An Application of Deep Hedging in Pricing and Hedging Caplets on the Prime Lending Rate

Keyur Patel

A dissertation submitted to the Faculty of Commerce, University of Cape Town, in partial fulfilment of the requirements for the degree of Master of Philosophy.

August 30, 2022

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

I declare that this dissertation is my own, unaided work. It is being submitted for the Degree of Master of Philosophy in the University of the Cape Town. It has not been submitted before for any degree or examination in any other University.

August 30, 2022

Abstract

Derivatives in South Africa are traded via an exchange, such as the JSE's derivatives markets, or over-the-counter (OTC). This dissertation focuses on the pricing and hedging of caplets written on the South African prime lending rate. In a complete market, caplets can be continuously hedged with zero risk. However, in the particular case of caplets written on the prime lending rate, market completeness ceases to exist. This is because the prime lending rate is a benchmark for retail lending and is not tradeable, in general. Since parametric models may not be specified and calibrated for such incomplete markets, the aim of this dissertation is to consider the deep hedging approach of Buehler *et al.* (2019) for pricing and hedging such a derivative.

First, a model dependent approach is taken to set a benchmark level of performance. This approach is derived using techniques outlined in West (2008) which rely heavily on interest rate pairs being cointegrated to use the market standard Black (1976) model. Thereafter, the deep hedging approach is considered in which a neural network is set up and used to price and hedge the caplets.

The deep hedging approach performs at least as well as the model dependent approach. Furthermore, the deep hedging approach can also be used to recover a volatility skew which is in fact, needed as an input in the model dependent approach. The approach has certain downsides to it: a rich set of historical data is required and it is more time consuming to conduct than the model dependent approach. The deep hedging approach, in this specific implementation, also has a limitation that only one hedge instrument is used. When this limitation is also applied to the model dependent approach, the deep hedging approach performs better in all cases. Therefore, deep hedging proves to be a sufficient alternative to pricing and hedging caplets on the prime lending rate in an incomplete market setting.

Acknowledgements

Throughout the writing of this dissertation, I have received an immense amount of support.

I would first like to thank my supervisor Obeid Mahomed for his support and guidance throughout this journey. Your insight and feedback has been extremely valuable in pushing my thinking and maintaining a high standard of work throughout the process.

To my parents, thank you for your continuous belief and support in what has been a long journey. You have always been there for me and it goes without say that none of this would have been possible without you.

Lastly, to my friends and classmates, thank you for all the encouragement you have given me in writing this dissertation. Going through this process in a remote learning environment has been extremely challenging but we have all managed to adapt to it and be there for one another.

Contents

1. Introduction . . . 1
- 1.1 Background . . . 1
- 1.2 Aim and Objectives . . . 2
- 1.3 Plan of Development . . . 3

2. The Model Dependent Approach . . . 4
- 2.1 The Black 76 Model . . . 4
  - 2.1.1 Pricing . . . 4
  - 2.1.2 Hedging . . . 5
- 2.2 Caplets Written on the Prime Lending Rate . . . 8
  - 2.2.1 Market Data and Caplet Parameters . . . 8
  - 2.2.2 The West Approach . . . 9
  - 2.2.3 Cointegration . . . 11
  - 2.2.4 Break-Even Volatility . . . 13
  - 2.2.5 Pricing and Hedging . . . 18

3. Benchmarking the Neural Network . . . 21
- 3.1 Theoretical Background . . . 21
  - 3.1.1 Pricing and Hedging Derivatives in Incomplete Markets . . . 21
  - 3.1.2 Machine Learning, Pricing and Hedging . . . 25
- 3.2 Implementation . . . 27
  - 3.2.1 Setting up the Neural Network . . . 27
  - 3.2.2 Training the Neural Network . . . 28
  - 3.2.3 Benchmarking Results . . . 31

4. Deep Hedging Caplets on the Prime Lending Rate . . . 35
- 4.1 Training the Neural Network . . . 35
  - 4.1.1 Input Data . . . 35
  - 4.1.2 Training Algorithm . . . 36
- 4.2 Post Processing . . . 37
- 4.3 Results . . . 38
  - 4.3.1 Optimising the Neural Network . . . 38
  - 4.3.2 Deep Hedging . . . 39

5. Conclusion . . . 41
- 5.1 Future Recommendations . . . 42

Bibliography 43

A. The Vasicek Model 45

B. The Libor Market Model 46

C. Statically vs Dynamically Hedging the ZCB Exposure 48

D. Pricing and Hedging Floorlets 49

E. Quadratic Risk Minimisation 52

F. Adam Optimisation Algorithm 54

List of Figures

2.1 Actual Market Data . . . 13
2.2 LMM BEV Skews . . . 15
2.3 Prime BEV Skews . . . 16
2.4 Prime BEV Skews . . . 17
2.5 Prime BEV Skews . . . 18
2.6 Bank Account and FRA Holdings . . . 20

3.1 Simple Neural Network Schematic . . . 27
3.2 LSTM Cell (Olah, 2015) . . . 28
3.3 Comparison of Black (1976) and Neural Network Delta Hedge Ratios 32
3.4 Average Relative Error between Deep Hedging and Black (1976) Delta Hedge Ratios . . . 33
3.5 Hedge PnL Distribution with 50% and 95% Confidence Levels . . . 34

4.1 Prime BEV Skew Comparison . . . 40

C.1 Hedge PnL Distribution with Static ZCB Hedge . . . 48
C.2 Hedge PnL Distribution with Dynamic ZCB Hedge . . . 48

List of Tables

2.1 Hedge PnL Distribution Parameter Comparison 8
2.2 ADF Test Statistics 12
2.3 No. of times a 6x9 prime caplet has a non-zero payoff 17

3.1 Benchmarking - Batch Size Optimisation 30
3.2 Benchmarking - Number of Epochs Optimisation (Batch Size = 500) . 31

4.1 Actual Market Data - Batch Size Optimisation 38
4.2 Actual Market Data - Number of Epochs Optimisation (Batch Size = 20) 38
4.3 Prime Caplet Hedge PnL Results 39

Chapter 1

Introduction

Derivatives in South Africa are traded via an exchange, such as the JSE's derivatives markets, or over-the-counter (OTC). In particular, interest rate derivatives are predominantly traded OTC given that most are non-standardised products. These derivatives include instruments such as forward rate agreements (FRAs), interest rate swaps, caps, floors and swaptions. They may be used for arbitrage, speculative or hedging activity, or to alter one's exposure to an underlying interest or reference rate.

This dissertation is centred around pricing and hedging caplets written on the prime lending rate. Interest rate derivatives written on the prime lending rate, which is a benchmark rate for retail lending, are always bespoke in nature. A key set of financial instruments that almost standardly reference the prime lending rate are preference shares.

1.1 Background

Caplets are the building blocks of interest rate caps and can be considered as the option version of a long FRA - they provide protection against rising interest rates. The payoff at time t_N of a τ-tenor caplet, initiated at time t_0, written on a simple underlying floating rate R_x with accrual period $[t_{N-1}, t_N]$, where $t_0 \leq t_{N-1}$ and x is the reference rate, is:

$$C(t_N) = A(R_x(t_{N-1}, t_N) - K)^+ \tau, \tag{1.1}$$

where A is the nominal, τ is the year fraction between any successive t_i and t_{i+1} for $i \in \{1, 2, \ldots, N-1\}$, $R_x(t_{N-1}, t_N)$ is the underlying rate at time t_{N-1} and K is the strike rate. The market standard for pricing caplets is to use the Black (1976) model which yields closed-form solutions for pricing and hedging. Using this model gives rise to the common problem of determining the correct volatility to use in the pricing formula. This can be dealt with by either using market observed option prices

to imply volatilities or through computing a break-even volatility skew as shown by Dupire (2006). In a complete market, these instruments can be continuously hedged with zero risk. However, in the particular cases of caplets written on the prime lending rate, market completeness ceases to exist. This is because the prime lending rate is a benchmark for retail lending and is not tradeable, in general.

West (2008) proposes dealing with this incompleteness by transforming the prime caplet into an equivalent caplet written on the 3-month Johannesburg Interbank Average Rate (JIBAR) which is the underlying rate for vanilla interest rate derivatives in South Africa. JIBAR rates are money market rates that represent the average rates at which South African banks buy and sell money in the interbank market. This transformation rests upon the historically observed cointegration relationship between 3-month JIBAR and the prime lending rate. More generalised approaches towards incomplete markets where the underlying is not tradeable are investigated by Hulley and McWalter (2015) which include utility indifference pricing and quadratic risk minimisation. Buehler *et al.* (2019) recast these classical approaches of incomplete market pricing and hedging within a deep learning framework. These approaches however, have not been compared to the approach set out by West (2008) to price and hedge caplets written on the prime lending rate.

1.2 Aim and Objectives

Since parametric models may not be specified and calibrated for such incomplete markets without some form of mathematical manipulation; the aim of this dissertation is to consider the deep hedging approach of Buehler *et al.* (2019) for pricing and hedging caplets written on the prime lending rate as an alternative. The following objectives have been set towards achieving the aim of the dissertation:

1. Perform a cointegration analysis between the prime lending rate and 3-month JIBAR to use in the approach set out by West (2008).

2. Using the results from the cointegration analysis, implement the West (2008) model dependent approach to set a benchmark of pricing and hedging caplets written on the prime lending rate in an incomplete market.

3. Set up a neural network and implement the deep hedging approach of Buehler *et al.* (2019) to price and hedge caplets written on the prime lending rate.

It is found that in most cases, the deep hedging approach performs at least as well as the model dependent approach set out by West (2008). Hence, it provides a suitable, model independent alternative for traders to price and hedge caplets written

on the prime lending rate. Some of the practical limitations of this approach are that it is data-driven - therefore, a rich set of historical data is required in its implementation. Furthermore, it may be more time consuming to use because the neural network needs to be retrained if a caplet with a different tenor, moneyness level or maturity needs to be priced.

1.3 Plan of Development

First, the model dependent approach set out by West (2008) is presented in chapter 2 along with the results from its implementation. This sets up a benchmark to compare the results from the deep hedging approach to. Next, a benchmarking analysis, conducted to ensure the correct functionality of the neural network for the deep hedging approach, is presented in chapter 3. The deep hedging approach towards pricing and hedging caplets written on the prime lending rate is then presented in chapter 4. The report concludes in chapter 5 with an overview of the entire project along with recommendations to take the research further.

Chapter 2

The Model Dependent Approach

The purpose of this chapter is to present the theoretical background relevant to the West (2008) approach along with the results from the implementation thereof. Section 2.1 provides a theoretical background to pricing and hedging caplets in the context of the Black (1976) model. This is followed by section 2.2 which details the West (2008) approach and the results of its implementation.

2.1 The Black 76 Model

2.1.1 Pricing

As mentioned previously, a caplet, whose payoff is shown by equation (1.1), is a derivative that protects holders against rising interest rates. Caplets are standardly priced using the Black (1976) model which assumes that interest rates possess a lognormal distribution under the respective T-forward measure. Considering the same caplet from section 1.1, the Black (1976) price at time t_0 is given by:

$$C(t_0) = AZ(t_0, t_N)\Big(R_x(t_0; t_{N-1}, t_N)\Phi(d_+) - K\Phi(d_-)\Big)\tau, \tag{2.1}$$

with

$$d_\pm = \frac{\ln\left(\frac{R_x(t_0;t_{N-1},t_N)}{K}\right) \pm \frac{1}{2}\sigma^2(t_{N-1} - t_0)}{\sigma\sqrt{t_{N-1} - t_0}}, \tag{2.2}$$

where $t_0 \leq t_{N-1}$, σ is the model volatility, $Z(t_0, t_N)$ is the risk-free discount factor applicable for $[t_0, t_N]$, $R_x(t_0; t_{N-1}, t_N)$ is the fair $[t_{N-1}, t_N]$ simple forward rate at t_0 and $\Phi(x)$ is the standard normal cumulative distribution function evaluated at x. For ease of reading, the fair $[t_{N-1}, t_N]$ forward rate $R_x(t_0; t_{N-1}, t_N)$ will be denoted as $R_{x,t_{N-1}}(t_0)$. Note that the time increments in these equations e.g., $t_{N-1} - t_0$, are in year fractions.

A market investigation by Gupta and Subrahmanyam (2005) on caps and floors illustrates that the lognormal distribution assumption matches stylised features of

interest rate markets. In particular, it results in a 'smaller skew in pricing errors across strike rates' compared to models which assume other distributions. Lognormal models do not allow for negative interest rates; however, this is not a limitation for an emerging market like South Africa, where interest rates have always been positive and the future likelihood of non-positive rates is negligible.

2.1.2 Hedging

Under continuous trading in a complete market, ignoring market frictions, Musiela and Rutkowski (2005) have shown that caplets can be replicated using appropriate positions in FRAs and a zero-coupon bond (ZCB) under the Black (1976) model. These FRAs reference the same underlying rate as the caplet. The time s value of a long FRA written on a simple underlying rate x, with strike K, nominal A, tenor τ, reset time t_{N-1} and settlement time t_N is given by:

$$V_{FRA}(s) = AZ(s, t_N)\Big(R_{x,t_{N-1}}(s) - K\Big)\tau. \tag{2.3}$$

where $s < t_{N-1}$. If caplets can be replicated in a complete market under continuous trading, this means that there exists a hedge portfolio with value process $(V_t)_{t\geq 0}$ such that $V_{t_N} = C(t_N)$ in all states of the world. Hence, caplets can be hedged with zero risk. Note that trading in the hedge portfolio only occurs until the expiry, t_{N-1}, of the caplet given that the final payoff, $C(t_N)$, becomes known at that time.

In a real-world context, however, continuous trading is not possible. Therefore, the hedge portfolio can only be rebalanced at discrete time-steps prior to t_{N-1}. At expiry, the hedge profit/loss (PnL) is calculated as the difference between the caplet and hedge portfolio values. For all delta hedging implementations in this dissertation, the hedge portfolio is rebalanced daily given that daily historical market data is available. Furthermore, it is assumed that we are sellers of the caplet and hence take a long position in the hedge portfolio.

The following sensitivities at any time $s \in [t_0, t_{N-1})$ are defined and used to compute the required positions in the hedge portfolio for a caplet:

$$\Delta_{FRA,s} := \frac{\partial V_{FRA}(s)}{\partial R_{x,t_{N-1}}(s)},$$
$$\Delta_{ZCB,s} := \frac{\partial C(s)}{\partial Z(s, t_N)},$$
$$\Delta_{C,s} := \frac{\partial C(s)}{\partial R_{x,t_{N-1}}(s)}.$$

These sensitivities are calculated using the equations below:

$$\Delta_{FRA,s} = A\tau Z(s, t_N), \tag{2.4}$$

$$\Delta_{ZCB,s} = A\big(R_{x,t_{N-1}}(s)\Phi(d_+^1) - K\Phi(d_-^1)\big)\tau, \tag{2.5}$$

$$\Delta_{C,s} = AZ(s,t_N)\left(\Phi(d_+^1) + R_{x,t_{N-1}}(s)\phi(d_+^1)\frac{\partial d_+^1}{\partial R_{x,t_{N-1}}(s)} - K\phi(d_-^1)\frac{\partial d_-^1}{\partial R_{x,t_{N-1}}(s)}\right)\tau, \tag{2.6}$$

where $\phi(x)$ is the standard normal probability density function evaluated at x and

$$\frac{\partial d_+^1}{\partial R_{x,t_{N-1}}(s)} = \frac{1}{R_{x,t_{N-1}}(s)\sigma\sqrt{t_{N-1}-s}} = \frac{\partial d_-^1}{\partial R_{x,t_{N-1}}(s)}.$$

For simplicity, the FRAs used in the hedge portfolio are chosen to match the parameters of the caplet i.e., nominal A, tenor τ, reset time t_{N-1} and settlement time t_N. Furthermore, the FRAs used are fair value FRAs on the day the caplet is sold i.e., with fixed rate $K = R_{x,t_{N-1}}(t_0)$ and hence, a market value of zero on the day the caplet is sold. Given pricing equation (2.1), there are two risk factors that emerge: the forward rate $R_{x,t_{N-1}}(t_0)$ and discount factor $Z(t_0,t_N)$ (or zero-coupon bond). As per Musiela and Rutkowski (2005), delta hedging is done by taking a static hedge position in the zero-coupon bond, Δ_{ZCB,t_0}, on the day the caplet is sold, and dynamically hedging the forward rate exposure using FRAs. Note that we also introduce a bank account holdings process $(b_t)_{t\geq t_0}$ i.e., an amount deposited or borrowed, which grows at the risk-free overnight rate. This bank account is used to fund the daily changes in FRA positions or deposit any excess cash arising from these changes.

