Break-Even Volatility for Caps, Floors and Swaptions

Wade Cresswell

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 31, 2019

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 to the University of Cape Town. It has not before been submitted for any degree or examination.

Signed by candidate

Wade Cresswell

August 31, 2019

Abstract

This dissertation investigates break-even volatility in the context of the South African interest rate market. Introduced by Dupire (2006), break-even volatility is a retrospective measure defined as the volatility that ensures the profit or loss from a delta hedged option position is zero. Break-even volatility sheds light on the inner structure of the market and is a promising investigatory tool.

Insurance houses in South Africa are interested in modelling long-dated interest rate derivatives embedded within their liabilities. In pursuit of this goal, some are currently calibrating the Lognormal Forward-LIBOR Market Model to market prices. They rarely directly trade in said derivatives, but merely delta hedge their risk daily. In this case, break-even volatility surfaces become more relevant than recovering market prices (which incorporate the banks risk premium and profit margin) as it should better represent the historical cost of replicating the option under consideration. This dissertation ultimately assesses the use of the Lognormal Forward-LIBOR Market Model in the South African interest rate market using break-even volatility.

It is found that several caps and swaptions are trading at volatilities that differ significantly from their break-even volatility estimates. Furthermore, through an investigation into the calibration of the Lognormal Forward-LIBOR Market Model to break-even volatilities, an argument that the underlying dynamics of the model are incompatible with that of the South African interest rate market is developed.

Acknowledgements

I would firstly like to thank my supervisor, Obeid Mahomed. His support and guidance through what has been a chaotic period was invaluable.

I would also like to acknowledge the White House digs members Henry, Maurice, Hugh, Harambe and Tyrone for their constant encouragement, patience and understanding.

Finally, to my parents, I must express my profound gratitude for providing unwavering support. Without them, none of this would have ever been possible. This dissertation is dedicated to them. Thank you.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In recent years, South Africa's derivatives market has continued to grow and develop. South Africa has the second largest economy on the continent and has benefited from a strong supervisory framework. Primary financial markets exist for bonds and equity while derivative financial markets exist for swaps, futures and options written on various underlying asset classes. The interbank interest rate market is a financial market in which participants are able to trade vanilla funding products as well as interest rate derivatives.

Within the South African interbank interest rate market, the 3-month Johannesburg Interbank Agreed Rate (JIBAR) is the main reference rate used for the floating interest rate component of derivatives such as Caps, Floors and Swaptions. As noted in Mahomed (2015), JIBAR rates represent the average mid-deposit rate for specific investment horizons. They are simple rates using the ACT/365 day-count convention and are rounded to three decimal places. The London Interbank Offered Rate (LIBOR) is the United States equivalent of the JIBAR.

This dissertation investigates break-even volatility in the context of the South African interest rate market.

1.1 Break-Even Volatility

If one is to sell an option for a premium (quoted as some Black-Scholes implied volatility σ), and delta hedge this position with reference to the same volatility σ, the resultant profit or loss (P&L) can be expressed in terms of said volatility σ.

Introduced by Dupire (2006), break-even volatility (BEV) is a retrospective measure defined as the volatility that ensures the profit or loss from a delta hedged option position is zero. Thus, BEV should represent the historical cost of replicating said option. This volatility may not be unique, but there always exists a unique strictly positive solution. Interestingly, a single historical price path yields an entire BEV skew. In the context of interest rate instruments, a 'single historical price path' refers to daily yield curve information. Importantly, BEV being a historical measure precludes it from making inference about implied volatility. BEV is a largely practitioner driven idea and as such the academic literature available is sparse. However, Dupire (2006) notes that BEV sheds light on the inner structure of the market and is a promising investigation tool. In some markets, traders and market makers have little idea about how implied volatilities should differ across strike and maturity. In this case, BEV could provide meaningful insight. In situations where the underlying market is illiquid, the case for BEV strengthens.

There exist several variations of BEV. Firstly, given plentiful historical price information, one has to decide how to segment data into time windows on which to calculate BEV. This can be done by either considering overlapping increments, non-overlapping increments or a statistical bootstrapping technique whereby samples of overlapping time windows are taken with replacement. For example, if one has 4020 days of market data with which to compute the BEV of an option with a term of 30 days - the data could be segmented into 3990 overlapping time windows or 133 non-overlapping time windows. This choice will affect the statistical properties of the resultant BEV's and will be discussed further. Dupire (2006) addresses the issue of time window aggregation. In the first approach, BEV is found by averaging the volatilities that neutralise the P&L in each individual time window. In the second, BEV is found as the single volatility that neutralises the average P&L of all the time windows. He notes that the second approach seems to yield smoother results. However, this aggregation gives little indication of the spread of BEVs that exist in the sample set. A useful addition to this measure is a confidence level estimate of BEV. This would ensure that practitioners making use of BEV could be a certain level of confident that, based on past events over a specified time horizon, a transaction would be profitable.

