minimum and maximum values generated by the Heston (1993) model compared to the S&P500.

Tab. 5.2: Heston model daily statistics for the S&P500

Statistic	S&P500 index	Andersen *et al.*	Our implementation
Mean	0.0301%	0.0231%	0.0235%
Std dev	0.8346%	0.7498%	0.6880%
Skewness	-2.0220	-0.0179	-0.0066
Kurtosis	60.0830	4.0254	3.9219
Minimum	-0.2290	-0.0379	-0.0349
Maximum	0.0871	0.0319	0.0310

The first two moments of the S&P500 (mean and standard deviation) are captured relatively well by the Heston (1993) model. However, the Heston (1993) model substantially underestimates the skewness and kurtosis over the period 2 January 1953 to 31 December 1996. As explained by Andersen et al. (2002), stochastic volatility on its own is not adequate to describe the observed S&P500 returns.

Next, we extended the results of Andersen et al. (2002) by calibrating the Heston (1993) model over different periods, spanning approximately 20 years, to test the stability of the model parameters. Table 5.3 below shows the first four statistical moments and the minimum and maximum values over each period:

Tab. 5.3: S&P500 daily statistics over different periods

Statistic	1960-1980	1970-1990	1980-2000	1990-2010	2000-2022
Mean	0.0117%	0.0264%	0.0520%	0.0224%	0.0174%
Std dev%	0.7508%	0.9861%	0.9996%	1.1724%	1.2474%
Skewness	0.04194	-2.5083	-2.6223	-0.1980	-0.4009
Kurtosis	7.5054	64.9265	62.9018	12.1684	13.5220
Minimum	-0.0691	-0.2290	-0.2290	-0.0947	-0.1277
Maximum	0.0490	0.0871	0.0871	0.1096	0.1096

The periods 1970 to 1990 and 1980 to 2000 both contain *Black Monday* and show substantially different values for the skewness and kurtosis compared to the other periods. Not even the global financial crisis of 2008 or the COVID-19 pandemic came close to the crash of 19 October 1987.

Table 5.4 below shows the calibrated parameters and goodness-of-fit statistic for the Heston (1993) model using the EMM for each period.

Tab. 5.4: Calibrated Heston model parameters (annualised) for S&P500

Parameter	1960-1980	1970-1990	1980-2000	1990-2010	2000-2022
μ	0.0277	0.0670	0.1359	0.0891	0.0858
α	0.0534	0.0769	0.0774	0.0602	0.1137
β	4.8925	5.0140	5.0555	3.0968	4.9395
σ_v	0.2341	0.2015	0.2266	0.2292	0.3747
$\rho_{x,v}$	-0.5703	0.0347	-0.4144	-0.8891	-0.8801
$T\hat{\theta}$	2.9665	6.7840	22.8005	18.1143	7.2800
$\chi^2_{0.05}$	9.4877	9.4877	9.4877	9.4877	9.4877

The expected return, μ, can vary substantially between periods and follows a similar trend to the mean of the S&P500 returns in Table 5.3. The Heston (1993) stochastic volatility parameters (α, β, σ_v) are relatively stable over time. The correlation parameter, $\rho_{x,v}$, shows a strong negative relationship between returns and asset volatility except for the period 1970 to 1990. Interestingly, $\rho_{x,v}$ changes from 0.0347 in the period 1970 to 1990 to -0.4144 in 1980 to 2000. Both these periods include *Black Monday*. This instability suggests that the estimation of $\rho_{x,v}$ is sensitive to outliers.

The goodness-of-fit statistic shows interesting results. From Table 5.4, the goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Heston (1993) model is not rejected for the periods 1960 to 1980, 1970 to 1990, and 2000 to 2022. However, this hypothesis is rejected for the periods 1980 to 2000 and 1990 to 2010, which indicates that the calibration is not stable over time.

Next, we calibrate the Bates (1996) SVJ model to historical S&P500 returns.

Bates Stochastic Volatility Jump Model

The Bates (1996) SVJ model adds three additional jump parameters, λ, μ_S, and σ_S to the standard Heston (1993) model. Andersen et al. (2002) explain that the mean jump parameter, μ_S, is of less importance and poorly identified in general. Therefore, we follow Andersen et al. (2002) by imposing the restriction $\mu_S = 0$. The Bates (1996) SVJ model in Eq. (5.2.4) then becomes:

$$\begin{aligned} dx(t) &= (\mu - \lambda\mu_J - \frac{1}{2}v(t))dt + \sqrt{v(t)}dW_x(t) + \log(1+J)dN(t), \\ dv(t) &= (\alpha - \beta v(t))dt + \sigma_v\sqrt{v(t)}dW_v(t), \\ dW_x(t)dW_v(t) &= \rho_{x,v}dt, \end{aligned} \tag{5.3.1}$$

where

$$\log(1+J) \sim \mathcal{N}\Big(-0.5\sigma_S^2, \sigma_S^2\Big).$$

Table 5.5 below compares the Bates (1996) SVJ parameters based on our implementation of the EMM to that of the results in Andersen et al. (2002) over the period 2 January 1953 to 31 December 1996:

Tab. 5.5: Comparison of annualised Bates SVJ parameters for S&P500

Parameter	Andersen *et al.*	Our implementation
μ	0.0766	0.0815
α	0.0438	0.0481
β	3.0240	3.6782
σ_v	0.1792	0.2133
$\rho_{x,v}$	-0.6220	-0.4926
λ	5.0904	4.0147
σ_S	0.0134	0.0151
$T\hat{\theta}$	14.9000	2.3648

Based on our implementation, the goodness-of-fit statistic decreases from 15.6099 under the Heston (1993) model to 2.3648 under the Bates (1996) SVJ model. This indicates that both stochastic volatility and jumps are important factors to consider when modelling S&P500 returns.

An interesting observation is that the goodness-of-fit statistic based on our implementation is significantly smaller than Andersen et al. (2002), which suggests a better fit to the S&P500.

Figures 5.11 and 5.12 compare the densities generated by the Bates (1996) SVJ model:

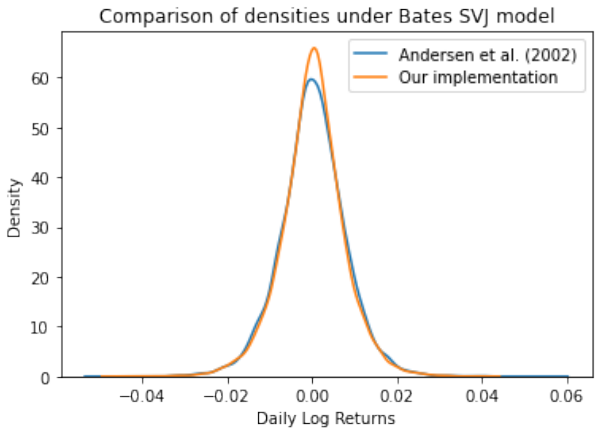

Fig. 5.11: Comparison of densities under the Bates SVJ model for S&P500

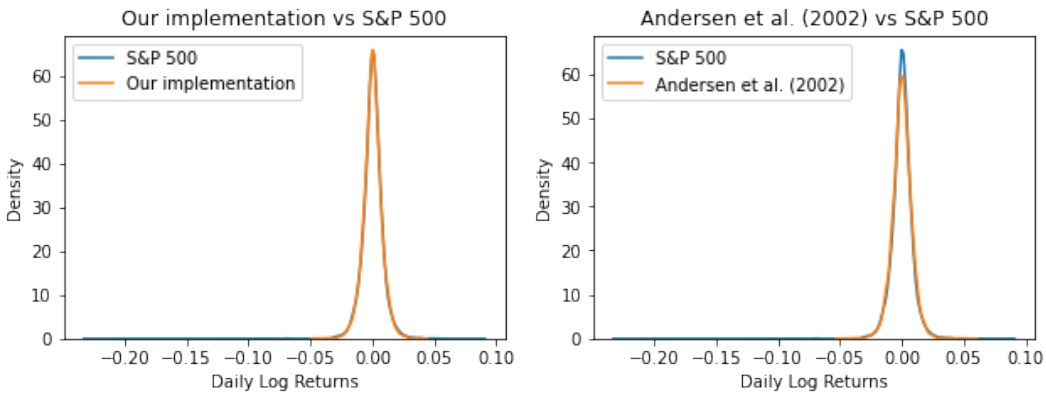

Fig. 5.12: Comparison of Bates and S&P500 densities

Figure 5.11 shows that the distribution generated by the Bates (1996) SVJ model yields similar results to Andersen et al. (2002). Figure 5.12 shows that our implementation captures the peak of the S&P500 returns slightly better than Andersen et al. (2002).

Table 5.6 below compares the first four statistical moments as well as the minimum and maximum values generated by the Bates (1996) SVJ model with the S&P500:

Tab. 5.6: Bates model daily statistics for the S&P500

Statistic	**S&P500 index**	**Andersen *et al.***	**Our implementation**
Mean	0.0301%	0.0239%	0.0209%
Std dev	0.8346%	0.7717%	0.7433%
Skewness	-2.0220	-0.0133	-0.0923
Kurtosis	60.0830	4.6294	4.5946
Minimum	-0.2290	-0.0498	-0.0460
Maximum	0.0871	0.0566	0.0407

Adding jumps to the return process improves the results for the skewness and kurtosis, but not nearly enough to capture the severe market shock of 19 October 1987.

In the next section, we calibrate the Heston (1993) and Bates (1996) SVJ models to FTSE/JSE Top40 returns to test which model captures risk and return best in the South African market.

5.3.2 FTSE/JSE Top40

The FTSE/JSE Top40 is an index consisting of the 40 largest publicly traded companies by market capitalisation in South Africa. Figure 5.13 below shows daily historical closing prices for the FTSE/JSE Top40 from 30 June 1995 to 30 June 2022:

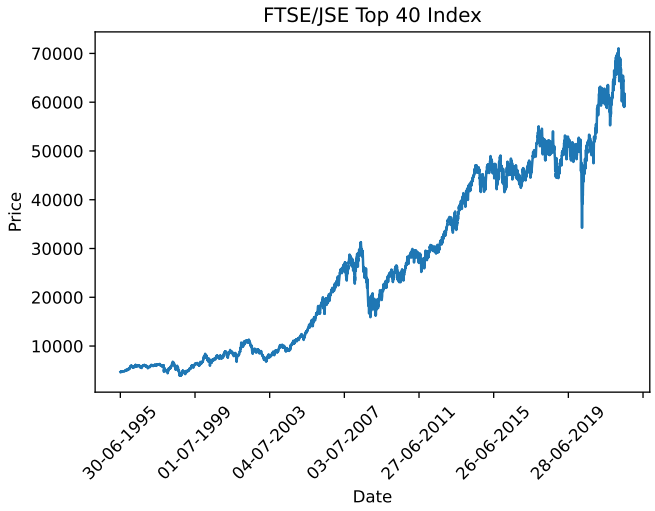

Fig. 5.13: FTSE/JSE Top40 historical closing prices

Figure 5.14 below shows the daily log returns for the FTSE/JSE Top40 from 30 June 1995 to 30 June 2022:

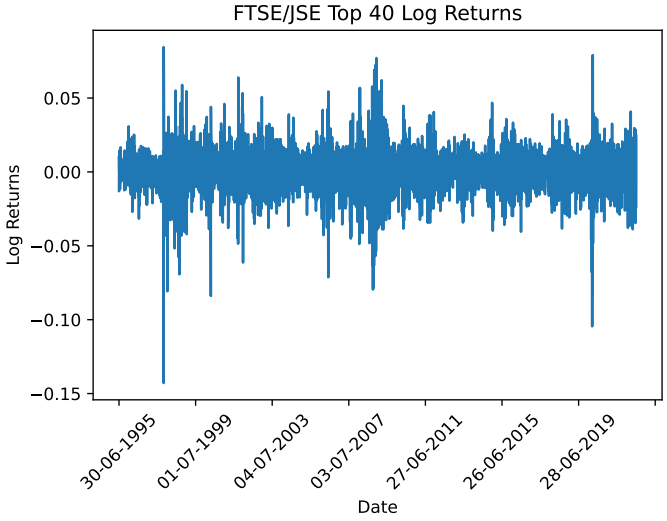

Fig. 5.14: FTSE/JSE Top40 log returns

Figure 5.15 below shows the normal distribution superimposed over the daily FTSE/JSE Top40 returns:

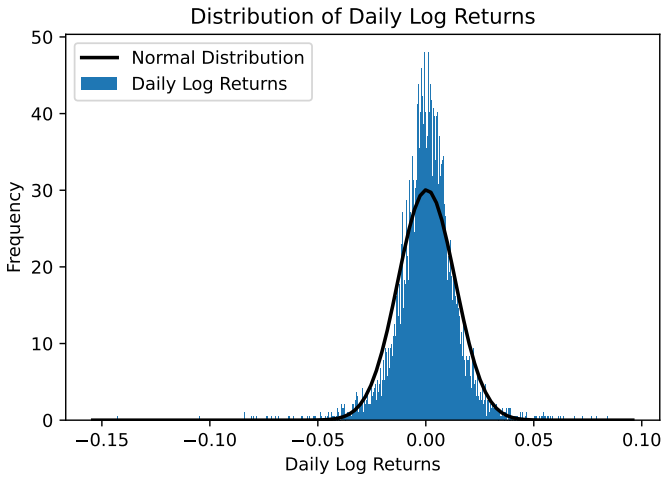

Fig. 5.15: FTSE/JSE Top40 and normal density

Note that the normal distribution is not able to capture the high peak and fat tails observed in the empirical distribution of daily FTSE/JSE Top40 returns. Next, we calibrate the Heston (1993) model to daily FTSE/JSE Top40 returns.