The delta hedge ratio, δ_s, for caplets i.e., the position to take in FRAs at time $s \in [t_0, t_{N-1})$, is a function of the sensitivity of the caplet value to changes in the underlying forward rate and the sensitivity of the value of a FRA to changes in the underlying forward rate:

$$\delta_s := \frac{\Delta_{C,s}}{\Delta_{FRA,s}}. \tag{2.7}$$

The hedge portfolio is constructed as follows:

1. At t_0, receive $C(t_0)$ from the sale of the caplet, purchase Δ_{ZCB,t_0} many ZCBs maturing at t_N and purchase δ_{t_0} many FRAs with the properties relevant to equation (2.3) at zero cost given that they are fair FRAs i.e., with the strike being the fair forward rate on the day the caplet is sold.

2. For all $s \in [t_0+1, t_{N-1}-1]$, adjust the FRA position from the previous day, δ_{s-1}, to ensure the hedge portfolio has δ_s many FRAs while adjusting the holding in the bank account b_s to fund this if necessary or deposit any excess cash that may arise.

3. At time t_{N-1}, the hedge PnL is calculated as:

$$PnL = b_{t_{N-1}-1}C^{ON}(t_{N-1}-1, t_{N-1}) + \delta_{t_{N-1}-1}V_{FRA}(t_{N-1}) - Z(t_{N-1}, t_N)C(t_N) + \Delta_{ZCB,t_0}Z(t_{N-1}, t_N),$$

where $b_{t_{N-1}-1}$ is the amount in the bank account at $t_{N-1}-1$ and $C^{ON}(t_{N-1}-1, t_{N-1})$ is a risk-free overnight capitalisation factor over the period $[t_{N-1}-1, t_{N-1}]$.

This delta hedging routine differs from a conventional one where one would expect to dynamically hedge the exposures to both $R_{x,t_{N-1}}(s)$ and $Z(s, t_N)$ instead of just $R_{x,t_{N-1}}(s)$ where $s \in [t_0, t_{N-1})$. In such a case, step 1 in the routine above would remain the same but steps 2 and 3 would be modified as follows:

2*. For all $s \in [t_0 + 1, t_{N-1} - 1]$, adjust the FRA and ZCB positions from the previous day, δ_{s-1} and $\Delta_{ZCB,s-1}$ respectively, to ensure the hedge portfolio has δ_s and $\Delta_{ZCB,s}$ many FRAs and ZCBs respectively while adjusting the holding in the bank account b_s to fund this if necessary or deposit any excess cash that may arise.

3*. At time t_{N-1}, the hedge portfolio PnL is calculated as:

$$PnL = b_{t_{N-1}-1}C^{ON}(t_{N-1}-1, t_{N-1}) + \delta_{t_{N-1}-1}V_{FRA}(t_{N-1}) - Z(t_{N-1}, t_N)C(t_N) + \Delta_{ZCB,t_{N-1}-1}Z(t_{N-1}, t_N).$$

Comparing Statically and Dynamically Hedging the Zero-Coupon Bond Exposure

Prior to committing to using the delta hedging methodology set out in Musiela and Rutkowski (2005), a comparison between statically and dynamically hedging the ZCB exposure is performed. This is done in a single-curve framework using simulated data from a single-factor Libor Market Model (Brace *et al.*, 1997) with volatility parameter $\sigma_{LMM} = 0.2$ (which becomes the volatility to use in the Black (1976) model too). This model forces discrete forward rate dynamics into lognormal form which then makes the approximations made in the Black (1976) model exact. To initialise an arbitrary yield curve, which is needed as input to the Libor Market Model, the Vasicek (1977) short-rate model is used with the following parameters:

$$r_0 = 0.07,\ a = 0.15,\ b = 0.09,\ \sigma_{vas} = 0.02$$

where r_0 is the initial short-rate, and a, b and σ_{vas} represent the rate of mean reversion, the mean reversion level and the Vasicek (1977) model volatility respectively.

Note that the parameters chosen for the Libor Market Model and Vasicek (1977) model are completely arbitrary. These models are only used to simulate interest rate data so that the pricing and hedging methodologies set out in sections 2.1.1 and 2.1.2 can be tested and compared. Further detail of how the Vasicek (1977) and Libor Market models are simulated can be found in Appendix A and B respectively.

A set of $100,000$ forward rate time series is generated with time increments of 1 day over a period of 9 months. Thereafter, 6x9 caplets are priced and hedged using the methodologies set out in sections 2.1.1 and 2.1.2. That is, 3-month tenor caplets initiated at t_0, with expiry 6 months later at t_{N-1}, and payoff 9 months after t_0, at t_N. An arbitrary moneyness level of 110% and nominal of $A = 1000$ is used for the caplets.

This analysis shows that the static ZCB hedge methodology proposed by Musiela and Rutkowski (2005) performs marginally better than one where a dynamic hedge is used. Therefore, the hedging methodology set out by Musiela and Rutkowski (2005) is used in this dissertation. The hedge PnL distribution parameters for the two different methods are shown in the table below:

Tab. 2.1: Hedge PnL Distribution Parameter Comparison

Hedging Method	Mean Hedge PnL	Hedge PnL Standard Deviation
Static ZCB Hedge	0.011	0.142
Dynamic ZCB Hedge	0.013	0.166

Histograms of the hedge PnL distributions from the two methods are shown in Appendix C for a pictorial representation.

2.2 Caplets Written on the Prime Lending Rate

2.2.1 Market Data and Caplet Parameters

To perform the necessary analyses, historical interest rate data spanning from 2 January 2004 to 28 June 2019 is used. As with the simulated data analysis above, 6x9 caplets with a nominal of $A = 1000$ are priced and hedged i.e., $\tau = 0.25$-tenor (or 3-month tenor) caplets which are initiated at a time t_0 with expiry 6 months later at t_{N-1} and settlement 9 months after t_0, at t_N. Due to the limited amount of historical data, overlapping time windows of data with increments of one day are created to maximise the number of sample paths available for pricing and hedging. Thus, there are 3740 sample paths of interest rates over which to price and hedge the 6x9 caplets. While using overlapping windows is inferior to using independent, non-overlapping windows, it is a compromise that is made to maximise the historical

data set which becomes extremely important in the deep hedging approach given that it is a data-driven approach.

As per the current standard approach in South Africa, the implied 3-month JIBAR nominal swap zero curve is the single, risk-free curve used for discounting in all of the analyses performed using real-world data in this dissertation. Furthermore, the SAFEX overnight (ON) rate, which is the best proxy for an overnight risk-free rate, is used to grow any overnight lent or borrowed bank account positions during the daily rebalancing of the hedge portfolio.

Normalisation of Market Data

Prior to the pricing and hedging analyses, the sample paths are normalised - this ensures that the volatility of the rates is captured and that the analyses can be done in terms of moneyness levels as opposed to absolute strikes. All caplet analyses are done for moneyness levels ranging from 80% to 120%, in increments of 10%. The normalisation is done using ratios between rates at each successive time-step to prevent negative rates from arising. The starting prime lending rate, to which the sample paths of ratios are applied, is 11.5% - this is the first value of the prime lending rate in the historical time series. To explain this process by way of example - if there are sample paths of prices P_i^n where $i \in \{1, 2, \ldots, 10\}$ denotes the time-step number and $n \in \{1, 2, \ldots, 1000\}$ denotes the sample path number, the normalisation takes place as follows:

1. The ratios, y_i^n, are determined using $y_i^n = \frac{P_i^n}{P_{i-1}^n}$ where $i \in \{2, 3, \ldots, 10\}$ for all $n \in \{1, 2, \ldots, 1000\}$.

2. If we choose the starting price to be P_1^1, then the normalised price paths are given by $P_i^n = P_1^1 \prod_{j=2}^{i} y_j^n$ for all $n \in \{1, 2, \ldots, 1000\}$ where $i \in \{2, 3, \ldots, 10\}$.

2.2.2 The West Approach

Following the same notation as in equation (1.1), the payoff at time t_N for a τ-tenor caplet written on the prime lending rate would be:

$$C_P(t_N) = A(R_P(t_{N-1}, t_N) - K)^+ \tau, \tag{2.8}$$

where R_P is a simple τ-tenor rate at t_{N-1}. For the caplets priced in this dissertation, the market prime lending rate at t_{N-1} is assumed to be a τ-tenor rate applicable over $[t_{N-1}, t_N]$. According to equation (2.1), under the Black (1976) model we would need a forward prime rate, $R_{P,t_{N-1}}(t_0)$, to compute the caplet price. However, given that the prime lending rate is not directly tradeable and there are no

market traded prime instruments such as FRAs and swaps, a prime rate curve does not exist. This in turn implies that a forward prime curve does not exist, hence we have an incomplete market.

To deal with this, West (2008) proposes transforming the caplet on prime to an equivalent caplet on 3-month JIBAR given that this is the underlying rate for vanilla interest rate derivatives in South Africa. This is done using a cointegration relationship i.e., linear relationship, between the prime and 3-month JIBAR rates of the following form:

$$R_P(t_{N-1}, t_N) = \alpha_1 + \beta_1 \times R_J(t_{N-1}, t_N), \tag{2.9}$$

where $R_J(t_{N-1}, t_N)$ is the 3-month JIBAR rate at t_{N-1} applicable over $[t_{N-1}, t_N]$ (3 months). West (2008) uses the cointegration relationship to tweak the underlying rate definition in the derivative payoff i.e., equation (2.8), while preserving the structure of the interest rate model that models the risk-free system for discounting (the 3-month JIBAR swap zero curve). This allows for the replacement of the prime rate with a 3-month JIBAR rate in the payoff which in turn results in the replacement of prime forward rates with 3-month JIBAR forward rates that exist - hence, Black (1976) caplet prices can be calculated. Substituting equation (2.9) into the prime caplet payoff, equation (2.8), we get:

$$\begin{aligned} C_P(t_N) &= A(\alpha_1 + \beta_1 \times R_J(t_{N-1}, t_N) - K)^+ \tau \\ \therefore C_P(t_N) &= A\beta_1 \left(R_J(t_{N-1}, t_N) - \frac{K - \alpha_1}{\beta_1} \right)^+ \tau, \end{aligned} \tag{2.10}$$

which represents the tweaked payoff, or 3-month JIBAR equivalent caplet. This transformation is hinged upon the integrity of the cointegration relationship between prime and 3-month JIBAR rates. While it is unrealistic to expect such relationships to exactly represent actual rate movements, Alexander (2008) notes that 'no financial systems have higher cointegration than term structures' lending merit to this approach of pricing prime caplets. Thus, using the Black (1976) pricing equation (2.1) and equation (2.10), the t_0 price of a prime caplet is given by:

$$C_P(t_0) = A\beta_1 Z(t_0, t_N) \left(R_{J,t_{N-1}}(t_0) \Phi(d_+^1) - \frac{K - \alpha_1}{\beta_1} \Phi(d_-^1) \right) \tau, \tag{2.11}$$

where

$$d_\pm^1 = \frac{\ln\left(\frac{\beta_1 R_{J,t_{N-1}}(t_0)}{K - \alpha_1} \right) \pm \frac{1}{2}\sigma^2 (t_{N-1} - t_0)}{\sigma \sqrt{t_{N-1} - t_0}}, \tag{2.12}$$

where $R_{J,t_{N-1}}(t_0)$ is the fair forward 3-month JIBAR rate at t_0 applicable over $[t_{N-1}, t_N]$.

2.2.3 Cointegration

As described in Alexander (2008): Roughly, cointegration is a measure of the long-run relationship between two or more variables. To formally define cointegration, we define stationary processes and discuss the concept of order of integration:

Stationary Processes

A discrete time stochastic process $(X_t)_{t\in\{1,2,\ldots,T\}}$ is stationary if:

1. $\mathbb{E}[X_t]$ is a finite constant for all t.
2. $\mathbb{V}[X_t]$ is a finite constant where $\mathbb{V}[\cdot]$ is the unconditional variance operator.
3. The joint distribution of (X_t, X_s) depends only on $t - s$.

Statistical hypothesis tests called unit root tests are used to test for stationarity (Dickey and Fuller, 1981).

Order of Integration

A time series process is said to be integrated of order n, and denoted $I(n)$, if it is not stationary and n is the smallest number of times the process must be differenced in order to achieve a stationary process. Thus, a stationary process is commonly referred to as an $I(0)$ process.

Now more formally: Two integrated processes $(X_t)_{t\in\{1,2,\ldots,T\}}$ and $(Y_t)_{t\in\{1,2,\ldots,T\}}$ are cointegrated if there exists a process $(Z_t)_{t\in\{1,2,\ldots,T\}}$ that is a linear combination of these variables and is stationary i.e.,:

$$Z_t = X_t - \beta Y_t - \alpha,$$

where β is a cointegrating factor and α is a constant. This can be extended to analyses involving more than two processes in which case β will be a vector.

The two most common cointegration methodologies are those proposed by Engle and Granger (1987) and Johansen and Juselius (1990). The Engel-Granger methodology is used in this dissertation due to its simplicity and given that only two variables are being analysed. Furthermore, the criterion of minimum variance implied by ordinary least squares (OLS) regression, which is used in this methodology, is advantageous in financial applications. As outlined by West (2008), the Engle-Granger methodology applied to two time series, $(X_t)_{t\in\{1,2,\ldots,T\}}$ and $(Y_t)_{t\in\{1,2,\ldots,T\}}$, works as follows:

1. Check that $(X_t)_{t\in\{1,2,\ldots,T\}}$ and $(Y_t)_{t\in\{1,2,\ldots,T\}}$ are not $I(0)$.
2. Check that $(X_t)_{t\in\{1,2,\ldots,T\}}$ and $(Y_t)_{t\in\{1,2,\ldots,T\}}$ are $I(1)$.
3. Estimate α and β by performing an OLS regression $Y_t = \alpha + \beta X_t + \epsilon_t$.
4. Check that the residuals, $(\epsilon_t)_{t\in\{1,2,\ldots,T\}}$, are $I(0)$.