Heston Stochastic Volatility Model

Table 5.7 below shows the calibrated Heston (1993) parameters to daily FTSE/JSE Top40 returns over the period 30 June 1995 to 30 June 2022 using the EMM:

Tab. 5.7: Annualised Heston parameters for FTSE/JSE Top40

Parameter	Estimate
μ	0.0982
α	0.2342
β	6.9424
σ_v	0.4782
$\rho_{x,v}$	-0.9364
$T\hat{\theta}$	13.4381
$\chi^2_{0.05}$	9.4877

Note that the correlation, $\rho_{x,v}$, for the FTSE/JSE Top40 is much more pronounced (negative) than the correlation for the S&P500. The goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Heston (1993) model is rejected at a 5% level of significance.

Figure 5.16 compares the Heston (1993) and FTSE/JSE Top40 densities:

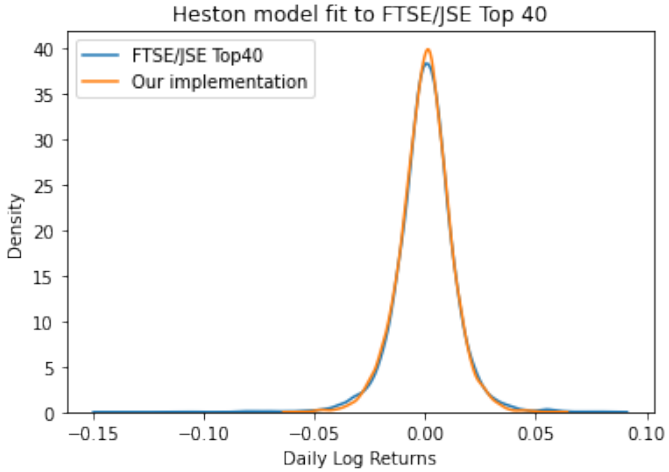

Fig. 5.16: Comparison of densities under the Heston model for FTSE/JSE Top40

The Heston (1993) model fits the FTSE/JSE Top40 density well. Table 5.8 below compares the first four moments as well as the minimum and maximum values generated by the Heston (1993) model with the daily FTSE/JSE Top40:

Tab. 5.8: Heston model daily statistics for the FTSE/JSE Top40

Statistic	FTSE/JSE Top40 index	Our implementation
Mean	0.0385%	0.0270%
Std dev	1.3290%	1.1679%
Skewness	-0.4369	-0.1313
Kurtosis	9.4344	4.1614
Minimum	-0.1429	-0.0585
Maximum	0.0845	0.0588

The Heston (1993) model captures the mean and standard deviation of the FTSE/JSE Top40 well. However, the model underestimates the observed skewness and kurtosis. Next, we calibrate the Bates (1996) SVJ model to historical FTSE/JSE Top40 returns.

Bates Stochastic Volatility Jump Model

Table 5.9 below shows the calibrated Bates (1996) SVJ parameters to daily FTSE/JSE Top40 returns over the period 30 June 1995 to 30 June 2022:

Tab. 5.9: Annualised Bates parameters for FTSE/JSE Top40

Parameter	Estimate
μ	0.0893
α	0.1815
β	5.4469
σ_v	0.4328
$\rho_{x,v}$	-0.8705
λ	3.9973
σ_J	0.0147
$T\hat{\theta}$	4.7879
$\chi^2_{0.05}$	5.9915

Note that the jump parameter, λ, indicates that jumps occur approximately four times per year. The goodness-of-fit statistic suggests that the hypothesis that the observed data are realised from the calibrated Bates (1996) SVJ model is not rejected at a 5% level of significance. The Bates (1996) SVJ model is, therefore, a plausible data-generating model for the FTSE/JSE Top40.

Figure 5.17 compares the density generated by the Bates (1996) SVJ model to the FTSE/JSE Top40 density:

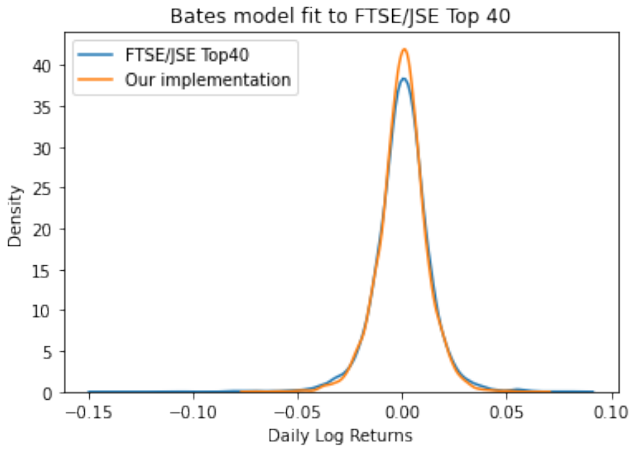

Fig. 5.17: Comparison of densities under the Bates SVJ model for FTSE/JSE Top40

The two densities align well. Table 5.10 below compares the first four statistical moments as well as the minimum and maximum values generated by the Bates (1996) SVJ model with the daily FTSE/JSE Top40 returns:

Tab. 5.10: Bates model daily statistics for the FTSE/JSE Top40

Statistic	FTSE/JSE Top40 index	Our implementation
Mean	0.0385%	0.0246%
Std dev	1.3290%	1.1532%
Skewness	-0.4369	-0.2039
Kurtosis	9.4344	4.9364
Minimum	-0.1429	-0.0714
Maximum	0.0845	0.0652

The Bates (1996) SVJ still underestimates the skewness and kurtosis observed in the FTSE/JSE Top40, but improves the fit compared to that of the Heston (1993) model.

Given that the Bates (1996) SVJ model captures the higher order moments better than the Heston (1993) model, we conclude that both stochastic volatility and jumps are required to characterise equity returns in the South African market.

In the next section, we show how the real-world Bates (1996) SVJ model can be used in practice by considering a simple volatility targeting strategy.

5.4 Volatility Targeting

In this section, we extend the work of Khuzwayo and Maré (2014) by implementing a simulation-based approach to assess the risk and return of various volatility targeting strategies in the South African equity market.

As explained by Khuzwayo and Maré (2014), volatility targeting is an asset allocation strategy that aims to keep the volatility of a portfolio stable by updating the allocation between a risky asset and cash on a regular basis.

Let $\Pi(t)$ denote the time t value of a portfolio consisting of a single risky asset (an equity index) and cash. Mathematically, the change in the value of the portfolio over the interval dt can be written as:

$$d\Pi(t) = w_S(t)\Pi(t)\frac{dS(t)}{S(t)} + rw_C(t)\Pi(t)dt + qw_S(t)\Pi(t)dt, \quad \Pi(0) > 0,$$

where r is the continuously compounded rate earned on cash, q is the continuously compounded dividend yield per annum; w_S and w_C are the equity and cash weights given by:

$$w_S(t) = \frac{\sigma_{Target}}{\sigma_{Model}(t+1)}, \quad w_C(t) = 1 - w_S(t).$$

Following Khuzwayo and Maré (2014), we impose the restriction $w_S \leq 1$ so that gearing (borrowing funds to increase equity exposure) is not allowed.

At each time t, we generate a volatility forecast for time $t+1$. To do this, we fit a GARCH(1,1) model (see, for example, Brownlees et al., 2012 and Venter et al., 2022) to each return series simulated from the Bates (1996) SVJ model.

To simplify matters and focus on equity, we assume $r = 0\%$ to ignore the effect of interest compounding. For the dividend yield, we set $q = 2.5\%$, roughly representing the long-term dividend yield of the equity market.

Figures 5.18 to 5.22 below illustrate a 10% volatility targeting strategy by simulating two paths for the equity index from the Bates (1996) SVJ model using an Euler Monte Carlo scheme over a period of one year (252 trading days):

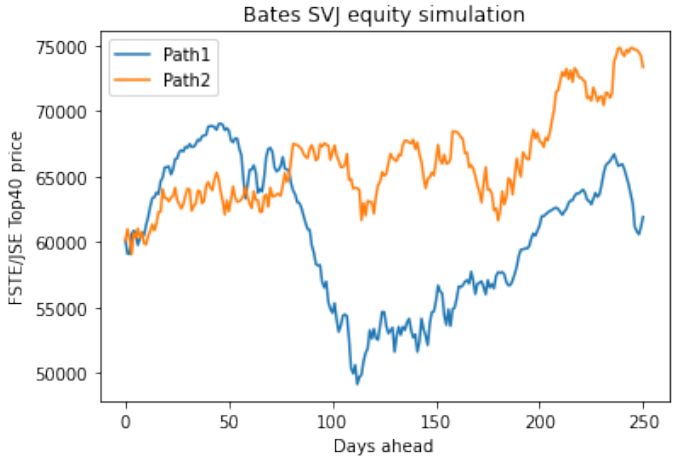

Fig. 5.18: Bates SVJ equity forecast

For each return series simulated from the Bates (1996) SVJ model, we fit a GARCH(1,1) model to forecast volatility. The GARCH(1,1) model is calibrated once every month to a rolling

1000-day history of returns and used to predict 1-week volatility as illustrated in Figure 5.19:

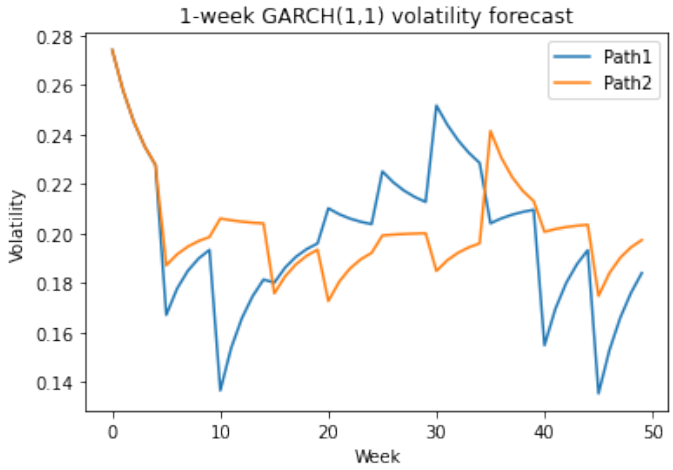

Fig. 5.19: Weekly volatility forecast

At the start of each week, we use the volatility forecast to calculate the equity/cash weights and track the performance of an initial investment of R100. This is illustrated in Figures 5.20 to 5.22:

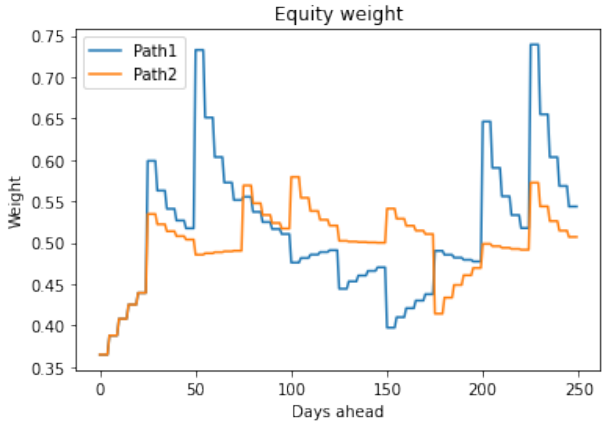

Fig. 5.20: Equity weight based on 10% volatility target

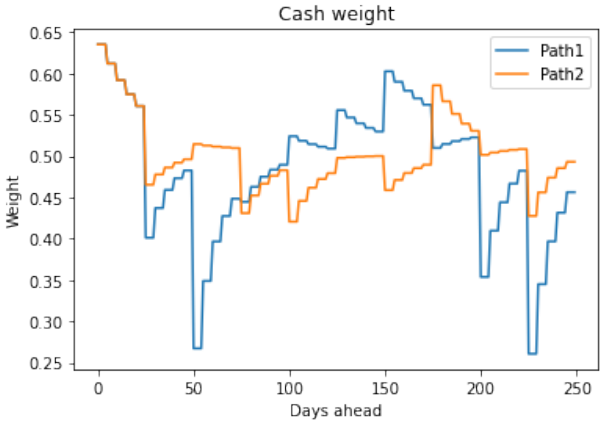

Fig. 5.21: Cash weight based on 10% volatility target

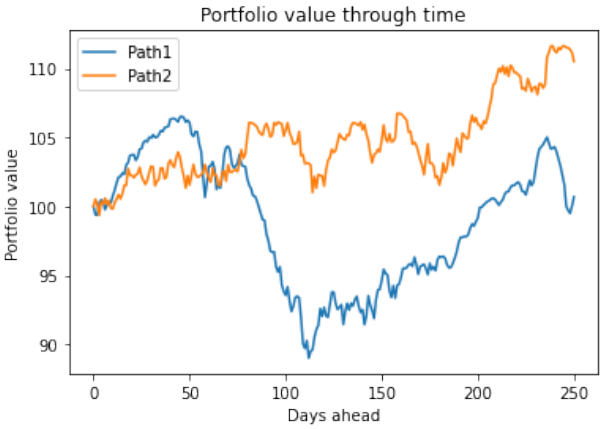

Fig. 5.22: Portfolio value through time

Note that the equity and cash weights are updated at the start of each week and held constant for a one-week period.

Next, we scaled the number of simulations to 1000 and tested the performance of different volatility targeting strategies over different investment horizons. We considered volatility targets of 10%, 15%, and 20% and investment horizons of one, three, and five years. We then compared the performance of the portfolio to that of an equity-only holding strategy. Note that the simulation-based approach allows us to analyse a distribution of returns. The distribution and statistics for each investment horizon and volatility targeting strategy are shown below.