To test for the order of integration, and hence stationarity, the augmented Dickey-Fuller (ADF) unit root test is used (Dickey and Fuller, 1981). This is done by importing the required time series of data into Matlab and using the **adftest** function. These tests are done at both 5% and 1% significance levels. Unit root tests are left-tailed tests with the following null and alternate hypotheses:

$$H_0 : (X_t)_{t\in\{1,2,\ldots,T\}} \sim I(1)$$
$$H_1 : (X_t)_{t\in\{1,2,\ldots,T\}} \sim I(0)$$

Hence, if the test statistic lies left of the critical values based on the relevant significance level, one can reject the null hypothesis and conclude that the time series is stationary. When testing if a time series is $I(1)$, if the null hypothesis cannot be rejected, this does not imply that it is true. It is necessary to then perform a unit root test on the first difference of the time series, $(\Delta X_t)_{t\in\{1,2,\ldots,T\}}$. If the null hypothesis can be rejected on this test, one can then conclude that the time series is $I(1)$.

Implementation

The table below shows the test statistics from the ADF tests carried out on the time series of 3-month JIBAR rates and prime rates:

Tab. 2.2: ADF Test Statistics

# of differences taken	3-month JIBAR	Prime	Conclusion
0	-0.4655	-0.5563	The time series are not $I(0)$
1	-56.62	-62.18	The time series are $I(1)$

The Engle and Granger (1987) critical values at the 1% and 5% significance levels are -2.5684 and -1.9416 respectively. The test statistics in the first row do not lie in the critical region hence, the time series are not $I(0)$. Once the first differences are taken, the test statistics drop significantly into the critical region, hence we can conclude that the time series are $I(1)$.

An OLS regression carried out between the prime and 3-month JIBAR rates gives a relationship as shown in equation (2.9) with the following coefficients:

$$\alpha_1 = 0.036088532 \qquad \beta_1 = 0.959163857$$

It can be seen that the relationship between prime and 3-month JIBAR is almost one-to-one with an addition of approximately 3.5% to 3-month JIBAR. Finally, the time series of residuals from the regression is tested for stationarity using the ADF test. The ADF test statistic for stationarity of the residuals is -4.9477. Therefore, we conclude that the residuals are $I(0)$ and that the cointegration relationship holds at both 1% and 5% significance levels. For illustrative purposes, Figure 2.1 below shows the actual and cointegrated time series of interest rates respectively based on the above results:

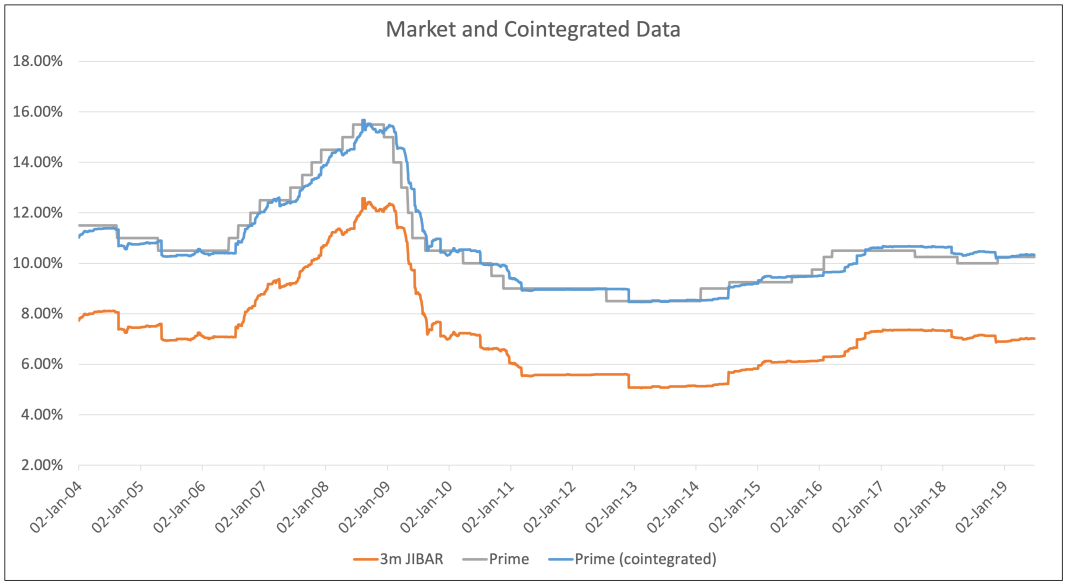

Fig. 2.1: Actual Market Data

From this graph, it can be seen that the cointegrated prime rate time series matches the actual prime rate time series closely. The cointegrated prime rate time series does not explicitly reflect the actual jumps in interest rates that we see in actual financial markets. This is expected given that 3-month JIBAR is being used to predict prime and interest rate movements typically get priced into 3-month JIBAR in advance, resulting in a smoother path.

2.2.4 Break-Even Volatility

Theoretical Background

A key input in the Black (1976) model, seen in equation (2.12), is the volatility parameter, σ. This can be computed using market observed caplet prices to imply volatilities across moneyness levels using equation 2.11. However, this approach

becomes problematic in illiquid caplet markets such as South Africa. Dupire (2006) introduced the data-driven break-even volatility (BEV) concept as another way of computing a volatility skew, or even the fair volatility for a single option, for pricing. The BEV of a caplet under the Black (1976) model for pricing is defined as the volatility, σ, which on average, results in zero PnL when the caplet is delta hedged within the context of the model. Using historical interest rate data, a BEV skew can be computed and used as input to the Black (1976) model for pricing in an illiquid market.

Using historical price data allows one freedom to choose how to segment the time series of data - in particular, to use overlapping or non-overlapping time windows. As highlighted in section 2.2.1, overlapping windows are used to maximise the data set. This also leads to smoother results due to the larger data set.

Given a series of time windows containing historical market data, there are two ways of arriving at one final volatility skew:

1. Compute the volatility skews for each time window and simply take an average of the skews.

2. For each moneyness level, calculate the BEV as the volatility which on average, gives a hedge PnL of zero across all time windows.

Dupire (2006) notes that the second method is more computationally efficient and may yield smoother skew results.

Implementation

Due to the computational efficiency and potential for smoother results, the second method is used. This method entails solving for a single volatility which when used in the Black (1976) model for pricing and hedging, results in an average hedge PnL of zero across the 3740 sample paths, for each moneyness level. The pricing and delta hedging routines used are those described in sections 2.1.1 and 2.1.2, tweaked for prime caplets as per section 2.2.5 below.

To determine a confidence interval around the BEV at each moneyness level, the 10^{th} and 90^{th} percentile BEVs are computed. To do this; at each moneyness level, the BEV is computed and then used to price and hedge the prime caplets. Thereafter, the 10^{th} and 90^{th} percentiles of the hedge PnL distribution are computed. For the two sample paths which correspond to these 10^{th} and 90^{th} percentile hedge PnLs, the Black (1976) volatilities that force the hedge PnLs to zero for these respective paths are solved for and taken to be the 10^{th} and 90^{th} percentile BEVs.

Prior to performing the BEV analysis using the real-world data, a BEV analysis is conducted using the simulated market data from section 2.1.2. This is done

for both 6x9 caplets and 6x9 floorlets with parameters as per section 2.1.2. The correct functionality of the code is implied if the Libor Market Model volatility of $\sigma_{LMM} = 0.2$ can be recovered for both caplets and floorlets across all moneyness levels - this constant volatility assumption is a property of the Black (1976) model. Furthermore, by put-call parity, the volatility skews for caplets and floorlets are expected to be similar. The relevant Black (1976) equations for pricing and hedging floorlets (both vanilla and under the West (2008) approach for prime floorlets) are shown in Appendix D for completeness. Figure 2.2 below shows the BEV skews that are recovered for 6x9 caplets and floorlets using the simulated Libor Market Model data:

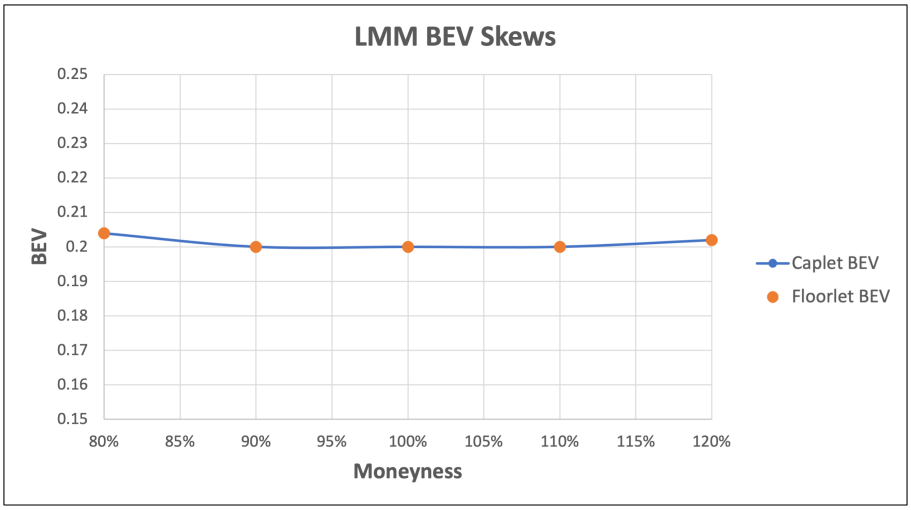

Fig. 2.2: LMM BEV Skews

It can be seen that the BEV analysis is able to recover flat BEV skews at a level of approximately 0.2 indicating the correct functionality of the code. Using the real-world interest rate data, vanilla 6x9 caplets and floorlets with parameters as per section 2.2.1, but written on the 3-month JIBAR rate are first considered in the BEV analysis. This is an intermediary step to test the BEV concept on vanilla products before prime caplets. Pricing and hedging these instruments is almost identical to those written on prime - the pricing and hedging equations given in sections 2.2.2 and 2.2.5, and Appendix D are used and tweaked by fixing $\alpha_1 = 0$ and $\beta_1 = 1$. Furthermore, in the final payoff equation of the caplets and floorlets, the t_{N-1} prime rate, $R_P(t_{N-1}, t_N)$, is replaced with the t_{N-1} 3-month JIBAR rate ,$R_J(t_{N-1}, t_N)$. The BEV skews for caplets and floorlets written on the 3-month JIBAR rate, along with the 10^{th} and 90^{th} percentile caplet BEV skews, are shown in Figure 2.3 overleaf:

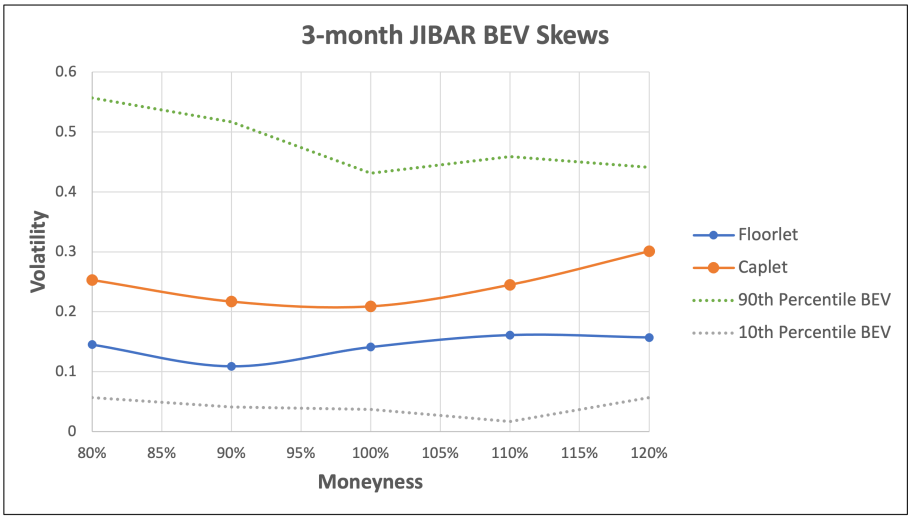

Fig. 2.3: Prime BEV Skews

It can be seen that when using the real-world data, there are some differences between the caplet and floorlet BEVs. Possible reasons for these differences are:

1. The small sample set of historical data - a smaller sample set means that fewer scenarios of interest rate fluctuations are accounted for. This can cause a significant amount of variation and skew in the results between caplets and floorlets. The 10^{th} and 90^{th} percentile caplet BEVs, which enclose the caplet and floorlet skews, indicate a large amount of variation and uncertainty in the BEVs.

2. The use of the SAFEX ON rate to grow bank account positions. Within a typical single-curve, lognormal Libor Market Model framework under the spot measure, the overnight rate is computed using interpolation on the single curve of interest rates (in this case, 3-month JIBAR). However, when using real-world data, the SAFEX ON rate needs to be used given that it is the best proxy for a risk-free overnight rate in South Africa. Thus, this adds another risk factor to the model.

3. The lognormal distribution assumption under the Black (1976) model does not fit well to the real-world historical interest rate data. Hence, a situation arises where the correct historical data is used, but the Black (1976) may not be the best model to use in a South African context. If the data were lognormal, one would expect to see results similar to those shown in Figure 2.2.

Figure 2.4 overleaf shows the prime caplet and floorlet BEV skews that are computed along with the 10^{th} and 90^{th} percentile caplet skews:

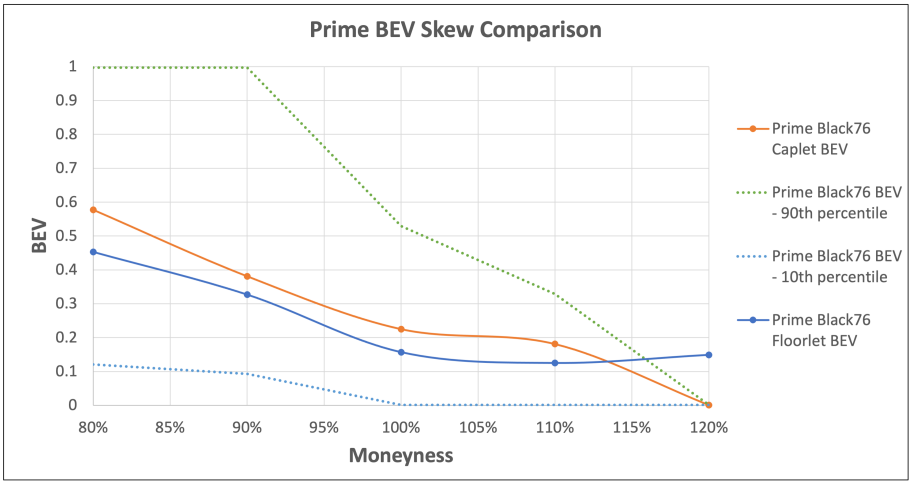

Fig. 2.4: Prime BEV Skews

The BEVs from the caplet skew (orange skew) are then used in the Black (1976) model for pricing and hedging at moneyness levels ranging from 80% to 120% in increments of 10%. Once again, there is some variation between the caplet and floorlet BEVs which can be attributed to the reasons given above. There is a fair amount of variability around the BEVs for in-the-money caplets i.e., at moneyness levels less than 100%. This variability generally reduces as the moneyness level increases. Notably, the BEVs become extremely small for out-the-money caplets i.e., moneyness levels from 110% onwards. This may be due to out-the-money caplets (moneyness levels greater than 100%) rarely having non-zero payoffs at expiry t_{N-1} due to the market prime rate being less than the caplet strike rate at expiry. The table below shows how many instances, out of the 3740 sample paths, end up with the caplet having a non-zero payoff at expiry at different moneyness levels:

Tab. 2.3: No. of times a 6x9 prime caplet has a non-zero payoff

Moneyness Level	80%	90%	100%	110%	120%
No. of Instances	3653	3518	1255	169	0

A possible explanation for this is the fact that at South African central bank meetings, it is very rare to have interest rate hikes greater than $0.25\% - 0.50\%$. Thus, given that these meetings typically take place every two months, the probability of the prime rate increasing enough to result in a non-zero payoff on a caplet with a moneyness level greater or equal to 110% is extremely low. The histogram overleaf shows the distribution of the relative changes in the prime lending rate over each 6-month interval.