One-year Performance

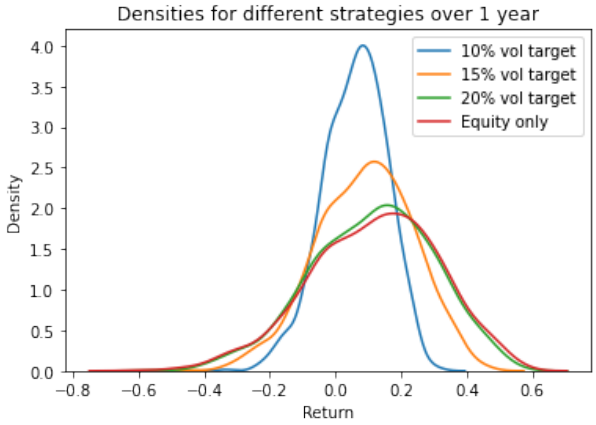

Fig. 5.23: 1-year returns

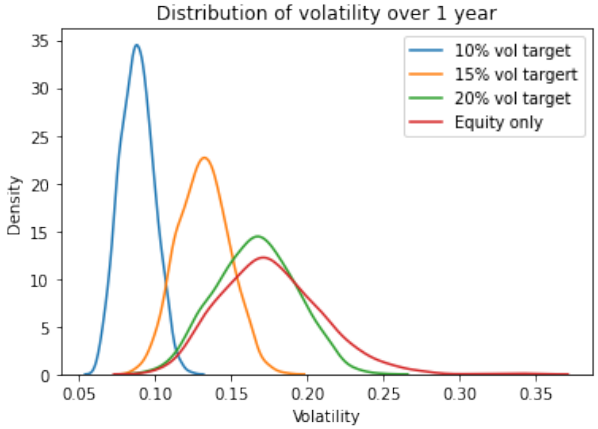

Fig. 5.24: 1-year volatility

Tab. 5.11: Statistics for volatility targeting strategies over one year

Statistic	10% target	15% target	20% target	Equity only
Mean return	5.9035%	8.9714%	11.2965%	11.8627%
Mean of volatility	8.8023%	13.1699%	16.6214%	17.6328%
Volatility of volatility	1.0873%	1.6691%	2.6237%	3.3337%
Skewness	-0.3488	-0.2320	-0.2628	-0.3214
Kurtosis	3.0587	2.8589	2.7879	2.9088
Minimum	-33.5001%	-46.0088%	-55.4725%	-60.3300%
Maximum	32.1040%	46.1876%	56.5016%	55.8077%

Three-year Performance

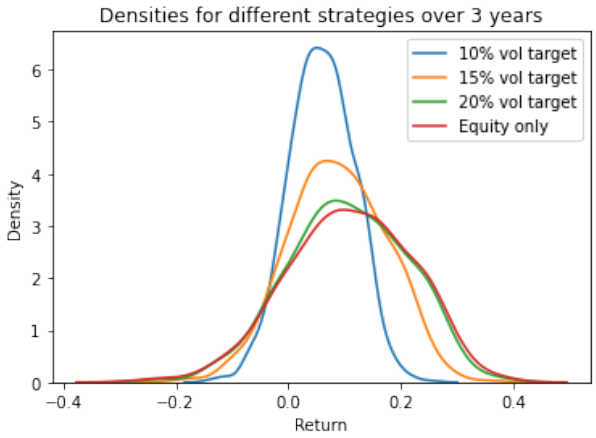

Fig. 5.25: 3-year returns

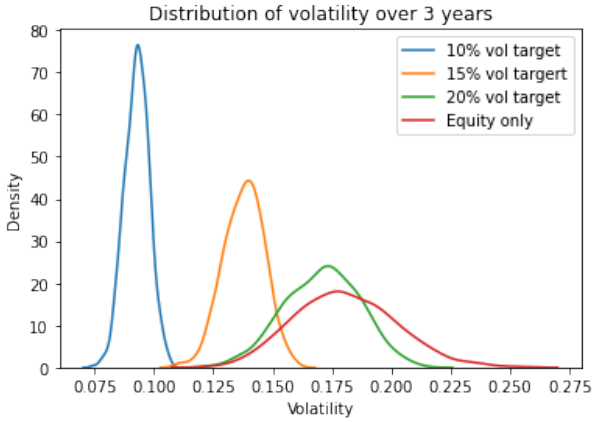

Fig. 5.26: 3-year volatility

Tab. 5.12: Statistics for volatility targeting strategies over three years

Statistic	10% target	15% target	20% target	Equity only
Mean return	5.9788%	8.7532%	10.5564%	10.9141%
Mean of volatility	9.2840%	13.7929%	17.0814%	18.0664%
Volatility of volatility	0.5179%	0.8613%	1.5903%	2.1366%
Skewness	-0.0584	-0.0509	-0.1559	-0.2347
Kurtosis	2.9797	2.8436	2.8178	2.9164
Minimum	-14.1057%	-20.7017%	-25.8108%	-29.3130%
Maximum	25.7383%	36.8075%	41.1602%	41.1486%

Five-year Performance

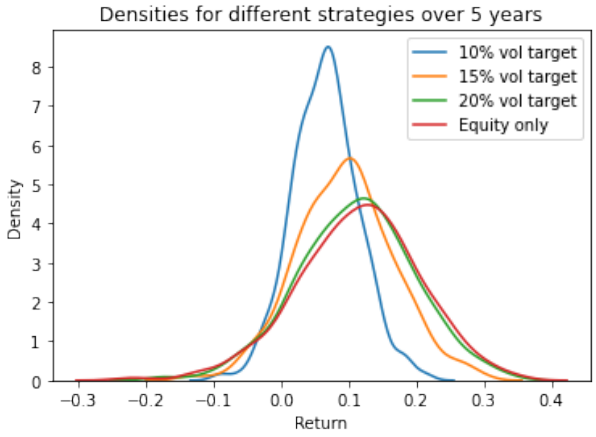

Fig. 5.27: 5-year returns

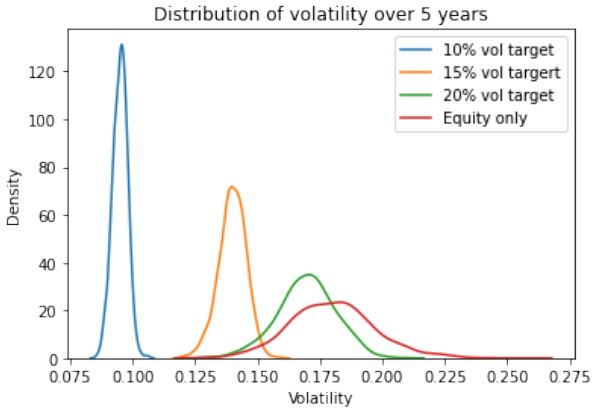

Fig. 5.28: 5-year volatility

Tab. 5.13: Statistics for volatility targeting strategies over five years

Statistic	10% target	15% target	20% target	Equity only
Mean return	6.5284%	9.4335%	11.0569%	11.4293%
Mean of volatility	9.4985%	13.9930%	16.9624%	18.0113%
Volatility of volatility	0.2990%	0.5538%	1.1693%	1.6878%
Skewness	0.0483	0.0028	-0.1638	-0.3107
Kurtosis	3.3085	3.1773	3.1763	3.4514
Minimum	-9.7357%	-14.5351%	-18.8936%	-23.4266%
Maximum	21.8382%	30.0869%	34.7505%	35.3795%

Results Discussion

There are a couple of interesting observations from the results. Firstly, the mean of the volatility estimate is close to the volatility target for each of the investment horizons. Furthermore, the volatility of volatility estimate decreases as the investment horizon increases. This indicates that the volatility targeting strategy improves for longer investment horizons. Investors can, therefore, expect a targeted volatility with a greater level of certainty for longer investment horizons.

Our results indicate that the volatility targeting strategy reduces the likelihood of extreme returns and reduces the volatility of volatility.

For all investment horizons, the risk and return of the portfolio increases as the volatility target increases. The 10% volatility target has the lowest risk when viewing the mean of volatility and volatility of volatility estimates, but also the lowest return. On the other hand, an equity only-holding strategy has the highest risk but also the highest expected return. Volatility targeting, therefore, gives investors an effective way of managing the downside risk of a portfolio, but limits the upside potential as shown by the minimum and maximum statistics.

Our findings are consistent with the results in Khuzwayo and Maré (2014). The simulation-based framework introduced in this chapter gives investors and fund managers a way of testing portfolio strategies for a wide variety of market conditions.

5.5 Conclusion

In this chapter, we calibrated the Heston (1993) and Bates (1996) SVJ models to historical S&P500 and FTSE/JSE Top40 returns using the EMM technique of Gallant and Tauchan (1996). First, we confirmed the accuracy of our implementation with Andersen et al. (2002) by calibrating the Heston (1993) and Bates (1996) SVJ models over the period 2 January 1953 to 31 December 1996. Our results confirmed that stochastic volatility and jumps are both required to characterise equity returns in the US equity market.

Next, we used the EMM to calibrate the Heston (1993) model over different periods to test the stability of the model parameters. Our results suggest that the calibration is sensitive to the input data used and not necessarily stable over time.

The EMM method was then used to calibrate the Heston (1993) and Bates (1996) SVJ models to FTSE/JSE Top40 returns over the period 30 June 1995 to 30 June 2022. Our calibration results suggest that both stochastic volatility and jumps are required to capture the behaviour of the South African equity market.

The final step was to show a practical application of the Bates (1996) SVJ model in the real-world measure. We performed a simulation-based study of various volatility targeting strategies and showed that the risk of a portfolio can be managed effectively by targeting a specific volatility and regularly allocating an investment between a risky asset and cash.

6 Static Hedging of Vanilla and Exotic Options in a South African Context

> *Never think that lack of variability is stability. Don't confuse lack of volatility with stability, ever.*
>
> — **Nassim Nicholas Taleb**

Keywords: Stochastic volatility double jump model · real-world measure · risk-neutral measure · calibration · replicating portfolio · static hedging

6.1 Introduction

This chapter[1] is dedicated to product development and considers the sale of vanilla and exotic financial derivatives in South Africa. We focus specifically on the trading of long-dated European call options and European spread call options for which no liquid market exists. Pricing these options is just one part of the challenge. Hedging, on the other hand, is an even bigger challenge.

Institutions wanting to sell long-dated European call options and European spread call options are faced with the challenge of buying assets to cover liabilities. In an ideal world, the portfolio manager will buy assets that match the risk sensitivities (Greeks) of the liabilities. Unfortunately, this is seldom the case as many liabilities have characteristics (longer maturities, for example) that cannot be matched perfectly by tradable assets. In South Africa, the equity derivatives market is typically very short-dated, which makes hedging a challenge.

Bowie and Carr (1994) and Derman et al. (1994) introduced the concept of static hedging, which replicates the value of the written option using standard exchange traded European options with varying strikes, maturities, and fixed portfolio weights. The advantage of static hedging

[1] This chapter is based on a paper (Levendis and Maré, 2023c) presented at the *AFRIC 2023: Actuarial, Finance, Risk and Insurance Congress*, Victoria Falls.

over traditional delta-hedging is that the hedging portfolio does not need to be rebalanced until one of the standard exchange traded options expires.

On whether market makers can sell long-dated European call options and European spread call options in South Africa and manage the risks effectively, we propose a simulation-based framework to test the performance of the static hedging program under numerous market conditions. We consider the stochastic volatility double jump (SVJJ) model of Duffie et al. (2000) to simulate the underlying equity prices under the real-world probability measure, $\mathbb{P}$, since the SVJJ model allows for jumps in both returns and volatility. Eraker et al. (2003) found strong evidence to support jumps in returns as well as volatility.

For each simulated path under $\mathbb{P}$, we price the vanilla European call options and European spread call option under the risk-neutral measure, $\mathbb{Q}$. To do this, we make use of the FFT of Hurd and Zhou (2010). Finally, we test two static hedging programs based on the work of Choie and Novomestky (1989) and Armstrong et al. (2018) to optimise the replicating portfolio weights.

The remainder of this chapter is structured as follows: Section 6.2 introduces the SVJJ model of Duffie et al. (2000). Section 6.3 focuses on the static hedging programs of Choie and Novomestky (1989) and Armstrong et al. (2018). Section 6.4 presents the static hedging results for the long-dated European call option and European spread call option, and Section 6.5 concludes the chapter.

6.2 Stochastic Volatility Double Jump Model

This section is split into two subsections. The first subsection introduces the SDE for the SVJJ model that will be used to simulate real-world equity price paths. The second subsection presents the characteristic function for the SVJJ model that will be used to price the vanilla European call options.

6.2.1 SVJJ Dynamics

The SVJJ model is an extension of the Bates (1996) SVJ model that adds correlated random jumps to the variance process. Under the $\mathbb{P}$-measure, the SVJJ model is given by the SDE:

$$dS(t) = (\mu - \lambda\mu_J)S(t)dt + \sqrt{v(t)}S(t)dW_x(t) + JS(t)dN(t),$$
$$dv(t) = (\alpha - \beta v(t)) + \sigma_v\sqrt{v(t)}dW_v(t) + ZdN(t),$$

$$dW_x(t)dW_v(t) = \rho_{x,v}dt,$$

where

$$\mu_J = \frac{\exp\left\{\mu_S + \frac{\sigma_S^2}{2}\right\}}{1 - \rho_J\mu_V} - 1,$$

and

$$Z \sim Exponential(\mu_V),$$

$$1 + J \mid Z \sim lognormal(\mu_S + \rho_J Z, \sigma_S^2),$$

with μ_V affecting the jump size of the variance, and ρ_J the correlation between the stock and variance jumps.