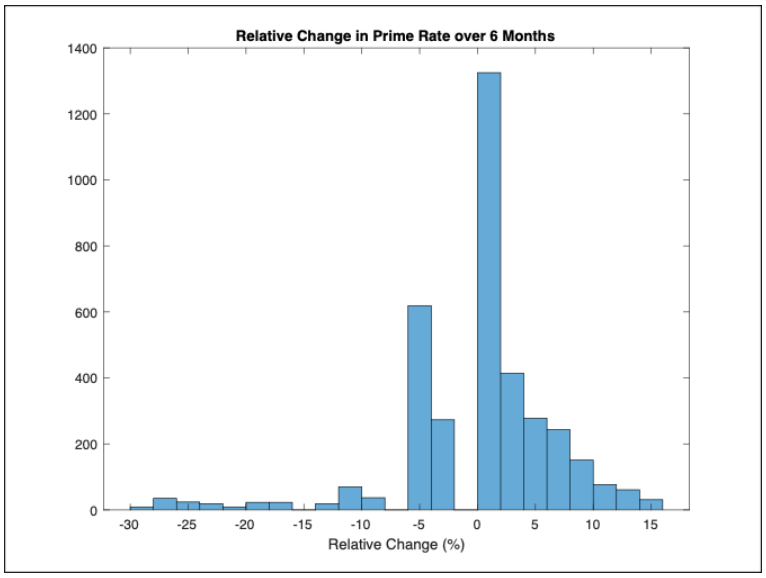

Fig. 2.5: Prime BEV Skews

The relative change can be interpreted as the move in moneyness level from 100% over the 6-month interval between when the caplet is sold and when it expires. As an example, a relative change of 10% would imply that any caplet bought with a moneyness level greater or equal to 110% at t_0 would have a payoff of zero at expiry. The distribution of relative changes is skewed towards increases in the prime lending rate over each 6-month interval. This is a possible reason for the higher variability in BEVs for moneyness levels less than 110% given that they almost always end up in-the-money. It can also be seen that the maximum relative change is 15% - this further consolidates that caplets bought with moneyness levels of 120% never have a non-zero payoff at expiry. Thus, there is no sensitivity to the volatility resulting in an extremely small caplet BEV being acceptable. Using higher volatilities, up to 0.15, is tested at the 120% moneyness level for prime caplets and it is found that the hedge PnL distribution parameters remain negligible (of order 10^{-5}). Using an extremely small volatility implies that a prime caplet at a moneyness level of 120% can be sold for zero premium given that the payoff is always zero using this dataset.

2.2.5 Pricing and Hedging

To price the prime caplets under the West (2008) approach across each of the 3740 sample paths, equation (2.11) is used. This makes use of the fair forward 3-month JIBAR rate, $R_{J,t_{N-1}}(t_0)$, which arises due to the transformation of the caplet payoff as highlighted in section 2.2.2. As per section 2.1.2, the delta hedging methodology

set out by Musiela and Rutkowski (2005) is used. The FRAs used for hedging reference 3-month JIBAR and have the same parameters as the caplet i.e., nominal A, tenor τ, reset time t_{N-1} and settlement time t_N. For caplets on the prime lending rate, the hedging methodology follows the same steps but the relevant sensitivities as per equations (2.5) and (2.6) are tweaked slightly due to the transformed payoff equation being different to a standard caplet. The tweaked sensitivities, where $s \in [t_0, t_{N-1})$, are shown below:

$$\Delta_{ZCB,s} := \frac{\partial C_P(s)}{\partial Z(s,t_N)} = A\beta_1 \left(R_{J,t_{N-1}}(s)\Phi(d_+^1) - \frac{K-\alpha_1}{\beta_1}\Phi(d_-^1) \right)\tau, \tag{2.13}$$

$$\begin{aligned}\Delta_{C_P,s} := \frac{\partial C_P(s)}{\partial R_{J,t_{N-1}}(s)} = A\beta_1 Z(s,t_N) \Bigg(& \Phi(d_+^1) + R_{J,t_{N-1}}(s)\phi(d_+^1)\frac{\partial d_+^1}{\partial R_{J,t_{N-1}}(s)} \\ & - \frac{K-\alpha_1}{\beta_1}\phi(d_-^1)\frac{\partial d_-^1}{\partial R_{J,t_{N-1}}(s)} \Bigg)\tau,\end{aligned} \tag{2.14}$$

where $\phi(x)$ is the standard normal probability density function evaluated at x and

$$\frac{\partial d_+^1}{\partial R_{J,t_{N-1}}(s)} = \frac{1}{R_{J,t_{N-1}}(s)\sigma\sqrt{t_{N-1}-s}} = \frac{\partial d_-^1}{\partial R_{J,t_{N-1}}(s)}.$$

Due to these tweaks, the delta hedge ratio, δ_s, as per equation 2.7 becomes:

$$\delta_s := \frac{\Delta_{C_P,s}}{\Delta_{FRA,s}}, \tag{2.15}$$

for prime caplets. The results from the first sample path of data are given for a 90% moneyness caplet that is priced and hedged. This caplet, with parameters as per section 2.2.1, is priced using the 90% moneyness BEV of 38%. The absolute strike at this moneyness level is $K = 10.35\%$ and the prime lending rate at expiry of the caplet is $R_P(t_{N-1}, t_N) = 11.5\%$. The static holding in the zero-coupon bond at t_0 is $\Delta_{ZCB,t_0} = 3.095$ and the hedge PnL at t_{N-1} is 0.0941. Figure 2.6 overleaf shows the bank account (b_s for $s \in [t_0, t_{N-1})$) and FRA holdings (δ_s for $s \in [t_0, t_{N-1})$) in the hedge portfolio:

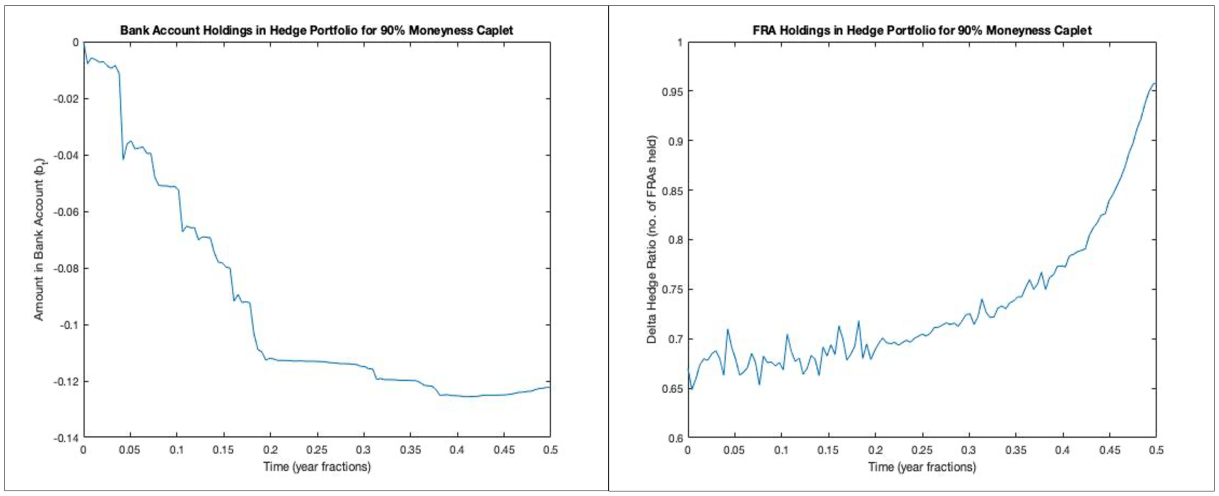

Fig. 2.6: Bank Account and FRA Holdings

For ease of comparison with the deep hedging approach, the hedge PnL distribution results across all moneyness levels are shown in section 4.3.2.

Chapter 3

Benchmarking the Neural Network

This chapter describes how the neural network is benchmarked to ensure that it is functioning correctly before using it to price and hedge caplets on the prime lending rate. First, the theoretical background of machine learning in pricing and hedging derivatives is presented in section 3.1. Thereafter in section 3.2, the implementation of this theory using simulated market data for the benchmarking exercise is shown.

The objective of the benchmarking analysis is to use the deep hedging approach to price and hedge the same 6x9 caplets analysed in section 2.1.2 where simulated data from a Libor Market Model is used. Given that simulated data from a log-normal model is used, we expect the deep hedging approach to yield delta hedge ratios that match that of the Black (1976) model - this would indicate that the neural network is functioning correctly.

3.1 Theoretical Background

This section serves to present the general theory around using machine learning to price and hedge derivatives. The theory behind pricing and hedging derivatives in incomplete markets is presented in section 3.1.1. This theory is used within the deep hedging framework which is then introduced in section 3.1.2.

3.1.1 Pricing and Hedging Derivatives in Incomplete Markets

Derivatives written on assets or processes which are not tradeable possess intrinsic risk which cannot be hedged away. Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a fixed time horizon $T \in (0, \infty)$ and a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ which satisfies the usual conditions. Furthermore, we assume a frictionless market with a single asset S,

and a bank account B which is a predictable process. Thus, the discounted price process of S is defined as $(X_t)_{0\leq t\leq T} := \frac{S_t}{B_t}$. In an incomplete market, there exist multiple equivalent martingale measures under which X is a local martingale. As per the Doob-Meyer decomposition, X_t can be represented as $X_t = X_0 + M_t + A_t$ where $X_0 \in \mathbb{R}$, M_t is a square integrable local martingale and A_t is predictable.

Suppose there exists an $\mathcal{F}_T$-measurable derivative with European style payoff denoted by H which we aim to price and hedge. To hedge against this derivative, we create a value process (portfolio) $(V_t)_{0\leq t\leq T}$ in which we trade in the asset S and the bank account. Thus, a hedging strategy is a pair of processes $(\xi_t, \eta_t)_{0\leq t\leq T}$ where $(\xi_t)_{0\leq t\leq T}$ is a predictable process and $(\eta)_{0\leq t\leq T}$ is adapted to the filtration $\mathbb{F}$. The value of the portfolio at time t is then given by:

$$V_t(\xi, \eta) = \xi_t S_t + \eta_t B_t, \tag{3.1}$$

i.e., ξ and η represent the holdings in S and B respectively.

The aim in pricing and hedging is to construct a hedging strategy such that $V_T(\xi_T, \eta_T) = H$ i.e., replicate the derivative. This has been shown to be possible in complete markets; however, in incomplete markets, we choose a hedging strategy that will minimise some measure of risk. In this regard, there are two overarching approaches to determine hedging strategies as compared by Hulley and McWalter (2015) - quadratic hedging and utility-based pricing.

Quadratic Hedging

The two main quadratic hedging approaches are local risk-minimisation and mean-variance hedging. As mentioned by Heath *et al.* (2001), the high-level difference between the two approaches is that local risk-minimisation gives simpler solutions for hedging strategies while mean-variance hedging allows for more control over total costs and risks. They compare the two approaches to analyse the costs and benefits of using the simpler solution and conclude that mean-variance hedging can be applied without a considerably larger amount of computational effort. This however, is under the assumption that the mean-variance hedging theory is well developed in the relevant area of interest.

Quadratic hedging techniques are a consistent extension of arbitrage-free pricing and hedging from complete markets to incomplete markets - they have been shown to recover the unique, arbitrage-free prices and hedging strategies in complete markets (Heath *et al.*, 2001). They also offer the advantage of yielding reasonably explicit solutions in very general models. A drawback of these techniques is that they go against traditional economic intuition. In particular, Heath *et al.*

(2001) mention the existence of an inconsistency in the connection between increasing utility of wealth and the fact that profits and losses are equally 'punished' by these methods. Thus, insatiable investors may not agree with these approaches. Furthermore, Hulley and McWalter (2015) note that these approaches rely on drift estimates of the asset price processes. For completeness, the two approaches are described as per Föllmer and Schweizer (1991) and Heath *et al.* (2001) in Appendix E.

Utility-Based Pricing

The idea behind utility-based pricing is to find a price p for a derivative such that an investor would be indifferent between paying nothing and not having the derivative H, and paying p now and receiving H at maturity T (Henderson and Hobson, 2004). This indifference is expressed in terms of the expected utility under optimal trading. More formally, if we define:

$$F(x,k) := \underset{X_T \in \mathcal{A}(x)}{\arg\max}\, \mathbb{E}[U(X_T + kH)],$$

where x is the investor's initial wealth, X_T represents the terminal wealth which can be generated from the initial wealth x, $\mathcal{A}(x)$ is the set of claims which can be replicated with initial wealth x, k is the number of derivative contracts and $U(.)$ is some concave utility function. The utility indifference price p is the solution to:

$$F(x-p,1) = F(x,0). \tag{3.2}$$

Thus, as shown by Monoyios (2004), in the application of hedging, we seek a hedging strategy (ξ^*, η^*) such that:

$$(\xi^*, \eta^*) = \underset{(\xi,\eta)}{\arg\max}\, \mathbb{E}[U(V_T(\xi,\eta) + kH)], \tag{3.3}$$

where V_T is the value process as defined in equation (3.1). Finding such a hedging strategy allows one to arrive at a time 0 price through the law of one price implied by the value process. Monoyios (2004) notes that equation (3.3) does not have a closed-form analytical solution and proposes using either a perturbation or cumulant expansion to compute a price and hedge parameters.

Utility indifference pricing has the advantage of being economically intuitive as noted by Henderson and Hobson (2004). It also allows the incorporation of coherent risk measures as shown by Buehler *et al.* (2019) and Xu (2006). Xu (2006) defines a coherent risk measure as follows:

Definition 3.1. Let X and $Y \in \mathcal{X}$ represent asset positions. A coherent risk measure $\rho : \mathcal{X} \to \mathbb{R}$ satisfies the following properties:

1. Subadditivity: $\rho(X+Y) \leq \rho(X) + \rho(Y)$.
2. Positive Homogeneity: If $\lambda \geq 0$, then $\rho(\lambda X) = \lambda\rho(X)$.
3. Monotone Decreasing: If $X \geq Y$, then $\rho(X) \leq \rho(Y)$.
4. Cash invariant: $\rho(X+c) = \rho(X) - c$ for $c \in \mathbb{R}$

Coherent risk measures are advantageous in banking due to regulatory requirements around risk management. By controlling the risk aversion parameters, one can also ensure that option prices recovered using this method are always positive. Henderson and Hobson (2004) also show that this approach recovers the unique prices and hedging strategies in complete market analyses.

On the contrary, this approach is limited in that it only allows for explicit calculations to be done in a few models consisting mainly of exponential utility functions (Henderson and Hobson, 2004). Furthermore, it is subjective in that one has to specify a utility function that market participants may not agree with. Another subjective issue is around the choice of a risk aversion parameter in the utility function. In particular, for exponential utility functions, Monoyios (2004) notes that this should be relatively small (around 0.01) to achieve low variances in the profit and loss distributions resulting from the optimal hedging strategies.

Xu (2006) shows that the concept of maximising utility in pricing can be replaced by minimising risk exposure with a focus on convex risk measures which are defined below:

Definition 3.2. Let $X, X_1,$ and $X_2 \in \mathcal{X}$ represent asset positions. Xu (2006) define $\rho : \mathcal{X} \to \mathbb{R}$ to be a convex risk measure if it is:

1. Monotone decreasing: if $X_1 \geq X_2$ then $\rho(X_1) \leq \rho(X_2)$.
2. Convex: $\rho(\alpha X_1 + (1-\alpha)X_2) \leq \alpha\rho(X_1) + (1-\alpha)\rho(X_2)$ for $\alpha \in [0,1]$.
3. Cash invariant: $\rho(X+c) = \rho(X) - c$ for $c \in \mathbb{R}$.

Furthermore, ρ is normalised if $\rho(0) = 0$.

Thus, a convex risk measure differs from a coherent risk measure in that the first two conditions (subadditivity and positive homogeneity) for a coherent risk measure are relaxed to require convexity. As per Buehler *et al.* (2019), given a convex risk measure ρ, we can interpret $\rho(H)$ as the minimum amount c that needs to be added to a position H in order for the position to be 'acceptable' to an investor. 'Acceptable' in this sense means that $\rho(H+c) \leq 0$. We define:

$$\pi(H) := \inf_{(\xi,\eta)\in\mathcal{H}} \rho(H + V_T(\xi,\eta)), \tag{3.4}$$

where $\mathcal{H}$ is the set of admissible trading strategies. Buehler *et al.* (2019) show that π is a convex risk measure if $\mathcal{H}$ is convex. Therefore, $\pi(H)$ is interpreted as the minimum amount required to make the terminal position, when taking a position in the derivative H and using a trading strategy (ξ_t, η_t), acceptable to an investor. Relating this to indifference pricing, Buehler *et al.* (2019) show that p is the indifference price of the derivative where $\pi(-H + p) = \pi(0)$. Given that π is a convex risk measure:

$$\pi(-H + p) = \pi(0) \implies \pi(-H) - p = \pi(0). \tag{3.5}$$

Buehler *et al.* (2019) show that under the above setting $\pi(0) = 0$, hence the problem reduces to finding:

$$p = \pi(-H). \tag{3.6}$$

3.1.2 Machine Learning, Pricing and Hedging

Deep Neural Networks

The theory in this section is drawn from Higham and Higham (2019). A neural network consists of multiple layers through which information is passed. The first and last layers are the input and output layers respectively, with the intermediate layers being called hidden layers. Each layer contains a specified number of nodes that are connected by non-linear functions. These connections are known as edges and each edge has a set of weights and bias parameters which are continuously adjusted as learning proceeds. Within a layer, each node receives one real value from every node in the previous layer and produces one real valued output passed to the next layer. Suppose we have an L layered neural network trying to model the function F^*, with each layer containing n_l nodes, where $l \in \{1, \dots, L\}$. Then, the neural network defines a mapping $F : \mathbb{R}_{n_1} \to \mathbb{R}_{n_l}$ and tries to learn the parameters of F to give the best approximation of F^*.

To train the neural network using training data $\mathbf{x}$, we specify a cost function that is minimised to drive F to F^*. This cost function will be a function of the aforementioned weights and biases, as well as the inputs to the neural network. If there is a set of labels $\mathbf{y} = F^*(\mathbf{x})$ that the network is targeting, this is supervised learning. If there is no set of labels fed to the network and it is left to learn from the data by itself, this is unsupervised learning and allows for more complex processing. A classical method to minimise the cost function by estimating the weights and biases of the network is the gradient descent method. This method seeks a local minimum of the cost function in an iterative manner by taking steps proportional to the negative of the gradient of the cost function at the current point of evaluation. This method becomes computationally expensive when there are a large number of parameters

and training points in the network. Hence, the stochastic gradient method is often preferred. In this method, a single training point is chosen at random from the entire training set at each iterative step. To solve for the gradient of the cost function, partial derivatives with respect to every weight and bias are needed. These derivatives need to be equal to zero to ensure that minimisation takes place and a back-propagation technique is used to ensure this. Pseudocode for the stochastic gradient and back-propagation techniques can be found in Higham and Higham (2019).

Deep Hedging

Deep hedging is a relatively novel concept that makes use of a neural network to obtain optimal hedging strategies for derivatives. The price of the derivative is a by-product of the neural network output. The optimisation of the weights and biases mentioned above essentially models a function for the optimal hedge ratios at each time-step. Therefore, the cost function in deep hedging tends to be a quadratic risk function or convex risk measure. Being a novel concept, the literature on deep hedging in mathematical finance is sparse. Two notable publications in this regard are Buehler *et al.* (2019) and Halperin (2020). Both publications highlight how deep hedging allows for the pricing and hedging of derivatives in a model-independent manner. This allows one to work in a $\mathbb{P}$-measure world rather than a $\mathbb{Q}$-measure one when pricing derivatives i.e., real-world data can be used directly. This also has implications that deep hedging is more effective in liquid markets given that a large amount of data is required to train the neural network. Being model independent, deep hedging has scope for application in incomplete markets. Furthermore, using a neural network also means that parameters such as volatility need not be directly estimated.

The approach by Halperin (2020) is more aligned to quadratic risk minimisation while that of Buehler *et al.* (2019) is based on indifference pricing under convex risk measures. Halperin (2020) notes that as the risk aversion parameter tends to zero in the indifference pricing framework, the quadratic risk minimisation and indifference pricing approaches yield the same results. Buehler *et al.* (2019) show that an advantage of the indifference pricing approach is that a trader has more control over the hedge profit and loss distributions through the specification of the risk aversion parameter in the convex risk measure which is used.

To incorporate risk preferences in the formulation shown in equation (3.6), Buehler *et al.* (2019) assumes that the agent has utility represented by the exponential utility function $U(x) = -\exp(-\lambda x)$. This function is commonly used in literature (Monoyios (2004),Henderson and Hobson (2004)). Buehler *et al.* (2019) demon-

strates that under the assumption of the exponential utility function, $\pi(-H)$ can be approximated by:

$$\pi^M(-H) = \frac{1}{\lambda} \log \inf_{(\xi,\eta)\in\mathcal{H}} J(\xi,\eta) \tag{3.7}$$

where

$$J(\xi,\eta) = \mathbb{E}[e^{-\lambda(-H+V_T(\xi,\eta))}] \tag{3.8}$$

3.2 Implementation

This section describes how the deep hedging approach is implemented for the benchmarking analysis. First, the setup of the neural network structure is described in section 3.2.1. Thereafter, the methodology used to train the network and conduct a deep hedging analysis on simulated market data is described in section 3.2.2. Finally, the results from the benchmarking analysis are presented in section 3.2.3.

3.2.1 Setting up the Neural Network

The neural network is created using the Tensorflow machine learning library in Python. The network is adapted from an existing Long-Short-Term-Memory (LSTM) network created by Groncki (2018) who uses it to investigate deep-hedging of equity derivatives. LSTM networks allow for past information to be recalled when making future predictions. This network differs slightly from that used by Buehler *et al.* (2019) who use a recurring feed-forward neural network. This feed-forward network is simpler in that it does not recall any past information when making future predictions. The network used consists of an input layer, three hidden layers and an output layer - a simple schematic of the network is shown in Figure 3.1 below:

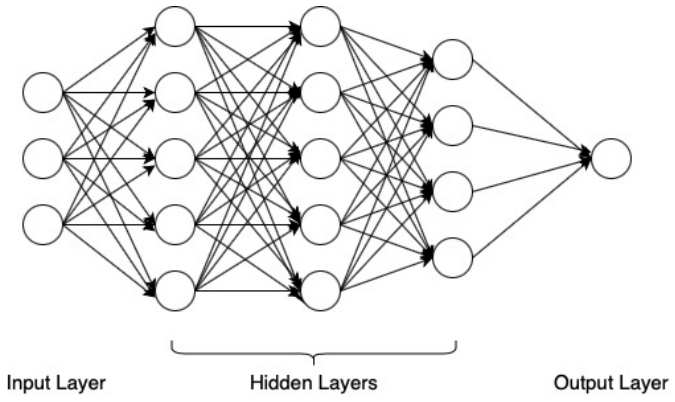

Fig. 3.1: Simple Neural Network Schematic

At each daily time-step, there are three inputs to the neural network: the market value of the hedge instrument (the FRA), the overnight risk-free rate used to grow the bank account, and the Black (1976) caplet value. The single output at each time-step is the optimal delta hedge ratio i.e., the required holding in the FRA as per section 2.1.2. In terms of the structure of the network, the number of nodes in each of the three hidden layers is fixed at 62, 48 and 24 respectively as per Groncki (2018).

LSTM network nodes can be represented by cells as shown in Figure 3.2 below where X_t would be the array of inputs and h_t is the optimal delta hedge ratio at time-step t:

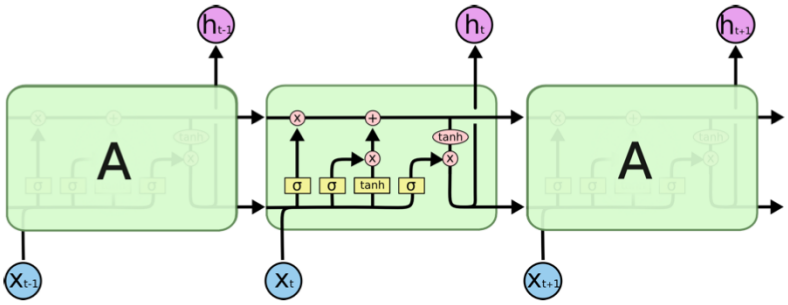

Fig. 3.2: LSTM Cell (Olah, 2015)

The following description is drawn from Olah (2015). The horizontal line running through the top of the cell is called the cell state. Information flows through it with potential linear interactions which may arise due to further information being introduced from the gates below it. An LSTM cell has three gates to 'protect and control' the cell state. The gates apply a sigmoid function to their inputs and then pointwise multiplication. The sigmoid function ensures that the outputs from the gates are between 0 and 1 where 0 allows no further information to be introduced and 1 allows the maximum amount of new information. The second and third gates contain $\tanh$ functions which respectively allow the cell to decide what information to keep and what to output (h_t). The use of the tanh function for the cell output ensures that the delta hedge ratios are constrained between $[-1, 1]$.

3.2.2 Training the Neural Network

Input Data

As mentioned in section 3.2.1, the three inputs to the neural network are the Black (1976) caplet value, the market value of the FRA used for hedging and the risk-free overnight rate at each time-step. The Black (1976) caplet values are used as an input in this exercise given that the goal is to check if the neural network can

output delta hedge ratios identical to the Black (1976) model. Typically, in a model independent deep hedging approach, only the final caplet payoffs would be used as an input as opposed to the Black (1976) caplet values at each time-step. Note that this specific implementation of a neural network has the limitation that only one hedge instrument can be used - hence, there is no static hedge in the zero-coupon bond as per the model dependent hedging methodology set out in section 2.1.2.

For the benchmarking analysis, the FRA used for hedging is identical to that used in the pricing and hedging analysis in section 2.1.2 using simulated Libor Market Model data. Hence from the $100,000$ sample paths of simulated data, the daily risk-free overnight rates, market values of the FRA and Black (1976) caplet values (with a nominal of $A = 1000$ and a moneyness level of 110%) are extracted for use as inputs to the neural network.

Training Algorithm

The following 3 step procedure is used to train the neural network using the above input data to price and hedge 6x9 caplets:

1. Create the neural network comprising of an input layer, three hidden layers and an output layer.
2. Specify a cost function for the neural network to minimise.
3. Feed in the input data and train the network to minimise the cost function.

The Adam optimisation algorithm, which is an extension of the stochastic gradient descent algorithm, is used to minimise the cost function in the network. This algorithm, developed by Kingma and Ba (2014), is said to work well for stochastic cost functions and only requires first-order partial derivatives of the cost function with respect to its parameters. 'The name Adam is derived from adaptive moment estimation' - that is, 'the method computes individual adaptive learning rates for the different parameters from estimates of first and second moments of the gradients' (Kingma and Ba, 2014). For completeness, the algorithm is shown in Appendix F.

Cost Functions

First, a quadratic risk cost function is used to train the neural network in the benchmarking analysis. The mean-variance approach, which aims to minimise the global risk (hedge PnL variance) in hedging the caplet, while assuming a mean hedge PnL

of zero, is used. This approach has the advantage of insisting that the value process of the hedge portfolio is self-financing. Therefore, a discrete-time version of equation (E.7) is used where the Black (1976) model caplet value is used for the initial price. This ensures that the neural network is best positioned to drive towards fitting the Black (1976) model delta hedge holdings function. The discrete-time version of the quadratic risk cost function used is of the following form:

$$\left(Z(t_{N-1},t_N)C(t_N) - C(t_0) - \sum_{s=0}^{N-1} h_s(V_{FRA}(t_{s+1}) - V_{FRA}(t_s))\right)^2 \tag{3.9}$$

where $Z(t_{N-1},t_N)C(t_N)$ is the terminal caplet payoff discounted to expiry t_{N-1}, $C(t_0)$ is the t_0 Black (1976) caplet price and h_s are the optimal delta hedge ratios which the neural network optimises for.

Thereafter, the findings of Buehler *et al.* (2019) that hedge PnL distributions can be controlled by traders under the utility indifference pricing framework is tested. The convex risk measure chosen is CVaR (ρ) which is defined as follows:

$$\rho(X) := \mathbb{E}[X|X > VaR_\alpha(X)] = \frac{1}{1-\alpha}\int_{-\inf}^{1-\alpha} VaR_\gamma(X)\, d\gamma \tag{3.10}$$

where X is the hedge PnL, $1-\alpha$ is the confidence level and $VaR_\gamma(X) = inf\{m \in \mathbb{R}^+ : \mathbb{P}(X < -m) \leq \gamma\}$. Confidence levels of 50% and 95% are used for comparison.

Network Optimisation

During the training of the neural network, a choice has to be made on the batch size and number of epochs to use. The batch size refers to the number of sample paths used at a time to train the network over. The number of epochs refers to the number of passes that are done over the entire data set for training. Thus, varying the batch size and number of epochs allows for some optimisation of the neural network performance.

Using the quadratic risk cost function, the batch size is first varied while keeping the number of epochs constant at 100. Once the optimal batch size is chosen, the number of epochs is varied. The tables below show the results from each optimisation process:

Tab. 3.1: Benchmarking - Batch Size Optimisation

Batch Size	Mean Hedge PnL	Hedge PnL Standard Deviation	Training Time (min)
500	-0.0058	0.0606	48
1000	-0.0056	0.0681	55
2500	-0.0056	0.1161	60
5000	-0.0057	0.1172	160

Tab. 3.2: Benchmarking - Number of Epochs Optimisation (Batch Size = 500)

Number of Epochs	Mean Hedge PnL	Hedge PnL Standard Deviation	Training Time (min)
100	-0.0058	0.0606	48
200	-0.0051	0.0601	89
400	-0.0044	0.0592	195
1000	-0.0040	0.0581	432

From Table 3.1, it is clear that the neural network optimises the mean hedge PnL extremely well using the simulated Libor Market Model data. The main deciding factor in choosing a batch size is the standard deviation of the hedge PnL distributions. Conveniently, the smaller batch sizes, which also require the least training time, produce the lowest standard deviations. Therefore, a batch size of 500 is chosen. Table 3.2 highlights that increasing the number of epochs produces better results in terms of the hedge PnL distributions. However, this comes at the cost of extremely long training times. Therefore, a balance has to be struck between optimal results and a reasonable training time given that the network would require training for each moneyness level at which a caplet is priced - this led to a choice of 200 epochs.