6.2.2 SVJJ Characteristic Function

The characteristic function for the SVJJ model is an extension of the characteristic function for the Heston (1993) and Bates (1996) models. From Poklewski-Koziell (2012), the characteristic function for the SVJJ model, defined under the $\mathbb{Q}$-measure, is given by the product of the Heston (1993) characteristic function and an independent jump component:

$$\phi_{SVJJ}(u) = \phi_H(u)\phi_J(u),$$

where

$$\phi_H(u) = e^{iu(x(0)+rT)+C(u,T)\frac{\alpha}{\beta}+D(u,T)v(0)},$$

and

$$C(u,T) = \beta\left[\left(\frac{Q - D_1}{2R}\right)T - \frac{2}{\sigma_v^2}\log\left(\frac{1 - Ge^{-D_1T}}{1 - G}\right)\right],$$

$$D(u,T) = \frac{Q - D_1}{2R}\left[\frac{1 - e^{-D_1T}}{1 - Ge^{-D_1T}}\right],$$

with

$$D_1 = \sqrt{Q^2 - 4PR},$$

$$G = \frac{Q - D_1}{Q + D_1},$$

$$P = \frac{-u^2 - iu}{2},$$

$$Q = \beta - \rho_{x,v}\sigma_v iu,$$
$$R = \frac{1}{2}\sigma_v^2.$$

Furthermore,

$$\phi_J(u) = e^{-\lambda T(1+iu\mu_J)+\lambda \exp\left\{iu\mu_S+\frac{\sigma_S^2(iu)^2}{2}\right\}\nu},$$

where

$$\nu = \frac{Q+D_1}{(Q+D_1)c-2\mu_V P} + \frac{4\mu_V P}{(D_1 c)^2-(2\mu_V P - Qc)^2}$$
$$\times \log\left[1-\frac{(D_1-Q)c+2\mu_V P}{2D_1 c}\left(1-e^{-D1T}\right)\right],$$
$$c = 1 - iu\rho_J\mu_V.$$

The SVJJ characteristic function will be used to price vanilla European call options under the $\mathbb{Q}$-measure using the FFT of Hurd and Zhou (2010). Note that the two-dimensional FFT of Hurd and Zhou (2010) for European spread call options reduces to the one-dimensional case for vanilla European call options when the second asset price is set to zero.

In the next section, we introduce the static hedging programs of Choie and Novomestky (1989) and Armstrong et al. (2018).

6.3 Static Hedging

This section is split into two subsections and introduces two static hedging programs that can be used to optimise the instrument weights in the replicating portfolio. The first subsection introduces the static hedging program of Choie and Novomestky (1989), and the second subsection focuses on the static hedging program of Armstrong et al. (2018).

6.3.1 Choie and Novometsky Optimisation

The static hedging program of Choie and Novomestky (1989) seeks to minimise the cost of setting up the replicating portfolio, subject to the value of the replicating portfolio being greater than or equal to the value of the target option at some future date. Mathematically, this can be expressed as:

$$\min_{\boldsymbol{B}} \sum_{i=1}^{n} C(i)B(i),$$

subject to

$$\sum_{i=1}^{n} F(ij)B(i) \geq Y(j), \quad j = 1, 2, ..., m,$$

where

$i = 1, 2, ..., n$:= the number of instruments in the replicating portfolio;

$j = 1, 2, ..., m$:= the price of the underlying asset at some future time;

$C(i)$:= the current price of the i^{th} instrument;

$B(i)$:= the number of units of the i^{th} instrument;

$F(ij)$:= the future price of the i^{th} instrument in state j; and

$Y(j)$:= the future price of the target option in state j.

In the next subsection, we introduce the static hedging program of Armstrong et al. (2018).

6.3.2 Armstrong et al. Optimisation

The static hedging program of Armstrong et al. (2018) seeks to minimise the difference between the value of the replicating portfolio and the target option at some future date, subject to the cost of the replicating portfolio being less than or equal to the initial wealth, *i.e.*, the premium received from the written option. Mathematically, this can be written as:

$$\min_{\boldsymbol{B}} \sum_{j=1}^{m} \Big(Y(j) - \sum_{i=1}^{n} F(ij)B(i) \Big)^2,$$

subject to

$$\sum_{i=1}^{n} C(i)B(i) \leq w,$$

where

$i = 1, 2, ..., n$:= the number of instruments in the replicating portfolio;

$j = 1, 2, ..., m$:= the price of the underlying asset at some future time;

$C(i)$:= the current price of the i^{th} instrument;

$B(i)$:= the number of units of the i^{th} instrument;

$F(ij)$:= the future price of the i^{th} instrument in state j;

$Y(j)$:= the future price of the target option in state j; and

$w :=$ the initial wealth, *i.e.*, the premium received.

In the next section, we present the static hedging results for a 5-year European call option and 1-year European spread call option based on the static hedging programs of Choie and Novomestky (1989) and Armstrong et al. (2018).

6.4 Results

This section is split into four subsections. The first subsection presents the calibration results for the SVJJ model to the FTSE/JSE Top40 index under the $\mathbb{P}$-measure. The second subsection contains the calibration results for the SVJJ model to the FTSE/JSE (Top40) implied volatility surface under the $\mathbb{Q}$-measure. The third subsection presents the static hedging results for a 5-year vanilla European call option written on the FTSE/JSE Top40 index; and, lastly, the fourth subsection shows the static hedging results for an arbitrary 1-year European spread call option.

6.4.1 SVJJ $\mathbb{P}$-Measure Calibration

The first step in setting up the simulation-based framework for static hedging is to calibrate the SVJJ model under the $\mathbb{P}$-measure to forecast future prices for the FTSE/JSE Top40.

Figure 6.1 below shows the historical closing prices for the FTSE/JSE Top40 index from 30 June 1995 to 30 October 2020:

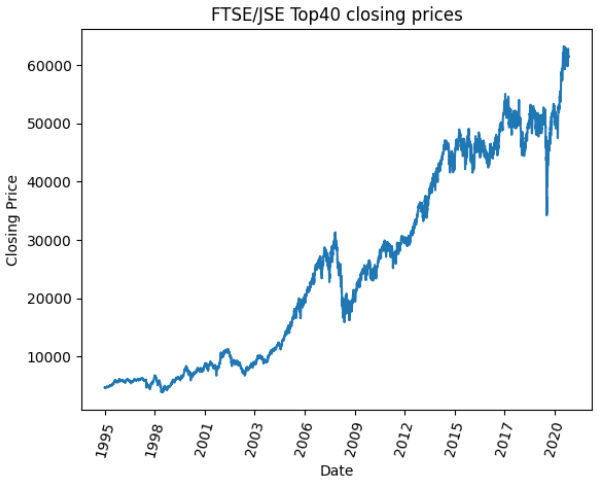

Fig. 6.1: FTSE/JSE Top40 closing prices

Using the EMM technique of Gallant and Tauchan (1996) that was discussed in Chapter 5, we calibrated the SVJJ model to daily returns from the FTSE/JSE Top40 index over the period 30 June 1995 to 30 October 2020. The calibrated parameters and goodness-of-fit statistic are shown in Table 6.1 below:

Tab. 6.1: SVJJ $\mathbb{P}$-parameters for FTSE/JSE Top40

Parameter	FTSE/JSE Top40
μ	0.1180
α	0.2888
β	6.0176
σ_v	0.4543
$\rho_{x,v}$	-0.9374
λ	4.7284
σ_S	0.0137
μ_V	0.0077
ρ_J	-0.3052
T$\hat{\theta}$	5.1022
$\chi^2_{0.05}$	5.9915

In the calibration, we set $\mu_S = 0$, since this parameter is generally insignificant and poorly identified as explained by Andersen et al. (2002).

From Table 6.1, the SVJJ model expects between four and five jumps per year. Furthermore, note the strong negative relationship between the stock and variance processes, and also the negative correlation between the stock and variance jumps. As a result, the model will produce a negative skew for the FTSE/JSE Top40.

Table 6.2 below compares the first four statistical moments as well as the minimum and maximum values from the SVJJ model with the daily returns from the FTSE/JSE Top40 index:

Tab. 6.2: SVJJ model daily statistics for the FTSE/JSE Top40

Statistic	FTSE/JSE Top40 index	SVJJ model
Mean	0.0385%	0.0406%
Std dev	1.3290%	1.1410%
Skewness	-0.4369	-0.2418
Kurtosis	9.4344	5.0463
Minimum	-0.1429	-0.0695
Maximum	0.0845	0.0592

The results indicate that the SVJJ model captures the mean and standard deviation well for the FTSE/JSE Top40 index, but underestimates the skewness and kurtosis. However, the goodness-of-fit statistic in Table 6.1 suggests that the SVJJ model is not rejected at a 5% level of significance. The SVJJ model is, therefore, a plausible data-generating model for the FTSE/JSE Top40 index.

Figure 6.2 below compares the distribution generated by the SVJJ model with the distribution of the FTSE/JSE Top40 index:

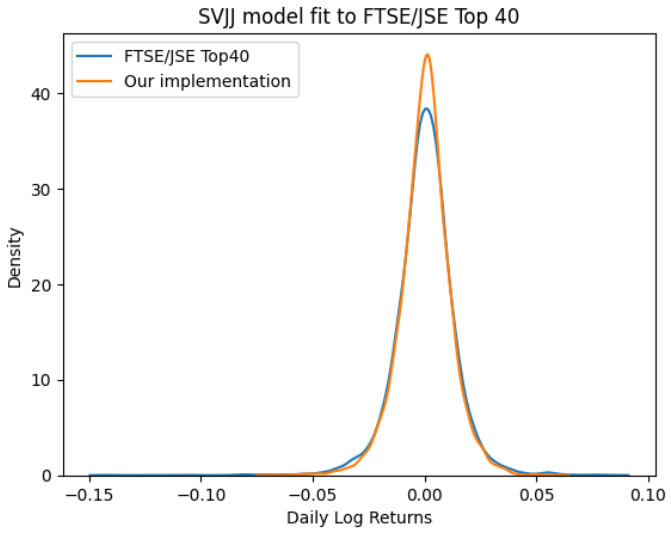

Fig. 6.2: SVJJ model versus FTSE/JSE Top40 distribution

The SVJJ model fits the historical distribution well. The $\mathbb{P}$-SVJJ model will be used to generate real-world sample paths for the FTSE/JSE Top40 index using Monte Carlo simulation. For each real-world path, the value of the written vanilla European call option, and values of the replicating options, must be calculated under the $\mathbb{Q}$-measure. This is the focus of the next subsection.

6.4.2 SVJJ $\mathbb{Q}$-Measure Calibration

For the purpose of this chapter, we assume a constant risk-free interest rate, $r = 7\%$. Using the SVJJ characteristic function in Section 6.2.2 and the FFT of Hurd and Zhou (2010) for European spread call options (reduced to one dimension by setting the second asset price to zero), The SVJJ model was calibrated to the FTSE/JSE (Top40) implied volatility surface on 16 November 2020. Table 6.3 below shows the calibrated parameters. Note that a tilde has been placed over each parameter to distinguish the $\mathbb{Q}$-parameters from the $\mathbb{P}$-parameters.

Tab. 6.3: SVJJ $\mathbb{Q}$-parameters for FTSE/JSE Top40

Parameter	FTSE/JSE Top40
r	0.0700
$\tilde{\alpha}$	0.0333
$\tilde{\beta}$	0.9995
$\tilde{\sigma_v}$	0.3827
$\tilde{\rho_{x,v}}$	-0.9205
$\tilde{\lambda}$	0.0583
$\tilde{\sigma_S}$	0.0058
$\tilde{\mu_V}$	0.0058
$\tilde{\rho_J}$	0.0097

Note that the $\mathbb{Q}$-parameters in Table 6.3 differ from the $\mathbb{P}$-parameters in Table 6.1. Grobler and Visagie (2019) explain that the returns distribution resulting from calibration to option prices can differ substantially from the historical returns distribution. The authors mention that a possible solution is to combine option prices and historical returns in the calibration procedure in order to minimise the discrepancy between the real-world and risk-neutral distributions. However, this generally leads to larger errors between the model prices and option prices.

Figure 6.3 below shows the fit of the SVJJ model to the FTSE/JSE (Top40) implied volatility surface on 16 November 2020:

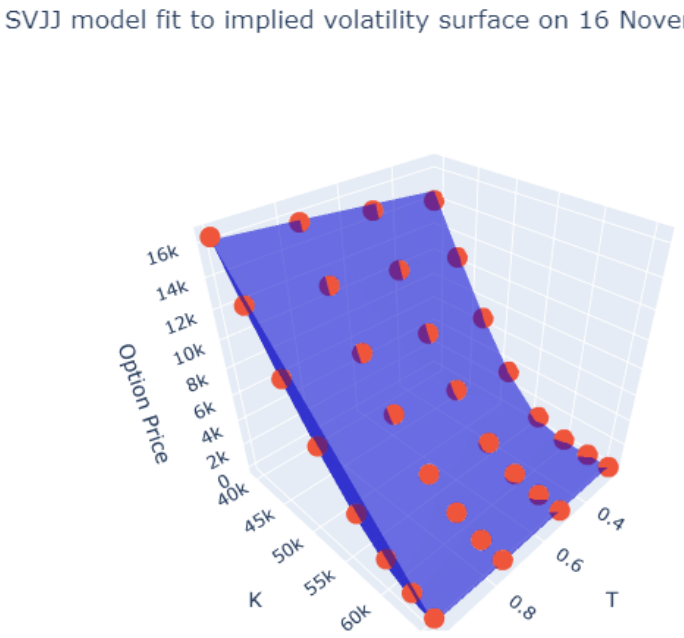

Fig. 6.3: SVJJ fit to FTSE/JSE Top40 implied volatility surface

The red dots represent the market quotes for FTSE/JSE Top40 European call options on 16 November 2020, and the blue surface represents the SVJJ model prices. Note that the SVJJ model reprices the exchange traded FTSE/JSE Top40 options well.

Before tackling the static hedging experiment, we first compared our implementation of the FFT for arbitrary European call options with the results obtained from a Monte Carlo simulation with 100,000 samples. Efficient pricing is important for the static hedging experiment, since option values must be calculated for multiple real-world paths. Monte Carlo simulation is computationally too expensive. Table 6.4 below compares the European call option prices for each numerical method:

Tab. 6.4: MC and FFT European call option prices under SVJJ model with $S(0) = 100,\ r = 0.1,\ v(0) = 0.04,\ \beta = 1,\ \alpha = 0.04,\ \sigma_v = 0.05,\ \rho_{x,v} = -0.5,\ \lambda = 5,\ \mu_S = 0,\ \sigma_S = 0.01,\ \rho_J = -0.3,\ \mu_V = 0.02,\ N = 256,\ \bar{u} = 40,\ \epsilon_1 = -3,\ \epsilon_2 = 1,\ T = 1.$

K	MC Price	FFT Price	Absolute Difference
20	81.906164	81.903234	0.002930
30	72.872779	72.855205	0.017574
40	63.773907	63.811614	0.037707
50	54.788865	54.795772	0.006907
60	45.908156	45.885217	0.022939
70	37.259369	37.253208	0.006161
80	29.160481	29.176272	0.015791
90	21.995527	21.975674	0.019853

Table 6.4 confirms that our implementation of the FFT was accurate. The FFT prices a single option in approximately 2.61 seconds, compared to 215.50 seconds in the case of Monte Carlo simulation.

For the static hedging experiment, we first consider the sale of a 5-year at-the-money vanilla European call option on the FTSE/JSE Top40 index on 16 November 2020. To hedge the sold option, we set up a static hedging portfolio consisting of 3-month, 6-month, 9-month, and 12-month exchange traded FTSE/JSE Top40 index options and cash. Note that the FTSE/JSE Top40 index option prices on 16 November 2020 are readily available from the option price surface in Figure 6.3. Their future prices can also be obtained by simulating real-world variations for the state variables from the $\mathbb{P}$-SVJJ model, and substituting these values in the $\mathbb{Q}$-SVJJ model.

The next subsection presents the static hedging results for the vanilla European call option.

6.4.3 Static Hedging Performance for Vanilla Call Option

The results in this subsection show the static hedging performance for a written 5-year at-the-money vanilla European call option. Since the longest maturity for the exchange traded FTSE/JSE Top40 options is generally 1 year, the replicating portfolio will need to be rolled as the options expire. As explained by Choie and Novomestky (1989), once the shortest dated option (3 months in our case) expires, the proceeds from the sale of the replicating portfolio will be used to purchase a new portfolio consisting of cash and 3-month, 6-month, 9-month, and 12-month FTSE/JSE Top40 options. This process repeats until the expiry of the written option; in this case, the 5-year FTSE/JSE Top40 European call option. It is important to note that each hedging interval is only for a period of 3 months.

Calculating the distribution of values for the 3-month option at the 3-month mark is simply $\max\left(S(0.25) - K, 0\right)$, where $S(0.25)$ are the real-world forecasts for the FTSE/JSE Top40 index from the $\mathbb{P}$-SVJJ model 3 months ahead, and K is the strike price of the option. The valuation of the 6-month, 9-month, and 12-month options at the 3-month mark is more complicated.

At the 3-month mark, the 6-month, 9-month, and 12-month options that were bought at inception have maturities of 3 months, 6 months, and 9 months respectively. Their values can be calculated by substituting the real-world forecasts, $S(0.25)$ and $v(0.25)$, into the SVJJ characteristic function along with the $\mathbb{Q}$-SVJJ parameters in Table 6.3. The FFT of Hurd and Zhou (2010) can then be used to calculate the option values.

The valuation of the 5-year FTSE/JSE Top40 index option at the 3-month mark follows a similar process, where the maturity of the option at this point is 4.75 years. This process gets repeated every quarter.

Table 6.5 below details the information for the written European call option on 16 November 2020:

Tab. 6.5: Market information for European call option on 16 November 2020

Option sale date	16 November 2020
Underlying	FTSE/JSE Top40 index
$S(0)$	52552
K	52552
T	5

Figure 6.4 below shows the real-world distribution for the FTSE/JSE Top40 index at $t = 0.25$ generated from the $\mathbb{P}$-SVJJ model in Table 6.1 with 10,000 Monte Carlo samples:

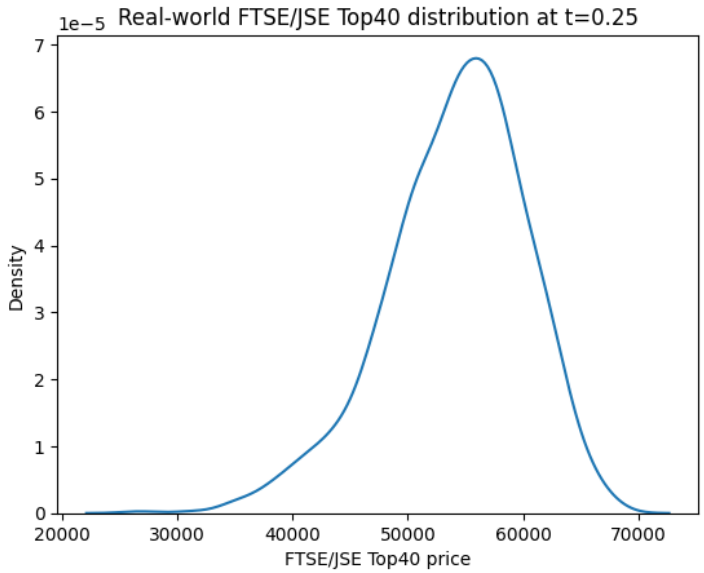

Fig. 6.4: Real-world FTSE/JSE Top40 distribution at $t = 0.25$

The option writer has the entire FTSE/JSE Top40 option price surface at his disposal when faced with the challenge of hedging the 5-year at-the-money European call option. The seller's aim is to find the optimal quantity for each exchange traded option on the option price surface to hedge his position at $t = 0.25$.

The static hedging results based on the optimisation routines of Choie and Novomestky (1989) and Armstrong et al. (2018) are shown below.

Choie and Novometsky Optimisation

Using the optimisation program of Choie and Novomestky (1989) discussed in Section 6.3.1, Figure 6.5 below shows the optimised quantities on 16 November 2020 for the exchange traded options based on the real-world distribution for the FTSE/JSE Top40 index at $t = 0.25$ in Figure 6.4.

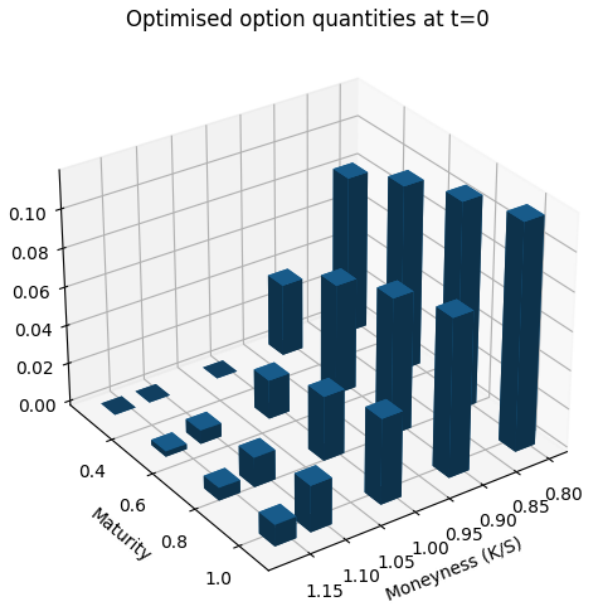

Fig. 6.5: Replicating option quantities based on Choie and Novometsky optimisation

Table 6.6 below shows the replicating option quantities for the first hedging interval, $t = 0$ to $t = 0.25$, expressed as a percentage:

Tab. 6.6: Replicating option quantity table based on Choie and Novometsky optimisation

Moneyness/Maturity	**0.25**	**0.5**	**0.75**	**1**
0.8	8.26%	9.70%	10.80%	11.70%
0.9	3.75%	5.75%	7.11%	8.18%
1	0.00%	2.02%	3.34%	4.41%
1.1	0.00%	0.69%	1.53%	2.39%
1.15	0.00%	0.27%	0.65%	1.10%

The replicating portfolio is skewed more towards in-the-money options, and the optimised cash balance was $R6,500$. The cost of setting up the replicating portfolio on 16 November 2020 was $R14,170$. On the other hand, the premium received from the sale of the 5-year at-the-money option was $R13,656$, which was calculated from the $\mathbb{Q}$-SVJJ model, *i.e.*, consistent with the market prices on 16 November 2020. Therefore, the option writer recorded an upfront loss of $R13,656 - R14,170 = -R514$.

Figure 6.6 below compares the value of the replicating portfolio based on the option quantities in Table 6.6 and a cash balance of $R6,500$ with the value of the target option at $t = 0.25$.

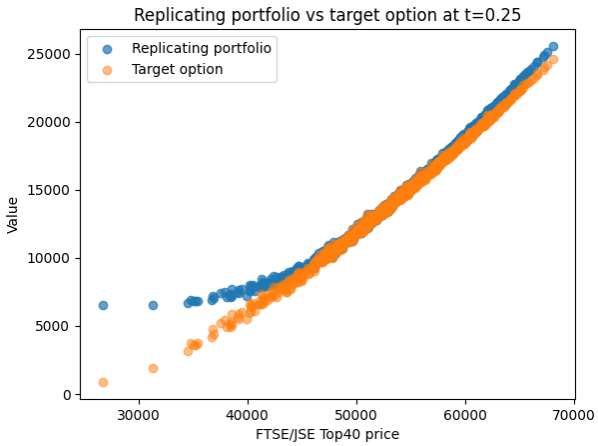

Fig. 6.6: Portfolio versus target option based on Choie and Novometsky optimisation

Next, we show the static hedging results based on the optimisation routine of Armstrong et al. (2018). The results will be discussed thereafter.

Armstrong et al. Optimisation

Using the optimisation routine of Armstrong et al. (2018) discussed in Section 6.3.2, Figure 6.7 below shows the optimised quantities for the exchange traded options based on real-world distribution for the FTSE/JSE Top40 index at $t = 0.25$ in Figure 6.4:

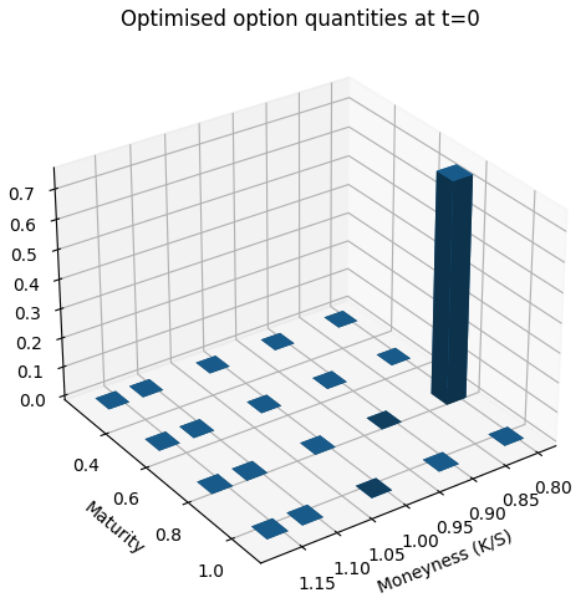

Fig. 6.7: Replicating option quantities based on Armstrong et al. optimisation

Once again, the option quantities were calculated for the first hedging interval, $t = 0$ to $t = 0.25$, and expressed as a percentage. Note that the optimisation returns a single option on the FTSE/JSE Top40 option price surface (80% moneyness and 9 months to maturity), which is significantly different from the results obtained by using the Choie and Novomestky (1989) routine.

The cost of setting up the replicating portfolio on 16 November 2020 was $R13,656$, which is exactly equal to the premium received. The optimised cash balance was $R6,425$.

Figure 6.8 below compares the value of the replicating portfolio based on the option quantities in Figure 6.7 and a cash balance of $R6,425$ with the value of the target option at $t = 0.25$:

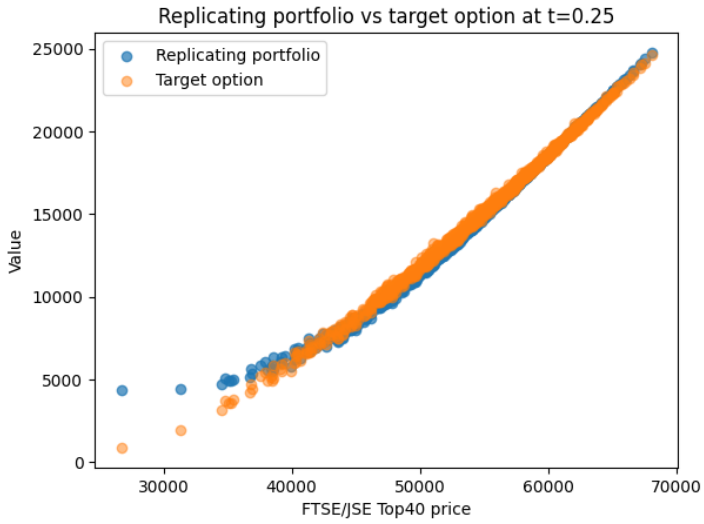

Fig. 6.8: Portfolio versus target option based on Armstrong et al. optimisation

The results are discussed next.