3.2.3 Benchmarking Results

As mentioned at the beginning of the chapter, the goal of the benchmarking exercise is to check if the neural network can predict optimal hedge ratios as close as possible to the lognormal Black (1976) model using the Libor Market Model simulated data from section 2.1.2. Using the quadratic cost function, the optimised neural network is trained and the following predicted delta hedge ratio time series are extracted from 3 random sample paths of the $100,000$:

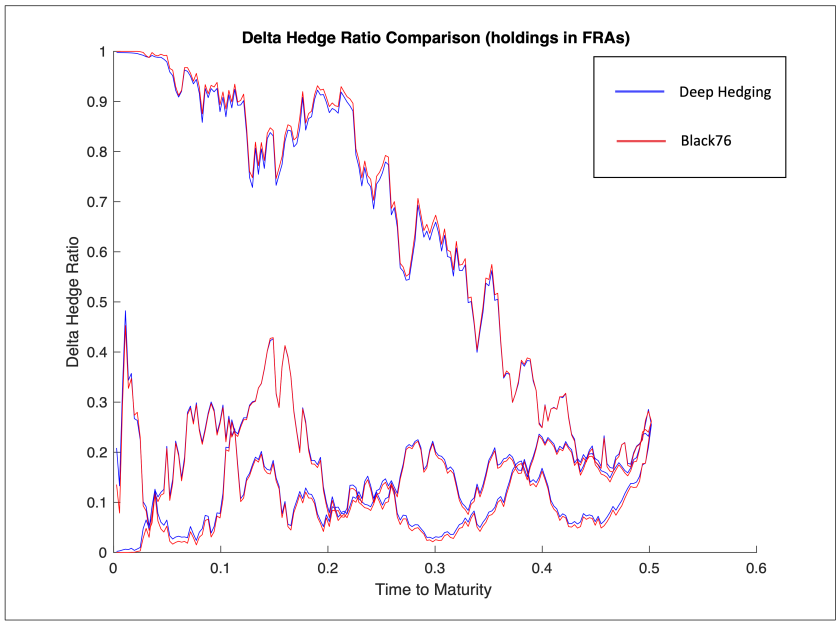

Fig. 3.3: Comparison of Black (1976) and Neural Network Delta Hedge Ratios

The deep hedging predicted delta hedge ratios are extremely close to those from the Black (1976) model. Under the Black (1976) model's pricing and hedging framework, it is assumed that interest rate data is lognormally distributed under the respective T-forward measure. Given that simulated lognormal Libor Market Model data is used in this exercise, the optimal hedging methodology is that specified by the Black (1976) model. Hence, these results indicate that the neural network is functioning correctly as it can fit the Black (1976) delta hedge ratio function to the simulated data while minimising the hedge PnL variance. Given that it is impractical to show the above plots for all sample paths, the average relative errors between the deep hedging and Black (1976) delta hedge ratios at each time-step are computed and shown in Figure 3.4 overleaf. It can be seen that the deep hedging delta hedge ratios on average, do not differ from the Black (1976) delta hedge ratios by more than 0.3% at each time-step.

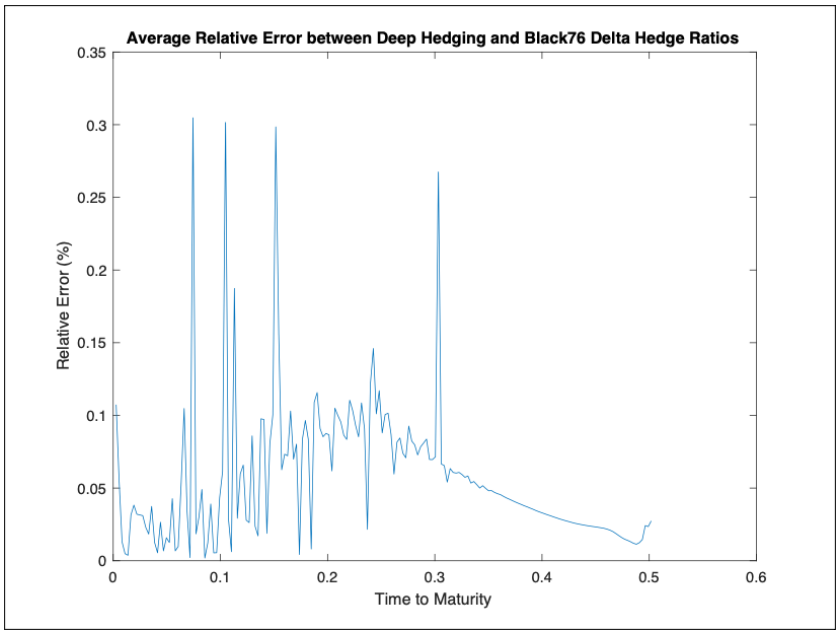

Fig. 3.4: Average Relative Error between Deep Hedging and Black (1976) Delta Hedge Ratios

Lastly, the utility indifference pricing framework is tested using a discretised version of equation 3.10 with two different risk aversion parameters. Figure 3.5 overleaf shows the hedge PnL distributions for confidence levels of 50% and 95% respectively, with a higher confidence level indicating less preference for risk:

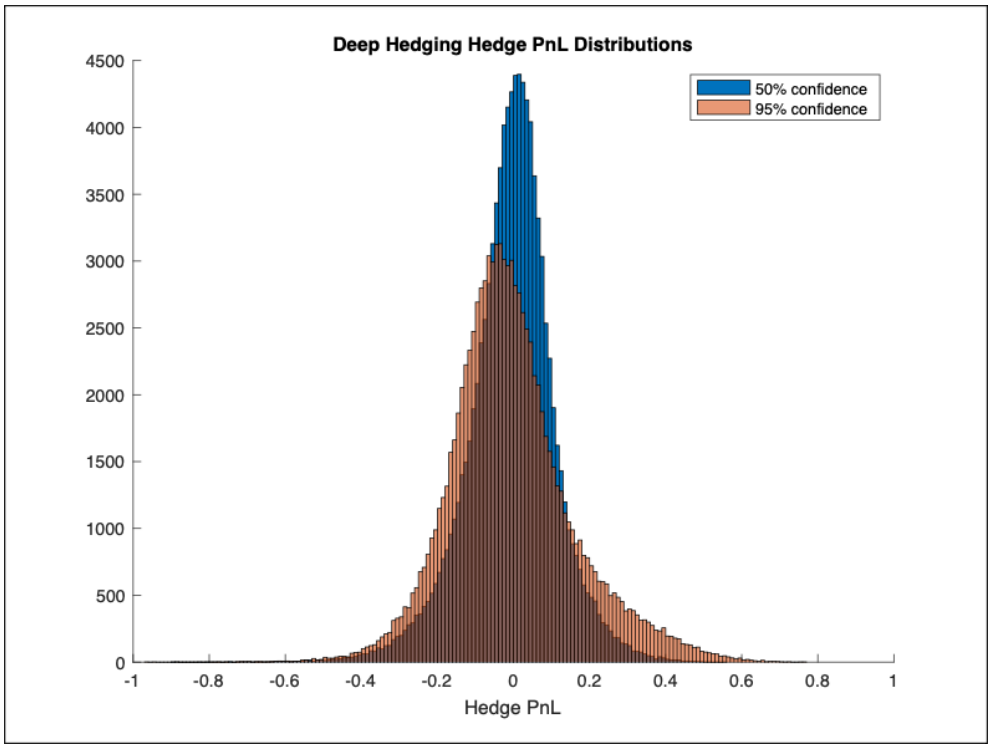

Fig. 3.5: Hedge PnL Distribution with 50% and 95% Confidence Levels

It is clear that under the utility indifference pricing framework, a trader can control their hedge PnL distribution by varying the risk aversion parameter when CVaR is used as a cost function in the deep hedging framework. The 95% confidence level distribution in Figure 3.5 is slightly skewed to favour positive hedge PnLs over negative ones.

Chapter 4

Deep Hedging Caplets on the Prime Lending Rate

The purpose of this chapter is to present the deep hedging approach to pricing and hedging caplets written on the prime lending rate. First, the methodology used to train the neural network is outlined in section 4.1. Thereafter, the manner in which the neural network outputs are analysed is presented in section 4.2. The results from the implementation of this methodology are then presented in section 4.3.

The prime caplets priced in the deep hedging approach are 6x9 caplets with the same parameters as per the model dependent approach. The structure of the neural network used to price and hedge prime caplets is identical to that outlined in section 3.2.1. That is, an LSTM network with an input layer, three hidden layers and an output layer. Again at each time-step, there are three inputs and one output (the optimal delta hedge ratio). The number of nodes in each of the three hidden layers is as per the neural network used in the benchmarking analysis.

4.1 Training the Neural Network

4.1.1 Input Data

The 3740 normalised sample paths of real-world interest rate data used in the model dependent approach as per section 2.2.5 are used to price and hedge the prime caplets. Again, this specific neural network implementation only allows for a single hedge instrument. Hence, there is no static hedge taken in the zero-coupon bond. The hedge instrument used is the same as that in the model dependent approach - a fair FRA (at t_0) on 3-month JIBAR with the same parameters as the caplet i.e., nominal $A = 1000$, tenor τ, reset time t_{N-1} and settlement time t_N.

Given the setup of the neural network, an array of three inputs is required at each time-step. Thus, two of the three inputs at each time-step are the mar-

ket value of the 3-month JIBAR FRA and the daily risk-free overnight rate used to grow the bank account. This application of the deep hedging approach differs from the benchmarking analysis in section 3.2 in that Black (1976) model caplet values are not used as an input. Instead, only the final caplet payoff is used. However, since an array of the three inputs is needed at each time-step and the caplet payoff only occurs at the final time-step; for every time-step prior to the final one, the caplet payoff value in the input array is fixed at zero. This makes the approach completely model independent.

4.1.2 Training Algorithm

The following 5 step procedure is used to train the neural network to price and hedge prime caplets. Note that steps 2-5 of this procedure are repeated for each moneyness level, from 80% to 120%, in increments of 10%:

1. Compute the market value of the FRA on 3-month JIBAR at each time-step until maturity of the caplet, for each sample path.
2. Compute the payoff of the caplet at maturity for each sample path.
3. Create a neural network comprising of an input layer, three hidden layers and an output layer.
4. Specify a cost function for the neural network to minimise.
5. Feed in the input data and train the network to minimise the cost function using the Adam optimiser.

Cost Function

A quadratic risk cost function is used to train the neural network. As per the benchmarking analysis, the mean-variance approach, which aims to minimise the hedge PnL variance, while assuming a mean hedge PnL of zero, is used. The quadratic cost function used in the benchmarking analysis is tweaked slightly in that no initial caplet price is specified. Instead, an initial caplet price is backsolved for using the optimised neural network output. Thus, the cost function takes the following form:

$$\left(Z(t_{N-1}, t_N) C_P(t_N) - \sum_{s=0}^{N-1} h_s (V_{FRA}(t_{s+1}) - V_{FRA}(t_s)) \right)^2 \tag{4.1}$$

where $Z(t_{N-1}, t_N) C_P(t_N)$ is the terminal prime caplet payoff discounted to expiry t_{N-1} and h_s are the optimal delta hedge ratios which the neural network optimises for.

4.2 Post Processing

The cost function specified by equation (4.1) allows for a model independent prime caplet price, p, to be backsolved for. It essentially implies that the caplet is sold with no initial premium and the neural network tries to find a way to optimally hedge the exposure. Using the optimised delta hedge ratios from the neural network output, the initial caplet price is backsolved for by making use of the delta hedging routine specified in section 2.1.2 as per Musiela and Rutkowski (2005). Taking into account that there is no static zero-coupon bond hedge in the deep hedging approach, the delta hedging routine needs to be modified as follows:

1. At t_0, receive p, the model independent price, from the sale of the caplet and purchase h_{t_0} many fair 3-month JIBAR FRAs with the properties as per section 4.1.1 at zero cost given that they are fair FRAs.

2. For all $s \in [t_0 + 1, t_{N-1} - 1]$, adjust the FRA position from the previous day, h_{s-1}, to ensure the hedge portfolio has h_s many FRAs while adjusting the holding in the bank account b_s to fund this if necessary or deposit any excess cash that may arise.

3. At time t_{N-1}, the hedge PnL is calculated as:

$$PnL = b_{t_{N-1}-1} C^{ON}(t_{N-1} - 1, t_{N-1}) + h_{t_{N-1}-1} V_{FRA}(t_{N-1}) - Z(t_{N-1}, t_N) C_P(t_N),$$

   where $C^{ON}(t_{N-1} - 1, t_{N-1})$ is a risk-free overnight capitalisation factor over the period $[t_{N-1} - 1, t_{N-1}]$ and $Z(t_{N-1}, t_N)C_P(t_N)$ is the terminal caplet payoff discounted to t_{N-1}.

To solve for an initial price p for each sample path, the PnL in step 3 is equated to zero. Thereafter, using the optimal hedge ratios from the neural network output, h_s for $s \in [t_0, t_{N-1} - 1]$, we backsolve for the bank account holdings from $t_{N-1} - 1$ to t_0. The model independent caplet price for each sample path is then given by $p = b_{t_0}$. Finally, the single caplet price for each moneyness level is assumed to be the average across all sample paths.

For each moneyness level, using this single initial model independent price and the deep hedging delta hedge ratios from each sample path, hedge PnLs are solved for and the hedge PnL distributions are compared with the model dependent approach. To allow for a direct comparison between the two approaches, the model dependent pricing and hedging process outlined in section 2.1.2 is performed for prime caplets without a static zero-coupon bond hedge i.e., $\Delta_{ZCB,t_0} = 0$. Lastly,

an implied Black (1976) model volatility is solved for using the initial model independent caplet price at each moneyness level to create a deep hedging implied volatility skew to compare with the model dependent BEV skew.

4.3 Results

4.3.1 Optimising the Neural Network

A similar neural network optimisation process to that outlined in section 3.2.2 is performed. The results from the batch size and number of epochs optimisations respectively are shown below. The optimisation process is done on prime caplets at the moneyness level of 100%.

Tab. 4.1: Actual Market Data - Batch Size Optimisation

Batch Size	Mean Hedge PnL	Hedge PnL Standard Deviation	Training Time (min)
10	1.47	2.64	28
20	1.34	2.51	21
50	1.63	2.89	16
100	1.72	2.77	12
200	1.78	2.99	9
500	2.18	3.22	7

Tab. 4.2: Actual Market Data - Number of Epochs Optimisation (Batch Size = 20)

Number of Epochs	Mean Hedge PnL	Hedge PnL Standard Deviation	Training Time (min)
100	1.34	2.51	21
200	0.71	1.71	42
400	0.77	1.64	87
1000	0.44	1.25	180

Interestingly, when training the network using the actual market data which has much fewer sample paths, using a larger batch size reduces the training time. A batch size of 20 is chosen given that this parameter produces the minimal mean hedge PnL and hedge PnL standard deviation. The training time of 21 minutes is acceptable for the performance improvement. The number of epochs chosen is 200. Any further increase in the number of epochs is deemed unnecessary given the significant increase in training time for just a marginal improvement in the hedge PnL distribution parameters.