Results Discussion

Recall that the static hedging program of Choie and Novomestky (1989) seeks to minimise the cost of setting up the replicating portfolio, subject to the value of the replicating portfolio being greater than or equal to the value of the target option at some future date; 3 months in our case. Figure 6.6 illustrates that the constraint was met - the value of the replicating portfolio was greater than or equal to the value of the target option in each of the real-world FTSE/JSE Top40 states at $t = 0.25$.

Figure 6.5 showed that the Choie and Novomestky (1989) optimisation produced replicating option quantities across most of the FTSE/JSE Top40 option price surface, with the quantities skewed more towards in-the-money options. A possible explanation for this is that the value of the target option is quite sensitive to price movements in the FTSE/JSE Top40, *i.e.*, delta. Hence, the replicating portfolio is skewed more towards in-the-money options since they have the highest delta.

The cost of setting up the replicating portfolio was slightly more expensive than the upfront premium received; $R14,170$ versus $R13,656$. The option writer, therefore, recorded an upfront loss.

The static hedging performance based on the Choie and Novomestky (1989) optimisation also deteriorated at the tails of the FTSE/JSE Top40 distribution.

The second routine tested was the optimisation of Armstrong et al. (2018). Recall that this optimisation seeks to minimise the difference between the value of the replicating portfolio and the target option at some future date (3 months in our case), subject to the cost of the replicating portfolio being less than or equal to the premium received from the written option. The optimisation returned a cost that matched the premium from the written option exactly.

Figure 6.7 showed that the optimisation of Armstrong et al. (2018) returned a single in-the-money option on the FTSE/JSE Top40 option price surface. Based on this option and a cash balance of $R6,425$, Figure 6.8 showed that there were instances where the value of the replicating portfolio was less than the value of the target option at $t = 0.25$.

In summary, the choice of optimisation routine can produce substantially different quantities for the instruments in the replicating portfolio. The replicating portfolio based on the Choie and Novomestky (1989) optimisation might be slightly more expensive to set up than the premium received, but ensures that the value of the replicating portfolio is greater than or equal to the value of the target option for the state variables considered at some future date. Alternatively, the cost of setting up the replicating portfolio based on the optimisation of Armstrong et al. (2018) is equal to the upfront premium received. The risk is that the value of the replicating portfolio might be less than the value of the target option at some future date.

Considering the complexity of hedging an option written on the FTSE/JSE Top40, which exhibits factors such as stochastic volatility and jumps, the static hedging approach is simple and shows promising results.

In the next subsection, we test the static hedging performance for an arbitrary 1-year European spread call option.

6.4.4 Static Hedging Performance for Spread Call Option

This subsection presents the static hedging results for an arbitrary 1-year European spread call option, hedged with vanilla FTSE/JSE Top40 European call options. Note that the underlying instrument used to hedge the spread option is not necessarily the same as the underlying instruments in the spread option.

Figure 6.9 below shows the real-world distribution for the spread (S_1-S_2) at $t = 0.25$, generated from the $\mathbb{P}$-SVJJ model with parameters shown in Table 6.7 below:

Tab. 6.7: SVJJ $\mathbb{P}$-parameters for S_1 and S_2

Parameter	$\boldsymbol{S_1}$	$\boldsymbol{S_2}$
$S_i(0)$ $(i = 1, 2)$	100	96
μ	0.13	0.11
α	0.2888	0.2888
β	6.0176	6.0176
σ_v	0.4543	0.4543
$\rho_{x,v}$	-0.9374	-0.9374
λ	4.7284	4.7284
σ_S	0.0137	0.0137
μ_V	0.0077	0.0077
ρ_J	-0.3052	-0.3052

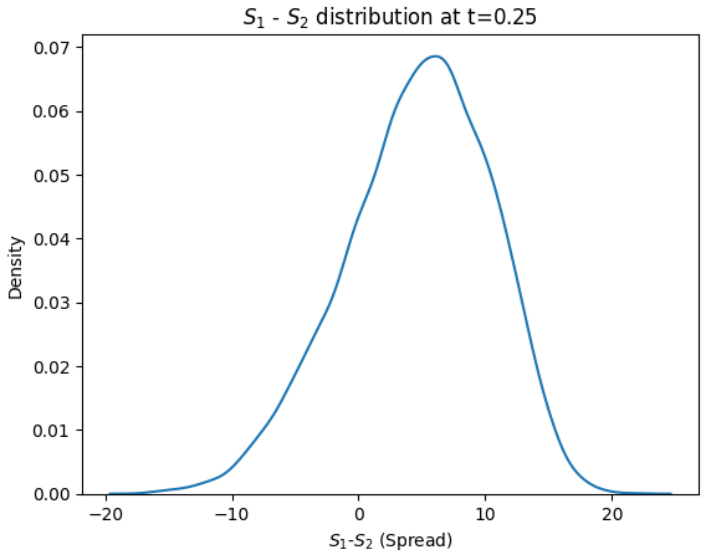

Fig. 6.9: Real-world distributions for S_1 - S_2 at $t = 0.25$

We further set the correlation between the Brownian motions driving S_1 and the FTSE/JSE Top40 equal to 1, and similar for S_2 and the FTSE/JSE Top40.

Figure 6.10 below shows the relationship between the FTSE/JSE Top40 price and the spread (S_1-S_2) based on 10,000 simulations at $t = 0.25$:

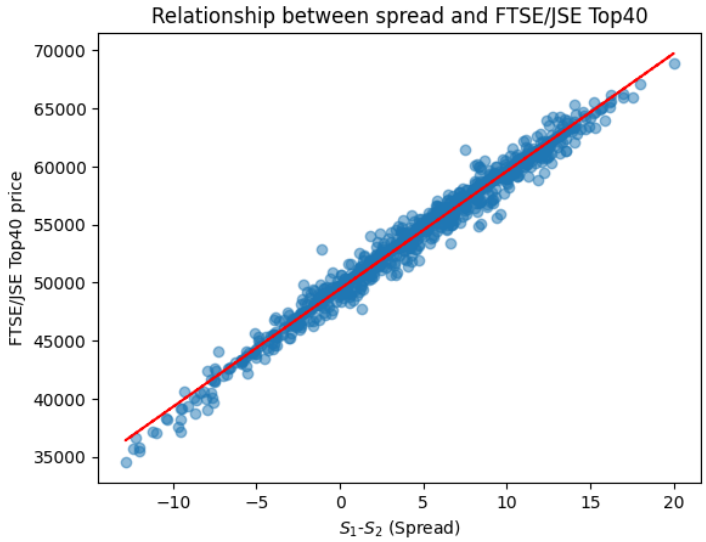

Fig. 6.10: Relationship between FTSE/JSE Top40 and $S_1 - S_2$ at $t = 0.25$

A linear relationship between the price of the underlying instrument used to hedge the spread option and the spread (S_1-S_2) must exist in order for the static hedge to work. Correlation is a key risk when hedging European spread call options with vanilla European call options.

Table 6.8 details the market information for the spread call option on 16 November 2020:

Tab. 6.8: Market information for European spread call option on 16 November 2020

Option sale date	16 November 2020
Underlying hedge instrument	FTSE/JSE Top40 index
$S_{Top40}(0)$	52552
$S_1(0)$	100
$S_2(0)$	96
K	3
T	1

The static hedging results for the European spread call option are discussed next.

Choie and Novometsky Optimisation

Using the optimisation routine of Choie and Novomestky (1989), Figure 6.11 below shows the optimised quantities for the FTSE/JSE Top40 options based on the distribution in Figure 6.9:

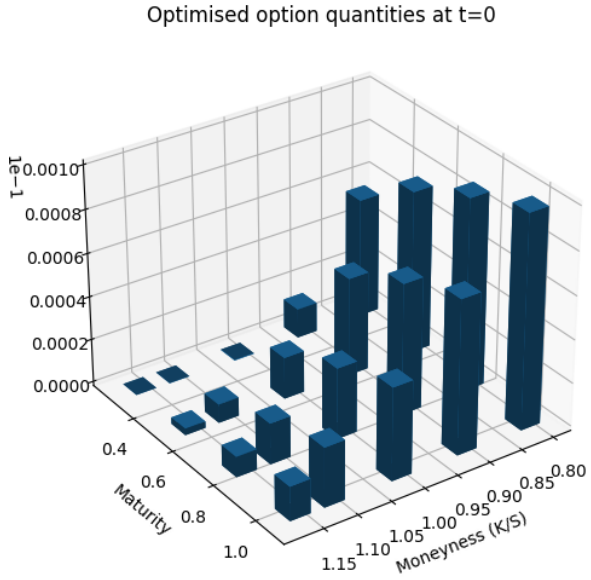

Fig. 6.11: Replicating option quantities based on Choie and Novometsky optimisation

Table 6.9 below shows the replicating option quantities for the first hedging interval, $t = 0$ to $t = 0.25$, expressed as a percentage:

Tab. 6.9: Replicating option quantity table based on Choie and Novometsky optimisation

Moneyness/Maturity	0.25	0.5	0.75	1
0.8	0.0055%	0.0075%	0.0089%	0.0099%
0.9	0.0014%	0.0045%	0.0060%	0.0071%
1	0.0000%	0.0019%	0.0033%	0.0043%
1.1	0.0000%	0.0009%	0.0019%	0.0027%
1.15	0.0000%	0.0003%	0.0009%	0.0016%

The replicating portfolio is skewed more towards in-the-money options, and the optimised cash balance was $R4$. The European spread call option value was calculated from the three-factor stochastic volatility model of Dempster and Hong (2002), and returned an option premium of $R7.48$. The cost of setting up the replicating portfolio based on the optimisation of Choie

and Novomestky (1989) was $R10.23$. Hence, the option writer recorded an upfront loss of $R7.48 - R10.23 = -R2.75$.

Figure 6.12 below compares the value of the replicating portfolio with the value of the European spread call option at $t = 0.25$:

Fig. 6.12: Portfolio versus target option based on Choie and Novometsky optimisation

Note that the value of the replicating portfolio is greater than or equal to the value of the European spread call option for all states at $t = 0.25$.

The static hedging results for the European spread call option based on the optimisation of Armstrong et al. (2018) are discussed next.

Armstrong et al. Optimisation

Using the optimisation program of Armstrong et al. (2018), Figure 6.13 below shows the optimised quantities for FTSE/JSE Top40 options based on the real-world spread distribution in Figure 6.9.

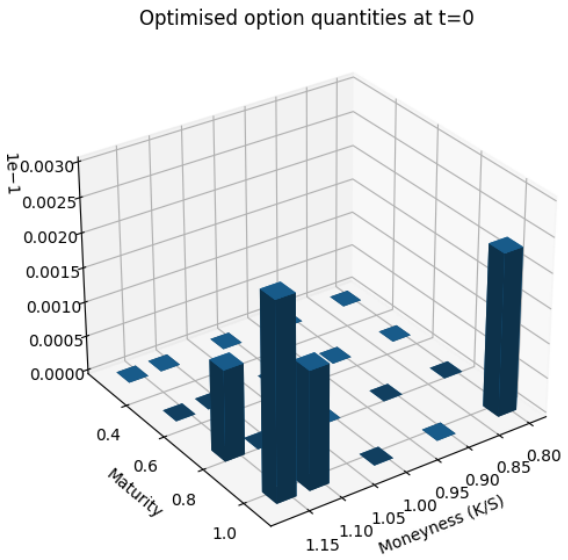

Fig. 6.13: Replicating option quantities based on Armstrong et al. optimisation

Again, the option quantities were calculated for the first hedging interval, $t = 0$ to $t = 0.25$. Note that the optimisation of Armstrong et al. (2018) returns only a small number of options on the FTSE/JSE Top40 surface (115% moneyness with 9 months to maturity and 80%, 110%, and 115% moneyness with 1 year to maturity). The optimised cash balance was $R3$, and the cost of setting up the replicating portfolio was $R7.48$, *i.e.*, equal to the premium received from the written European spread call option.

Figure 6.14 compares the value of the replicating portfolio with the value of the European spread call option at $t = 0.25$:

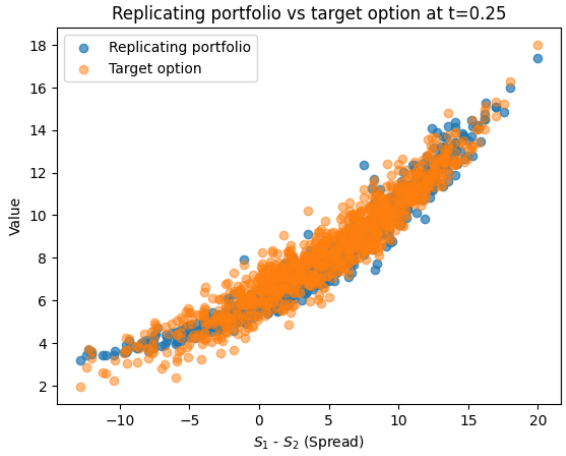

Fig. 6.14: Portfolio versus target option based on Armstrong et al. optimisation

Note that there are instances where the value of the replicating portfolio is less than the value of the European spread call option at $t = 0.25$.

The results are discussed next.

Results Discussion

For the European spread call option, the optimisation of Choie and Novomestky (1989) returned a replicating portfolio with option quantities spanning almost the entire FTSE/JSE Top40 option price surface. Conversely, the optimisation of Armstrong et al. (2018) returned only a small number of options on the FTSE/JSE Top40 option price surface.

Based on the optimisation of Choie and Novomestky (1989), the cost of setting up the replicating portfolio was $R10.23$. The premium received from the written option was $R7.48$, hence, the option writer recorded an upfront loss.

The cost of the replicating portfolio based on the optimisation of Armstrong et al. (2018) was exactly equal to the premium received. However, the optimisation did not guarantee a replicating portfolio value that was greater than or equal to the value of the European spread call option at a future date.

It is important to note that the static hedge will only work if the price of the hedging instrument is strongly and positively correlated with the spread generated by the two underlying instruments in the spread option. For European spread call options written on two stocks that form part of the FTSE/JSE Top40 index, this may very well be the case. We suggest this as an area for future research.

6.5 Conclusion

The purpose of this chapter was to link the $\mathbb{P}$- and $\mathbb{Q}$-probability measures for stochastic volatility models. To achieve this, we considered a portfolio risk management problem, *i.e.*, static hedging of a long-dated European call option and European spread call option in South Africa.

In risk management applications, the $\mathbb{P}$-measure is typically used to generate real-world events that can affect the value of a portfolio. Therefore, we calibrated the $\mathbb{P}$-SVJJ model to historical FTSE/JSE Top40 returns in order to generate share price and volatility shocks.

Pricing derivatives is done under the $\mathbb{Q}$-measure due to the principle of no arbitrage. Therefore, we calibrated the $\mathbb{Q}$-SVJJ model to the FTSE/JSE Top40 option price surface in order to price the options.

The link between $\mathbb{P}$ and $\mathbb{Q}$ was introduced when we simulated future option values. The process followed was to consider the current FTSE/JSE Top40 option price surface on 16 November 2020 and apply shocks to the underlying state variables (stock price and volatility), where the shocks were produced by the $\mathbb{P}$-SVJJ model over a 3-month period. We then revalued the options for each of the scenarios using the $\mathbb{Q}$-SVJJ model to produce a distribution of option values at $t = 0.25$. Note that this process is similar to VaR.

For the static hedge of the long-dated European call option, we applied the optimisation programs of Choie and Novomestky (1989) and Armstrong et al. (2018) to calculate the instrument weights in the replicating portfolio. The Choie and Novomestky (1989) optimisation produced weights across most of the FTSE/JSE option price surface and guaranteed that the value of the replicating portfolio was greater than or equal to the value of the long-dated European call option at a future date. However, the cost of setting up the replicating portfolio was more expensive than the premium received, leading to an upfront loss. The Armstrong et al. (2018) optimisation returned a single option on the FTSE/JSE option price surface with cost equal to the premium received, but did not guarantee a replicating portfolio value that was greater than or equal to the value of the long-dated European call option at a future date. The difference between the two optimisation routines is the timing of the loss.

The static hedging results for the European spread call option were similar to the results for the long-dated European call option. It is important to note that the key risk when hedging a European spread call option with vanilla European call options is correlation. An area for future research might be to consider a European spread call option that is struck on two stocks that form part of the FTSE/JSE Top40 index.

Static hedging shows promising results in the South African market and option writers may find that static hedging provides a cheaper and more effective solution than traditional delta-hedging.

7 Conclusion

In this thesis, we applied stochastic volatility models to three core problems in the financial services industry: pricing derivatives embedded in life insurance contracts; multi-asset option pricing; and portfolio risk management.

The research questions posed in Section 1.2 are answered below:

1. **What are the important risk factors to consider when pricing long-dated options**?

   In Chapter 2, we extended the Heston-Hull-White model of Grzelak and Oosterlee (2011) to include stochastic mortality and used this model to price long-dated GMMB and GMDB products. Our results showed that stochastic interest rates are the dominant risk driver when pricing these products. Furthermore, we proposed both a discrete-time AR(1)-ARCH(1) model and continuous-time CIR++ model to capture mortality risk - an important risk factor for products that are contingent on survival or death. We showed that the AR(1)-ARCH(1) model produced a good fit to historical mortality rates, where the CIR++ model was able to reproduce the input survival probability curve.

   In Chapter 4, we derived the two-asset Heston-Hull-White model and used this model to price long-dated spread options with the two-dimensional FFT of Hurd and Zhou (2010). Our results showed that stochastic interest rates have a significant impact on long-dated out-of-the-money European spread call options. The correlation between the stock and interest processes also plays an important role - positive correlation increases the option value, where negative correlation decreases the option value.

2. **Does the FFT lead to significant time saving, but similar accuracy, compared to Monte Carlo simulation for numerival evaluation of option values?**

   In Chapter 3, we showed that the two-factor gBm model and two-dimensional FFT of Hurd and Zhou (2010) produced accurate "worst-of-2" call option prices compared to the Stulz (1982) analytic price; a result that contradicts a previous finding in the literature. Furthermore, we implemented the three-factor stochastic volatility model of Dempster and Hong (2002) with the two-dimensional FFT and obtained "worst-of-2" call option prices that aligned closely with the results from a Monte Carlo simulation. Lastly, we showed that the two-dimensional FFT converges at an exponential rate, which leads to significant computational savings compared to Monte Carlo simulation.

In Chapter 4, we derived the characteristic function for the two-asset Heston-Hull-White model and implemented the two-dimensional FFT of Hurd and Zhou (2010) to price spread options with this model. We showed that the two-dimensional FFT is significantly faster than Monte Carlo simulation and yields similar accuracy. The two-dimensional FFT converged to the solution in approximately 3.23 seconds, where the Monte Carlo simulation took approximately 168 seconds to converge.

3. **Are jumps an important risk factor to consider when modelling equity returns?**

   In Chapter 5, we calibrated the Heston (1993) and Bates (1996) SVJ models to historical S&P500 and FTSE/JSE Top40 returns, i.e., the real-world probability measure. Our results showed that the goodness-of-fit statistic improved when jumps were added to the underlying model.

   In Chapter 6, we calibrated the SVJJ model of Duffie et al. (2000) to historical FTSE/JSE Top40 returns. The calibration was not rejected at a 5% level of significance, indicating that jumps are an important risk factor to consider when modelling equity returns.

4. **Is volatility targeting an effective trading strategy?**

   In Chapter 5, we proposed a simulation-based framework to test various volatility targeting strategies. Our results showed that volatility targeting is an effective risk management tool that can be used to manage downside risk. However, it was noted that lower volatility targets sacrifice upside return.

5. **Can we hedge long-dated options with a combination of other derivatives and cash?**

   In Chapter 6, we implemented a static hedging strategy for a long-dated European call option and European spread call option in South Africa. First, we calibrated the SVJJ model to historical FTSE/JSE Top40 returns, i.e., the real-world probability measure. We then calibrated the SVJJ model to the FTSE/JSE Top40 implied volatility surface (the risk-neutral probability measure) in order to price the options. Next, we simulated numerous paths for the underlying equity price, valued the options, and applied the static hedging programs of Choie and Novomestky (1989) and Armstrong et al. (2018). We showed that the static hedging programs replicated the target option well, but the upfront premium received from the written option did not necessarily ensure that the value of the replicating portfolio was at least equal to the value of the target option at a future date.

In this thesis, we showed that stochastic volatility models have numerous applications in the financial services industry. Combining stochastic volatility models in the $\mathbb{P}$- and $\mathbb{Q}$-measure, and applying the FFT and EMM numerical methods, may give market participants an edge to identify profitable opportunities and improve risk management practices.

As financial markets evolve and buzzwords such as “machine learning”, “AI”, and “blockchain” become more common, one must not forget to appreciate the forces (stochastic volatility) that drive them. There is truth in the following saying by Drishti Bablani: “*Of course things change course, but in the end, they circle back to the source.*”

Bibliography

Aït-Sahalia, Y., & Jacod, J. (2009). Testing for jumps in a discretely observed process. *Annals of Statistics*, *37*(1), 184–222 (cit. on p. 76).

Aït-Sahalia, Y., Jacod, J., & Li, J. (2012). Testing for jumps in noisy high frequency data. *Journal of Econometrics*, *168*(2), 207–222 (cit. on p. 76).

Alam, M., & Uddin, G. (2009). Relationship between Interest Rate and Stock Price: Empirical Evidence from Developed and Developing Countries. *International Journal of Business and Management*, *4*(3), 43–51 (cit. on p. 61).

Alfeus, M., & Schlögl, E. (2018). On Numerical Methods for Spread Options. *Research Paper Series 388, Quantitative Finance Research Centre, University of Technology, Sydney* (cit. on pp. 54, 69).

Andersen, T., Benzoni, L., & Lund, J. (2002). An Empirical Investigation of Continuous-Time Equity Return Models. *The Journal of Finance*, *57*(3), 1239–1284 (cit. on pp. 2, 6, 76, 78, 84, 86–91, 103, 111).

Andersen, T., Chung, H., & Sorensen, B. (1999). Efficient method of moments of a stochastic volatility model: A Monte Carlo study. *Journal of Econometrics*, *91*(1), 61–87 (cit. on pp. 82, 83).

Armstrong, J., Pennanen, T., & Rakwongwan, U. (2018). Pricing index options by static hedging under finite liquidity. *International Journal of Theoretical and Applied Finance*, *21*(6) (cit. on pp. 6, 106, 108–110, 116, 118, 120, 124–127, 129).

Balotta, L., Eberlein, E., Schmidt, T., & Zeineddine, R. (2020). Variable annuities in a Lévy-based hybrid model with surrender risk. *Quantitative Finance*, *20*(5), 867–886 (cit. on p. 10).

Bates, D. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *The Review of Financial Studies*, *9*(1), 69–107 (cit. on pp. 1–3, 6, 7, 76, 77, 79, 81, 82, 84, 86, 89–91, 94–97, 103, 104, 106, 107, 129).

Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the 1976 Meeting of the American Statistical Association, Business and Economic Statistics Section*, 177–181 (cit. on p. 32).

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654 (cit. on pp. 1, 7, 9, 40, 75).

Bowie, J., & Carr, P. (1994). Static Simplicity. *Risk*, *7*, 18–20 (cit. on p. 105).

Brigo, D., & Mercurio, F. (2001). On Deterministic Shift Extensions of Short Rate Models. *Available at SSRN: https://ssrn.com/abstract=292060 or http://dx.doi.org/10.2139/ssrn.292060* (cit. on pp. 5, 8, 10, 17).

Browne, B., Duchassaing, J., & Suter, F. (2009). Longevity: A 'Simple' Stochastic Modelling of Mortality. *British Actuarial Journal*, *15*(S1), 249–265 (cit. on p. 39).

Brownlees, C., Engle, R., & Kelly, B. (2012). A practical guide to volatility forecasting through calm and storm. *The Journal of Risk*, *14*(2), 3–22 (cit. on p. 97).

Cairns, A., Blake, D., & Dowd, K. (2008). Modelling and Management of Mortality Risk: A Review. *Scandinavian Actuarial Journal*, *2*(3), 79–113 (cit. on pp. 10, 17).

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, *2*(4), 61–73 (cit. on pp. 2, 5, 9, 10, 15, 29, 46, 49, 60).

Choie, K., & Novomestky, F. (1989). Replication of long-term with short-term options. *The Journal of Portfolio Management*, *15*(2), 17–19 (cit. on pp. 6, 106, 108, 110, 115, 116, 119, 120, 123, 126, 127, 129).

Cont, R. (2001). Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, *1*, 223–236 (cit. on pp. 7, 11, 48, 75).

Cooley, J., & Tukey, J. (1965). An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of Computation*, *19*(90), 297–301 (cit. on p. 46).

Cox, J., Ingersoll, J., & Ross, S. (1985). A Theory of the Term Structure of Interest Rates. *Econometrica*, *53*(2), 385–407 (cit. on pp. 8, 10, 12, 16, 17, 20, 48, 52, 64).

Dempster, M., & Hong, S. (2002). Spread Option Valuation And The Fast Fourier Transform. *Mathematical Finance-Bachelier Congress 2000*, 203–220 (cit. on pp. 6, 48, 51, 52, 57, 59–62, 71, 123, 128).

Derman, E., Ergener, D., & Kani, I. (1994). Forever Hedged. *Risk*, *7*, 139–145 (cit. on p. 105).

Duffie, D., Pan, J., & Singleton, K. (2000). Transform Analysis and Asset Pricing for Affine Jump-diffusions. *Econometrica*, *68*(6), 1343–1376 (cit. on pp. 2, 4, 6, 9, 13, 63, 106, 129).

Eberlein, E., Glau, K., & Papapantoleon, A. (2010). Analysis of Fourier Transform Valuation Formulas and Applications. *Applied Mathematical Finance*, *17*(3), 211–240 (cit. on pp. 46, 49–51, 56).

El Euch, O., Gatheral, J., & Rosenbaum, M. (2019). Roughening Heston. *Risk*, 84–89 (cit. on pp. 2, 76).

Eraker, B., Johannes, M., & Polson, N. (2003). The impact of jumps in volatility and returns. *The Journal of Finance*, *58*(3), 1269–1300 (cit. on p. 106).

Fang, F., & Janssens, B. (2007). Characteristic function of the hybrid Heston–Hull–White model. *In: Proceedings of the 58th European Study Group Mathematics with Industry. Utrecht. http://www.math.uu.nl/swi2007/proc2007.pdf*, 107–115 (cit. on p. 9).

Fang, F., & Oosterlee, C. (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, *31*(2), 826–848 (cit. on p. 10).

Feng, R. (2018). *An introduction to computational risk management of equity-linked insurance*. Taylor & Francis, Boca Raton, Florida. (Cit. on pp. 20, 37, 43).