4.3.2 Deep Hedging

The table below compares the model dependent and deep hedging (DH) hedge PnL distribution parameters across moneyness levels for prime caplets:

Tab. 4.3: Prime Caplet Hedge PnL Results

Moneyness	DH Mean Hedge PnL	DH Hedge PnL SDev	Black76 Mean Hedge PnL	Black76 Hedge PnL SDev	Black76 Mean Hedge PnL (no ZCB)	Black76 Hedge PnL SDev (no ZCB)
80%	-1.33	1.28	-0.001	2.62	-5.46	2.62
90%	-0.136	0.364	0.006	2.33	-2.92	2.33
100%	-0.012	0.057	0.003	0.972	-0.797	0.972
110%	6.38E-04	0.014	0.001	0.165	-0.089	0.165
120%	1.55E-08	5.41E-07	0	0	0	0

Based on the generally low hedge PnL means and standard deviations across moneyness levels, the West (2008) approach provides a good solution to dealing with market incompleteness when pricing caplets on the prime lending rate within a Black (1976) model framework. For caplets at moneyness levels $\geq 100\%$, both approaches perform extremely well with little to no mean hedge PnL and hedge PnL standard deviation.

The final two columns in Table 4.3 contain PnL distribution parameters from the Black (1976) model where the static hedge on the zero-coupon bond (ZCB) is excluded. The hedge PnL standard deviations are the same as those from the standard Black (1976) model results because the position taken in ZCBs is static - hence, there is no additional variance created. The deep hedging approach always produces hedge PnL distributions that have lower standard deviations and hence, variability.

Looking at the mean hedge PnLs, the standard Black (1976) hedging approach (including the ZCB) performs better than the deep hedging approach. However, as mentioned above, the deep hedging approach performs better than the standard Black (1976) model approach in terms of the variability of the hedge PnL. Looking at the more direct comparison between the deep hedging approach and Black (1976) model approach without ZCBs, the neural network outperforms the Black (1976) model in all cases.

Some limitations to keep in mind about the deep hedging approach are that only a single hedge instrument is used in this specific implementation and that the neural network needs to be retrained for each moneyness level. Furthermore, if caplets with different expiry dates and tenors need to be priced, the neural network would also have to be retrained. Therefore, the model independent approach can

become extremely time-consuming which could be disadvantageous on a trading floor.

The final comparison between approaches involves the model dependent BEV skew and the implied deep hedging volatility skew. Figure 4.1 below shows the deep hedging implied volatility skew along with the Black (1976) BEV skew from the model dependent approach:

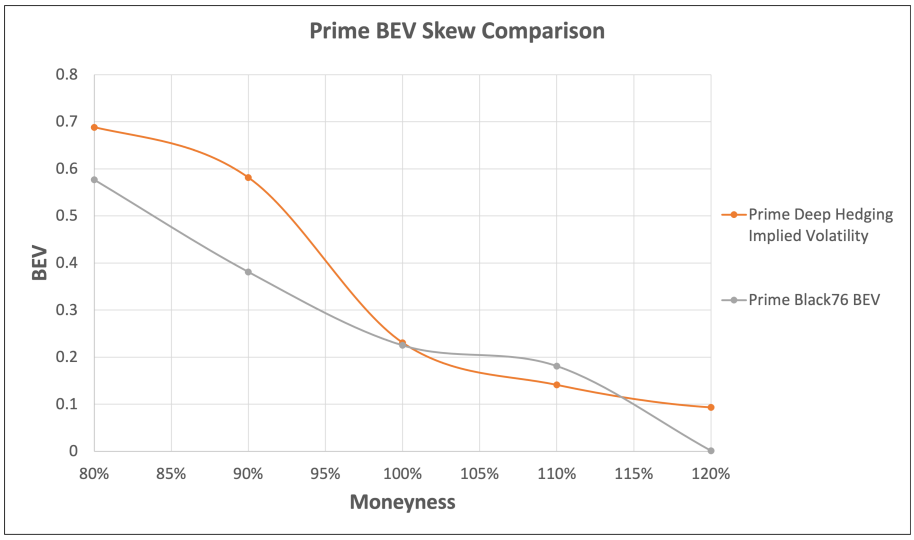

Fig. 4.1: Prime BEV Skew Comparison

The differences between the skews are larger for moneyness levels less than 100%. As a trader, one would be concerned with what volatility to use when pricing a prime caplet given the variations seen above. The West (2008) approach is a good first-order approach. But the deep hedging implied skew shows that one would need to take more care when pricing prime caplets at lower moneyness levels. This may be due to caplets at moneyness levels below 100% almost always ending up in-the-money as shown in section 2.2.4. At moneyness levels above 100%, the differences between the skews drop significantly with the difference at 120% moneyness technically irrelevant. This is because it is found that using a volatility as high as 0.15 results in hedge PnL distribution parameters of order 10^{-5}. Overall, this result also indicates that the deep hedging approach has the advantage of providing a comparable implied volatility skew which can be used in a model dependent analysis in pricing and hedging prime caplets.

Chapter 5

Conclusion

The aim of this dissertation, of investigating the effectiveness of deep hedging in pricing and hedging caplets on the prime lending rate, has been achieved.

A benchmark model dependent framework has been set up successfully using the approach set out by West (2008) to deal with the market incompleteness that arises. This approach, which allows for the use of the market standard Black (1976) model, proves to be sufficient because of the cointegration relationships between the prime and 3-month JIBAR rates. This relationship allows for the transformation of a prime caplet into an equivalent 3-month JIBAR caplet. The break-even volatility concept introduced by Dupire (2006) then allows for the computation of a volatility skew to use in the Black (1976) model for pricing and hedging where hedging using a ZCB and FRAs on 3-month JIBAR proves sufficient.

A neural network has been set up and benchmarked successfully to conduct the model independent, deep hedging analysis. The benchmarking analysis has shown that within the deep hedging framework, a trader's risk preferences can be incorporated in the pricing and hedging exercise. The neural network is able to price and hedge the prime caplets to an extent that is comparable to the model dependent approach. The limitation of only being able to use one hedge instrument (FRAs on 3-month JIBAR) means that the neural network does not always outperform the standard model dependent approach (which uses the static ZCB hedge). However, when the model dependent approach is restricted to only using the FRAs as a hedge instrument, the neural network always outperforms it. Lastly, it has been shown that the deep hedging framework can also be used to imply a Black (1976) volatility skew which can be considered when pricing and hedging within a model dependent framework.

5.1 Future Recommendations

Looking ahead, to build on this analysis, further research can be done on the optimisation of neural network performance in an attempt to further minimise the absolute hedge PnLs. For instance, the structure of the network can be optimised by varying the number of layers or nodes in each layer as this is not addressed in this dissertation. Furthermore, should time allow, more epochs can be used as this has already shown improved performance.

The results of both the model dependent and independent approaches may be improved using a richer data set of historical interest rate data. To this extent, data augmentation techniques to extend this data set may be investigated. Alternatively, sourcing data dating back further in time is also an option.

In terms of the West (2008) approach to pricing and hedging caplets on the prime lending rate, alterations may be made to the pricing model that is used. As seen in the BEV analysis, using the Black (1976) model does not result in the BEV skews of caplets and floorlets being exactly the same - which should be the case by put-call parity. A possible consideration in this regard could be to fit a multi-factor interest rate model to the real-world data which may better represent actual interest rate changes through time. Thereafter, this fitted model may be used to test the pricing and hedging performance in a model dependent world.

Lastly, the model independent approach can be further improved by tweaking the neural network architecture to allow for two hedge instruments to be used. This would allow for the possibility of including the static ZCB hedge which has already shown improved performance in the model dependent analysis.

Bibliography

Alexander, C. (2008). *Market Risk Analysis, Practical Financial Econometrics*, John Wiley & Sons.

Black, F. (1976). The pricing of commodity contracts, *Journal of Financial Economics* **3**(1-2): pp. 167–179.

Brace, A., Gatarek, D. and Musiela, M. (1997). The market model of interest rate dynamics, *An International Journal of Mathematics, Statistics and Finanical Economics* **7**(2): pp. 127–155.

Buehler, H., Gonon, L., Teichmann, J. and Wood, B. (2019). Deep hedging, *Quantitative Finance* **19**(8): pp. 1271–1291.

Dickey, D. A. and Fuller, W. A. (1981). Likelihood ratio statistics for autoregressive time series with a unit root, *Econometrica: Journal of the Econometric Society* pp. 1057–1072.

Dupire, B. (2006). Fair skew: Break-even volatility surface, *Bloomberg LP* .

Engle, R. F. and Granger, C. W. (1987). Co-integration and error correction: representation, estimation, and testing, *Econometrica: Journal of the Econometric Society* pp. 251–276.

Föllmer, H. and Schweizer, M. (1991). Hedging of contingent claims under incomplete information, *in* M. H. A. Davis and R. J. Elliott (eds), *Applied Stochastic Analysis*, Gorden and Breach Science Publishers.

Groncki, M. (2018). Option hedging with long-short-term-memory recurrent neural networks part i, WordPress.
URL: *https://ipythonquant.wordpress.com/2018/06/05/option-hedging-with-long-short-term-memory-recurrent-neural-networks-part-i/*

Gupta, A. and Subrahmanyam, M. G. (2005). Pricing and hedging interest rate options: Evidence from cap–floor markets, *Journal of Banking & Finance* **29**(3): pp. 701–733.

Halperin, I. (2020). Qlbs: Q-learner in the black-scholes (-merton) worlds, *The Journal of Derivatives* **28**(1): pp. 99–122.

Heath, D., Platen, E. and Schweizer, M. (2001). A comparison of two quadratic approaches to hedging in incomplete markets, *Mathematical Finance* **11**(4): pp. 385–413.

Henderson, V. and Hobson, D. (2004). Utility indifference pricing-an overview, Princeton University Press, Princeton, NJ, USA.

Higham, C. F. and Higham, D. J. (2019). Deep learning: An introduction for applied mathematicians, *SIAM Review* **61**(4): pp. 860–891.

Hulley, H. and McWalter, T. A. (2015). Quadratic hedging of basis risk, *Journal of Risk and Financial Management* **8**(1): pp. 83–102.

Hunter, C. J., Jäckel, P. and Joshi, M. S. (2001). Drift approximations in a forward-rate-based libor market model, *Getting the Drift* pp. pp. 81–84.

Johansen, S. and Juselius, K. (1990). Maximum likelihood estimation and inference on cointegration—with appucations to the demand for money, *Oxford Bulletin of Economics and Statistics* **52**(2): pp. 169–210.

Kingma, D. P. and Ba, J. L. (2014). Adam: A method for stochastic optimization, *ICLR*.

Monoyios, M. (2004). Performance of utility-based strategies for hedging basis risk, *Quantitative Finance* **4**(3): pp. 245–255.

Musiela, M. and Rutkowski, M. (2005). *Martingale Methods in Financial Modelling: Second Edition*, Springer-Verlag Berlin Heidelberg.

Olah, C. (2015). Understanding lstm networks.
URL: *https://colah.github.io/posts/2015-08-Understanding-LSTMs/*

Vasicek, O. (1977). An equilibrium characterization of the term structure, *Journal of Financial Economics* **5**(2): pp. 177–188.

West, G. (2008). Interest rate derivatives in the south african market based on the prime rate, *Studies in Economics and Econometrics* **32**(1): pp. 75–87.

Xu, M. (2006). Risk measure pricing and hedging in incomplete markets, *Annals of Finance* **2**(1): pp. 51–71.

Appendix A

The Vasicek Model

The Vasicek (1977) model is a short-rate interest rate model described by the following dynamics:

$$dr = a(b - r) + \sigma_{vas} dW_t \tag{A.1}$$

where r is the short rate, W_t is a standard Brownian motion and $a, b, \sigma_{vas} \in \mathbb{R}^+$ represent the rate of mean reversion, the mean reversion level and the volatility respectively. The following theorem about affine term structure models is used in simulating an initial yield curve for the benchmarking exercise in the model independent analysis:

Theorem: A short rate model with risk neutral dynamics $dr_t = \mu(t, r_t)dt + \sigma(t, r_t)dW_t$ (with μ, σ continuous functions) is an affine term structure model if and only if μ and σ^2 are affine functions of r. In that case, $Z(t,T) = e^{A(t,T) - B(t,T)r_t}$ where $Z(t,T)$ are bond prices/discount factors with period $[t,T]$ and $A(t,T)$, $B(t,T)$ satisfy the weakly coupled system of ordinary differential equations:

$$\begin{cases} \partial_t B(t,T) = -\alpha(t)B(t,T) + \frac{1}{2}\gamma(t)B^2(t,T) - 1 \\ B(T,T) = 0 \end{cases}$$

$$\begin{cases} \partial_t A(t,T) = \beta(t)B(t,T) - \frac{1}{2}\delta(t)B^2(t,T) - 1 \\ A(T,T) = 0 \end{cases}$$

where $\alpha(t), \gamma(t), \beta(t), \delta(t)$ are deterministic functions of t

This theorem applies to the Vasicek (1977) model where

$$A(t,T) = (\frac{\sigma_{vas}^2}{2a^2} - b)[(T-t) - B(t,T)] - \frac{\sigma_{vas}^2}{4a}B^2(t,T)$$

$$B(t,T) = \frac{1}{a}(1 - e^{a(T-t)})$$

Therefore, $Z(0,T) = e^{A(0,T) - B(0,T)r_0}$ is used to initialise a yield curve with $T \in [0, \frac{9}{12}]$ given that 6x9 caplets are being priced. This time horizon is discretised into a grid with 1-day increments as daily re-balancing is assumed in hedging the caplets.

Appendix B

The Libor Market Model

The Libor market model, also known as the BGM model after the authors who derived the approach Brace *et al.* (1997), can model actual traded instruments in the market. Typically, under this model, discrete-time forward interest rate dynamics are specified. The model is a lognormal market model where all rates specified are simple rates. Being a lognormal model ensures that interest rates can never go negative and makes the approximations used in deriving the Black (1976) model exact.

In its most general form, a single-factor Libor market model can be specified by the following dynamics:

$$dF(t;T_j,T_{j+1}) = F(t;T_j,T_{j+1})\mu_j(t)dt + F(t;T_j,T_{j+1})\sigma_j(t)dW_t \tag{B.1}$$

where W_t is a standard Brownian motion and $F(t;T_j,T_{j+1})$ is the fair forward rate at time t for maturity T_j with accrual period $[T_j,T_{j+1}]$. To simulate the Libor market model, the volatility is kept constant at $\sigma_j(t) = \sigma_{LMM} = 0.2$.