Flint, E., & Maré, E. (2017). Estimating option-implied distributions in illiquid markets and implementing the Ross recovery theorem. *South African Actuarial Journal*, *17*(1), 1–28 (cit. on p. 27).

Flint, E., Ochse, E., & Polakow, D. (2014). Estimating long-term volatility parameters for market-consistent models. *South African Actuarial Journal*, *14*(1), 19–72 (cit. on pp. 27, 31).

Gallant, A., & Nychka, D. (1987). Semi-nonparametric maximum likelihood estimation. *Econometrica*, *55*(2), 363–390 (cit. on p. 82).

Gallant, A., & Tauchan, G. (1996). Which Moments to Match? *Econometric Theory*, *12*(4), 657–681 (cit. on pp. 6, 76, 77, 82, 103, 111).

Gatheral, J. (2006). *The volatility surface: A practitioner's guide*. John Wiley & Sons, New York. (Cit. on pp. 75, 76).

Giacometti, R., Betocchi, M., Rachev, S., & Fabozzi, F. (2012). A comparison of the Lee-Carter model and AR-ARCH model for forecasting mortality rates. *Insurance: Mathematics and Economics*, *50*(1), 85–93 (cit. on p. 19).

Glasserman, P. (2003). *Monte carlo methods in financial engineering*. In: Stochastic Modelling; Applied Probability, Springer-Verlag, New York. (Cit. on p. 59).

González-Urteaga, A. (2012). Further empirical evidence on stochastic volatility models with jumps in returns. *The Spanish Review of Financial Economics*, *10*(1), 11–17 (cit. on p. 79).

Grobler, G., & Visagie, I. (2019). On the discrepancy between the objective and risk neutral densities in the pricing of European options. *ORiON*, *35*(1), 33–56 (cit. on pp. 76, 113).

Grzelak, L., & Oosterlee, C. (2011). On the Heston Model with Stochastic Interest Rates. *SIAM Journal on Financial Mathematics*, *2*, 255–286 (cit. on pp. 1, 5, 6, 9, 12–15, 28, 61, 62, 64, 65, 74, 128).

Guillaume, T. (2008). Making the best of best-of. *Review of Derivatives Research*, *11*(1), 1–39 (cit. on p. 45).

Gurrieri, S., Nakabayashi, M., & Wong, T. (2009). Calibration Methods of Hull-White Model. *Available at SSRN: https://ssrn.com/abstract=1514192 or http://dx.doi.org/10.2139/ssrn.1514192* (cit. on p. 28).

Hardy, M. (2003). *Investment guarantees modeling and risk management for equity-linked life insurance*. John Wiley & Sons, Hoboken, New Jersey. (Cit. on p. 11).

Heston, S. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, *6*(2), 327–343 (cit. on pp. 1–3, 6–8, 28, 29, 61, 75–82, 84, 86–91, 93, 94, 96, 103, 107, 129).

Hull, J., & White, A. (1990). Pricing Interest Rate Derivative Securities. *Review of Financial Studies*, *3*(4), 573–592 (cit. on pp. 2, 8, 10–12, 23, 28, 29, 61).

Hurd, T., & Zhou, Z. (2010). A Fourier Transform Method for Spread Option Pricing. *SIAM Journal on Financial Mathematics*, *1*(1), 142–157 (cit. on pp. 6, 46, 47, 49, 52, 54, 56, 59, 61, 68, 70, 71, 74, 106, 108, 112, 115, 128, 129).

Ignatieva, K., Song, A., & Ziveyi, J. (2016). Pricing and hedging of guaranteed minimum benefits under regime-switching and stochastic mortality. *Insurance: Mathematics and Economics*, *70*(100), 286–300 (cit. on pp. 10, 18, 21).

in 't Hout, K., Bierkens, J., van der Ploeg, A., & in 't Panhuis, J. (2007). A semi closed-form analytic pricing formula for call options in a hybrid Heston–Hull–White model. *In: Proceedings of the 58th European Study Group Mathematics with Industry. Utrecht. http://www.math.uu.nl/swi2007/proc2007.pdf*, 101–116 (cit. on p. 9).

Kammeyer, H., & Kienitz, J. (2012). The Heston-Hull-White Model Part I: Finance and Analytics. *Wilmott Magazine*, *2012*(57), 46–53 (cit. on pp. 7, 61).

Kammeyer, H., & Kienitz, J. (2009). An Implementation of the Hybrid-Heston-Hull-White Model. *Available at SSRN: https://ssrn.com/abstract=1399389 or http://dx.doi.org/10.2139/ssrn.1399389* (cit. on p. 9).

Khuzwayo, B., & Maré, E. (2014). Aspects of volatility targeting for South African equity investors. *South African Journal of Economic and Management Sciences*, *17*(5), 691–699 (cit. on pp. 96, 97, 103).

Klyueva, E. (2014). Pricing and hedging tools for spread option contracts. *Master's thesis, University of Toronto* (cit. on p. 45).

Lee, R., & Carter, L. (1992). Modeling and Forecasting U.S. Mortality. *Journal of the American Statistical Association*, *87*(419), 659–671 (cit. on pp. 10, 11).

Levendis, A., & Maré, E. (2022a). An Economic Scenario Generator for Embedded Derivatives in South Africa. *South African Actuarial Journal*, *22*(1), 79–118 (cit. on pp. 5, 7).

Levendis, A., & Maré, E. (2022b). Efficient Pricing of Spread Options with Stochastic Rates and Stochastic Volatility. *Journal of Risk and Financial Management*, *15*(11), 504 (cit. on pp. 6, 60).

Levendis, A., & Maré, E. (2023a). On the calibration of stochastic volatility models to estimate the real-world measure used in option pricing. *ORiON*, *39*(1), 65–91 (cit. on pp. 6, 75).

Levendis, A., & Maré, E. (2023b). Pricing Two-Asset Rainbow Options with the Fast Fourier Transform. *South African Statistical Journal*, *57*(1) (cit. on pp. 5, 45).

Levendis, A., & Maré, E. (2023c). Static hedging of vanilla and exotic options in south africa. *Accepted and to be presented at the Afric 2023: Actuarial, Finance, Risk and Insurance Congress, Victoria Falls* (cit. on pp. 6, 105).

Lin, T., Wang, C., & Tsai, C. (2015). Age-specific copula AR-GARCH mortality models. *Insurance: Mathematics and Economics*, *61*(100), 110–124 (cit. on p. 19).

Mandelbrot, B., & van Ness, J. (1968). Fractional brownian motions, fractional noises and applications. *SIAM Review*, *10*(4), 422–437 (cit. on p. 2).

Margrabe, W. (1978). The Value of an Option to Exchange One Asset for Another. *The Journal of Finance*, *33*(1), 177–186 (cit. on pp. 6, 48).

Maze, S. (2014). Efficient Implementation of the Heston-Hull & White Model. *Available at SSRN: https://ssrn.com/abstract=2378955 or http://dx.doi.org/10.2139/ssrn.2378955* (cit. on pp. 8, 9).

Merton, R. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, *4*(1), 141–183 (cit. on p. 1).

Ngugi, A., Maré, E., & Kufakunesu, R. (2015). Pricing variable annuity guarantees in South Africa under a variance-gamma model. *South African Actuarial Journal*, *15*(1), 131–170 (cit. on p. 11).

Olivieri, A., Thirurajah, S., & Ziveyi, J. (2022). Target Volatility Strategies For Group Self-Annuity Portfolios. *ASTIN Bulletin: The Journal of the IAA*, *52*(2), 591–617 (cit. on p. 77).

Oosterlee, C. (2007). The ING problem: A problem from the financial industry. *In: Proceedings of the 58th European Study Group Mathematics with Industry. Utrecht. http://www.math.uu.nl/swi2007/proc2007.pdf*, 91–92 (cit. on p. 9).

Ouwehand, P., & West, G. (2006). Pricing Rainbow Options. *Wilmott Magazine*, *5*, 74–80 (cit. on pp. 45, 46).

Patel, R. (2019). Approximating the Heston-Hull-White Model. *Master's thesis, University of Cape Town, http://hdl.handle.net/11427/30881* (cit. on pp. 9, 12–14, 16).

Poklewski-Koziell, W. (2012). Stochastic volatility models: Calibration, pricing and hedging. *Master's thesis, University of the Witwatersrand, http://hdl.handle.net/10539/11991* (cit. on pp. 76, 107).

Roberts, J. (2018). Fourier Pricing of Two-Asset Options: A Comparison of Methods. *Master's thesis, University of Cape Town, http://hdl.handle.net/11427/28126* (cit. on pp. 45–47, 53–57, 69).

Ruijter, M., & Oosterlee, C. (2012). Two-Dimensional Fourier Cosine Series Expansion Method for Pricing Financial Options. *SIAM Journal on Scientific Computing*, *34*(5), 642–671 (cit. on p. 56).

Stulz, R. (1982). Options on the Minimum or the Maximum of Two Risky Assets: Analysis and Applications. *Journal of Financial Economics*, *10*(2), 161–185 (cit. on pp. 45, 47, 54, 55, 57, 59, 128).

Syuhada, K., & Hakim, A. (2021). Stochastic modeling of mortality rates and Mortality-at-Risk forecast by taking conditional heteroscedasticity effect into account. *Heliyon*, *7*(10), e08083. doi: 10.1016/j.heliyon.2021.e08083 (cit. on pp. 16, 19, 26, 39).

Truter, G. (2012). The Valuation and Hedging of Default-Contingent Claims in Multiple Currencies. *Master's thesis, University of the Witwatersrand, http://hdl.handle.net/10539/11955* (cit. on p. 10).

van Dijk, M., de Graaf, C., & Oosterlee, C. (2018). Between $\mathbb{P}$ and $\mathbb{Q}$: The $\mathbb{PQ}$ Measure for Pricing in Asset Liability Management. *Journal of Risk and Financial Management*, *11*(4), 67 (cit. on pp. 29, 77).

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, *5*(2), 177–188 (cit. on p. 10).

Veilleux, P. (2016). On the impact of stochastic volatility, interest rates and mortality on the hedge efficiency of GLWB quarantees. *Master's thesis, Laval University, http://hdl.handle.net/20.500.11794/26668* (cit. on p. 11).

Venter, P., Levendis, A., & Maré, E. (2022). Collateralised option pricing in a South African context: A Univariate GARCH approach. *Cogent Economics & Finance*, *10*(1) (cit. on pp. 75, 97).

Wang, G. (2011). An Equity and Foreign Exchange Heston-Hull-White model for Variable Annuities. *Master's thesis, Delft University of Technology, http://resolver.tudelft.nl/uuid:b8b2276c-aeee-49ed-b9b7-0bc78b794d76* (cit. on p. 10).

Index

AR(1)-ARCH(1) model, 8, 16, 19, 20, 24–26, 34–36, 44

Bates SVJ model, 1, 3, 6, 7, 76, 77, 79, 81, 82, 84, 86, 89–91, 94–97, 103, 104, 106, 107, 129

characteristic function, 2, 9, 14, 15, 29, 46, 47, 49, 51, 52, 61, 63–65, 68, 106–108, 112, 115, 129
CIR++ model, 10, 17, 19, 24, 32, 34, 36, 44

EMM, 2, 6, 76, 77, 82–84, 86, 88, 90, 93, 103, 111, 129

FFT, 2, 3, 6, 9, 10, 46, 47, 49, 52–59, 61, 68–74, 106, 108, 112, 114, 115, 128, 129

GMDB, 3, 5, 7–9, 11, 16, 20–22, 27–32, 36–40, 44, 128
GMMB, 3, 5, 7–9, 11, 16, 20–22, 27–32, 36–44, 128

Heston model, 1, 3, 6–8, 28, 29, 61, 75–82, 84, 86–91, 93, 94, 96, 103, 107, 129
Heston-Hull-White model, 1, 5, 9–16, 20, 23, 28–32, 128
Heston-Hull-White-Mortality model, 8, 11, 20, 27, 28, 36, 39–41, 43, 44

rainbow option, 4, 5, 45–47, 56, 59
real-world measure, 2–4, 76, 77, 104, 106, 129
replicating portfolio, 106, 108, 109, 115, 117, 119, 120, 123–127, 129
risk-neutral measure, 2–4, 76, 77, 106, 129

spread option, 3, 4, 6, 61, 68–71, 73, 74, 121, 122, 126, 128, 129
static hedging, 4, 6, 105, 106, 108–110, 114–116, 118–122, 124, 126, 127, 129
stochastic interest rates, 3, 5–8, 11, 15, 16, 20, 44, 60, 61, 68–70, 73, 74, 128
stochastic mortality, 5, 8, 10, 11, 16, 20, 39, 44, 128
stochastic volatility, 2–5, 7, 8, 11, 19, 38, 44, 60, 61, 76, 77, 82, 83, 86, 88, 90, 96, 103, 120, 126, 128–130
SVJJ model, 4, 6, 106, 107, 110–117, 121, 126, 127, 129

three-factor stochastic volatility model, 3, 6, 47–52, 54, 57, 59, 61, 62, 71, 123, 128
two-asset Heston-Hull-White model, 3, 6, 61, 71, 128, 129
two-factor gBm model, 6, 46–52, 54, 57, 59, 128

volatility targeting, 4, 6, 77, 96, 97, 99, 103, 104, 129

worst-of-2 call option, 3, 5, 6, 46, 47, 49, 51–54, 56–59, 128

Colophon

This thesis was typeset with LaTeX 2$_\varepsilon$. It uses the *Clean Thesis* style developed by Ricardo Langner. The design of the *Clean Thesis* style is inspired by user guide documents from Apple Inc.

Download the *Clean Thesis* style at http://cleanthesis.der-ric.de/.