To initialise a set of forward rates at time t_0, the yield curve simulated as per appendix A is used where

$$F(t_0;T_j,T_{j+1}) = \frac{Z(t_0,T_j) - Z(t_0,T_{j+1})}{(T_{j+1}-T_j)Z(t_0,T_{j+1})}$$

Thereafter, forward rates with maturity and accrual period which match those of the 6x9 caplets are simulated for each day from t_00 until the expiry of the caplet. These forward rates are needed to calculate the value of the FRA (the hedge instrument) for each day until expiry of the caplets for each sample path. To simulate the model in discrete time with 1 day increments, the following discretisation of equation B.1 is used:

$$F(t_i;T_j,T_{j+1}) = F(t_{i-1};T_j,T_{j+1})\exp\left[(\mu_j(t_{i-1})-\frac{1}{2}\sigma^2_{LMM})(T_{j+1}-T_j)+\sigma_{LMM}\sqrt{T_{j+1}-T_j}Z_i\right] \tag{B.2}$$

where $j = i, i+1, \ldots, 120$ (i.e. we assume a month has 20 trading days), $Z_i \sim N(0,1)$ and

$$\mu_j(t_{i-1}) = \sum_{k=i}^{j} \frac{(T_{k+1}-T_k)F(t_k;T_{i-1},T_i)\sigma^2_{LMM}}{1+(T_{k+1}-T_k)F(t_k;T_{i-1},T_i)} \tag{B.3}$$

In order to reduce the discretisation error that arises due to the drift term, $\mu_j(t_{i-1})$ being state dependent, the predictor corrector method, proposed by Hunter *et al.*

(2001) is used. This method evolves the rates to the end of the period as normal, and then computes the terminal drift using the evolved rates. Using the same random variates used to estimate the terminal drift, the initial rates are then evolved using a drift computed as the average of the initial and terminal drift. The following 5 step procedure describes the method:

1. Initialise the forward rate curve using the initial yield curve.
2. Compute $\tilde{F}(t_i; T_j, T_{j+1}) = F(t_i; T_j, T_{j+1})$ for $j = i, i+1, \ldots, 120$ using equation B.2 where we replace $\mu_j(t_{i-1})$ with $\mu_j^{init}(t_{i-1})$ which is also given by equation B.3.
3. Using the intermediate values $\tilde{F}(t_i; T_j, T_{j+1})$, compute

$$\mu_j^{term}(t_{i-1}) = \sum_{k=i}^{j} \frac{(T_{k+1} - T_k)\tilde{F}(t_k; T_{i-1}, T_i)\sigma_{LMM}^2}{1 + (T_{k+1} - T_k)\tilde{F}(t_k; T_{i-1}, T_i)}$$

4. Compute the new rates

$$\bar{F}(t_i; T_j, T_{j+1}) = \bar{F}(t_{i-1}; T_j, T_{j+1}) \exp\big[\frac{1}{2}(\mu_j^{init}(t_{i-1}) + \mu_j^{term}(t_{i-1}) - \sigma_{LMM}^2)(T_{j+1} - T_j) + \sigma_{LMM}\sqrt{T_{j+1} - T_j}Z_i\big]$$

Appendix C

Statically vs Dynamically Hedging the ZCB Exposure

The histograms below show the hedge PnL distributions when statically and dynamically hedging the ZCB exposure using the simulated Libor Market Model data:

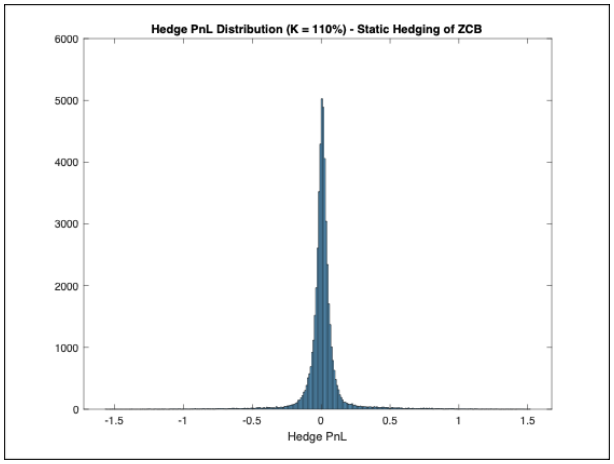

Fig. C.1: Hedge PnL Distribution with Static ZCB Hedge

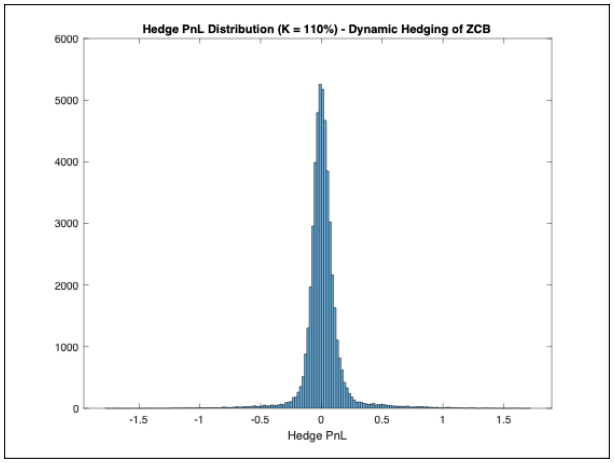

Fig. C.2: Hedge PnL Distribution with Dynamic ZCB Hedge

Appendix D

Pricing and Hedging Floorlets

The payoff of a floorlet is given by the equation below:

$$F(t_N) = A(K - R_x(t_{N-1}, t_N))^+ \tau,$$

where A is the nominal, τ is the year fraction between any successive t_i and t_{i+1} for $i \in \{1, 2, \ldots, N-1\}$, $R_x(t_{N-1}, t_N)$ is the underlying rate at time t_{N-1} and K is the strike rate.

Under the Black (1976) model; considering the same parameters from sections 2.1.1 and 2.1.2, the floorlet price at time t_0 is given by:

$$F(t_0) = AZ(t_0, t_N)\Big(K\Phi(-d_-) - R_{x,t_{N-1}}(t_0)\Phi(-d_+)\Big)\tau,$$

with

$$d_\pm = \frac{\ln\left(\frac{R_{x,t_{N-1}}(t_0)}{K}\right) \pm \frac{1}{2}\sigma^2(t_{N-1} - t_0)}{\sigma\sqrt{t_{N-1} - t_0}},$$

where $t_0 \leq t_{N-1}$, σ is the model volatility, $Z(t_0, t_N)$ is the risk-free discount factor applicable for $[t_0, t_N]$, $R_{x,t_{N-1}}(t_0)$ is the fair $[t_{N-1}, t_N]$ simple forward rate at t_0 and $\Phi(x)$ is the standard normal cumulative distribution function evaluated at x.

The following sensitivities at any time $s \in [t_0, t_{N-1})$ are defined and used to compute the required positions in the hedge portfolio for a floorlet:

$$\Delta_{ZCB_F,s} := \frac{\partial F(s)}{\partial Z(s, t_N)},$$

$$\Delta_{F,s} := \frac{\partial F(s)}{\partial R_{x,t_{N-1}}(s)}.$$

These sensitivities are calculated using the equations below:

$$\Delta_{ZCB_F,s} = A\big(K\Phi(-d_-^1) - R_{x,t_{N-1}}(s)\Phi(-d_+^1)\big)\tau,$$

$$\Delta_{F,s} = AZ(s, t_N)\left(K\phi(-d_-^1)\frac{-\partial d_-^1}{\partial R_{x,t_{N-1}}(s)} - \Phi(-d_+^1) - R_{x,t_{N-1}}(s)\phi(-d_+^1)\frac{-\partial d_+^1}{\partial R_{x,t_{N-1}}(s)}\right)\tau,$$

where $\phi(x)$ is the standard normal probability density function evaluated at x and

$$\frac{\partial d_+^1}{\partial R_{x,t_{N-1}}(s)} = \frac{1}{R_{x,t_{N-1}}(s)\sigma\sqrt{t_{N-1} - s}} = \frac{\partial d_-^1}{\partial R_{x,t_{N-1}}(s)}.$$

Thus, for floorlets the delta hedge ratio is defined as:

$$\delta_{F,s} := \frac{\Delta_{F,s}}{\Delta_{FRA,s}},$$

and the delta hedging routine is identical to that shown in section 2.1.2 with the final hedge PnL calculation from step 3 altered as follows:

3**. At time t_{N-1}, the hedge portfolio profit/loss (PnL) is calculated as:

$$PnL = b_{t_{N-1}-1}C(t_{N-1}-1, t_{N-1}) + \delta_{F,t_{N-1}-1}V_{FRA}(t_{N-1}) - Z(t_{N-1}, t_N)F(t_N) + \Delta_{ZCB_F,t_0}Z(t_{N-1}, t_N),$$

where $C(t_{N-1}-1, t_{N-1})$ is a risk-free overnight capitalisation factor over the period $[t_{N-1}-1, t_{N-1}]$.

For floorlets written on the prime lending rate, the West (2008) approach and above equations can be used. For a prime floorlet, the equivalent 3-month JIBAR floorlet payoff is given by:

$$F_P(t_N) = A\beta_1\left(\frac{K-\alpha_1}{\beta_1} - R_J(t_{N-1}, t_N)\right)^+\tau,$$

Thus, using the Black (1976) pricing equation, the t_0 price of a prime floorlet is given by:

$$F_P(t_0) = A\beta_1 Z(t_0, t_N)\left(\frac{K-\alpha_1}{\beta_1}\Phi(-d_-^1) - R_{J,t_{N-1}}(t_0)\Phi(-d_+^1)\right)\tau,$$

where

$$d_\pm^1 = \frac{\ln\left(\frac{\beta_1 R_{J,t_{N-1}}(t_0)}{K-\alpha_1}\right) \pm \frac{1}{2}\sigma^2(t_{N-1}-t_0)}{\sigma\sqrt{t_{N-1}-t_0}},$$

where $R_{J,t_{N-1}}(t_0)$ is the fair forward 3-month JIBAR rate at t_0 applicable over $[t_{N-1}, t_N]$.

For floorlets on the prime lending rate, the hedging methodology is identical to that of prime caplets but the relevant sensitivities as per equations (2.13) and (2.14) are tweaked due to the slightly different payoff of a floorlet. The tweaked sensitivities, where $s \in [t_0, t_{N-1})$, are shown below:

$$\Delta_{ZCB_F,s} := \frac{\partial F_P(s)}{\partial Z(s, t_N)} = A\beta_1\left(\frac{K-\alpha_1}{\beta_1}\Phi(-d_-^1) - R_{J,t_{N-1}}(s)\Phi(-d_+^1)\right)\tau,$$

$$\Delta_{F_P,s} := \frac{\partial F_P(s)}{\partial R_{J,t_{N-1}}(s)} = A\beta_1 Z(s, t_N)\left(\frac{K-\alpha_1}{\beta_1}\phi(-d_-^1)\frac{-\partial d_-^1}{\partial R_{J,t_{N-1}}(s)} - \Phi(-d_+^1) - R_{J,t_{N-1}}(s)\phi(-d_+^1)\frac{-\partial d_+^1}{\partial R_{J,t_{N-1}}(s)}\right)\tau,$$

where $\phi(x)$ is the standard normal probability density function evaluated at x and

$$\frac{\partial d_+^1}{\partial R_{J,t_{N-1}}(s)} = \frac{1}{R_{J,t_{N-1}}(s)\sigma\sqrt{t_{N-1}-s}} = \frac{\partial d_-^1}{\partial R_{J,t_{N-1}}(s)}.$$

Due to these tweaks, the delta hedge ratio for prime floorlets, $\delta_{F,s}$, becomes:

$$\delta_{F,s} := \frac{\Delta_{F_P,s}}{\Delta_{FRA,s}},$$

which is used in the delta hedging routine.

Appendix E

Quadratic Risk Minimisation

Local Risk-Minimisation

Referring back to the scenario introduced at the beginning of section 2.2, this approach insists that $V_T(\xi, \eta) = H$ at the cost of the requirement that the value process (or hedge portfolio) be self-financing. Local risk-minimisation entails finding a hedging strategy such that the conditional variance of the instantaneous cost process increments over time is minimised. More formally:

$$R_t(\xi, \eta) := \mathbb{E}[(C_T(\xi, \eta) - C_t(\xi, \eta))^2 | \mathcal{F}_t], \tag{E.1}$$

is minimised for $0 \leq t \leq T$. We define the cost process $(C_t)_{0 \leq t \leq T}$ as follows:

$$C_t(\xi, \eta) = V_t(\xi, \eta) - \int_0^t \xi_s dS_s. \tag{E.2}$$

This process describes the total costs incurred over $[0, t]$ due to price fluctuations in S which is traded. It can be shown that finding the hedging strategy which minimises $R_t(\xi, \eta)$ is equivalent to finding a decomposition of the derivative H known as the Föllmer-Schweizer (FS) decomposition:

$$H = H_0 + \int_0^T \xi_u^H dX_u + L_T^H, \tag{E.3}$$

where $H_0 \in \mathbb{R}$, $(\xi_t^H)_{0 \leq t \leq T}$ is a predictable process and L^H is a square-integrable martingale strongly orthogonal to M. Once we find this decomposition, our hedging strategy is given by:

$$\xi_t = \xi_t^H, \tag{E.4}$$

and

$$C_t(\xi, \eta) = H_0 + L_t^H, \tag{E.5}$$

for $0 \leq t \leq T$. We then use the value process to determine η_t:

$$V_t(\xi, \eta) = H_0 + \int_0^t \xi_u^H dX_u + L_t^H. \tag{E.6}$$

Mean-Variance Hedging

Mean-variance hedging deals with the global risk over the life of the derivative and insists that the value process is self-financing. Therefore, at maturity, a profit or loss is realised due to the difference between the value of the derivative, H and the value process. The optimal mean-variance hedging strategy is that which results in the profit and loss at maturity having the smallest variance, where we assume the mean profit and loss is zero. Thus, the mean-variance optimal strategy is the hedging strategy that minimises

$$\mathbb{E}\left[(H - V_0 - \int_0^T \xi_u dX_u)^2\right], \tag{E.7}$$

where V_0 is said to be the approximation price of H. The bank account holdings are then determined using $\eta_t = V_0 + \int_0^t \xi_u dX_u - \xi_t X_t$.

Appendix F

Adam Optimisation Algorithm

The Adam algorithm, as described in Kingma and Ba (2014), is given by the pseudo-code below:

Algorithm 1 Adam Optimisation

Require: α: Stepsize
Require: $\beta_1, \beta_2 \in [0, 1)$: Exponential decay rates for the moment estimates
Require: $f(\theta)$: Cost function with parameters θ
Require: θ_0: Initial parameter vector
  $m_0 \leftarrow 0$ (Initialise 1st moment vector)
  $v_0 \leftarrow 0$ (Initialise 2nd moment vector)
  $t \leftarrow 0$ (Initialise timestep)
  while θ_t not converged **do**
    $t \leftarrow t + 1$
    $g_t \leftarrow \nabla_\theta f_t(\theta_{t-1})$ (Get gradients w.r.t. cost function at timestep t)
    $m_t \leftarrow \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t$ (Update biased first moment estimate)
    $v_t \leftarrow \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2$ (Update biased second raw moment estimate)
    $\hat{m}_t \leftarrow m_t / (1 - \beta_1^t)$ (Compute bias-corrected first moment estimate)
    $\hat{v}_t \leftarrow v_t / (1 - \beta_2^t)$ (Compute bias-corrected second raw moment estimate)
    $\theta_t \leftarrow \theta_{t-1} - \alpha \cdot \hat{m}_t / (\sqrt{\hat{v}} + \epsilon)$ (Update parameters)
  end while
  return θ_t (Resulting parameters)

As suggested by Kingma and Ba (2014), the default parameters $\alpha = 0.001, \beta_1 = 0.9, \beta_2 = 0.999$ and $\epsilon = 10^{-8}$ are used in the algorithm.