Secondly, one can either calculate BEVs based on relative or absolute strike levels. The choice is dependent on beliefs surrounding the dependency of volatility behaviour. If one calculates BEVs across absolute strikes, there may be a complication in that the current relevant strikes surround the current price level. In this case, the sample size (containing common absolute strikes) will be diminished. It is important to maximise the sample size in BEV calculations in order to get smoother, more reliable results. However, in the case of the interest rate market in South Africa, absolute strikes remain constant throughout the entire sample set considered.

Thirdly, one could also consider weighting time windows by date in an attempt to capture the importance of current market conditions. Alternatively, one could weight according to price level.

Lastly, this measure can also be altered to include the effects of market frictions, trading costs, profit margins and so forth.

Insurance houses in South Africa are interested in modelling long-dated interest rate derivatives embedded within their liabilities. In pursuit of this goal, some are currently calibrating the Lognormal Forward-LIBOR Market Model to market prices. They rarely directly trade in said derivatives, but merely delta hedge their risk daily. In this case, BEV surfaces become more relevant than recovering market prices (which incorporate the bank's risk premium and profit margin) as it should better represent the historical cost (fair value) of replicating the option under consideration. This dissertation ultimately assesses the use of the Lognormal Forward-LIBOR Market Model in the South African interest rate market using BEV.

Chapter 2 discusses various hedging techniques for caps, floors and swaptions. This includes a delta and rho hedging technique, as well as a proposed modification to the traditional delta. Using simulated data from the Vasicek and G2++ short-rate models, the modified delta is found to perform better than the traditional and is used as the technique of choice throughout the rest of the dissertation.

Chapter 3 investigates BEV with simulated market data. At-the-money BEV estimates are generated for a particular caplet and swaption over several sample sizes and time window segmentation methods. Although the results differ minimally between segmentation methods, it is argued that non-overlapping segments are preferred for their statistical properties. A single historical price path yields a well structured BEV skew with both Vasicek and G2++ simulated data.

Chapter 4 assesses BEV within the South African interest rate market. It is found that several caps and swaptions are trading at volatilities that differ significantly from their BEV estimates. This is indicative of an inefficient market.

Lastly, Chapter 5 calibrates the Lognormal Forward-LIBOR Market Model with South African BEVs using co-terminal swaptions. An adequate model fit is attainable, but only when the instantaneous correlation matrix is allowed to become nonsensical. Using this, it is argued that the underlying dynamics of the South African interest rate market are incompatible with those underlying the Lognormal Forward-LIBOR Market Model.

Chapter 2

Hedging of Caps, Floors and Swaptions

Hedging is a risk management technique used to offset potential losses and gains due to a particular risk exposure. Taking positions in derivatives such as caps, floors and swaptions leads to interest rate exposure, which may be hedged in various ways. A perfect hedge that removes the risk exposure completely is known as replication, which is the fundamental basis for option pricing theory.

2.1 Delta Hedging

The delta of a particular option is simply the first derivative of the valuation formula with respect to the underlying asset. In this way, it represents a first-order sensitivity of the option value to price movements of the underlying asset.

Thus, BEV is model specific. This dissertation will consider Black's model (1976) for Caps, Floors and Swaptions. Black (1976) introduced this model in order to extend the Black-Scholes (1973) model for the case of commodity futures. Interestingly, this model became the standard option pricing model for vanilla interest rate derivatives before the rigorous mathematical theory to justify its use had been developed.

The payoff of standard European caplets and floorlets, initiated at time t with nominal value P, strike K, expiring at time $t_m > t$, written on the future spot simple floating rate $F(t_m, t_n)$ (JIBAR in this specific case), with $t_m < t_n$ is given by:

$$\text{Payoff} = \frac{P \max\{\eta(F(t_m, t_m, t_n) - K), 0\}\tau}{1 + F(t_m, t_m, t_n)\tau}$$

while the corresponding Black 1976 formula for pricing standard European caplets and floorlets is given by:

$$V(t) = P\tau Z(t, t_n)\eta[F(t; t_m, t_n)\Phi(\eta d_1) - K\Phi(\eta d_2)] \tag{2.1}$$

with:

$$
\begin{aligned}
\tau &= t_n - t_m \\
F(t; t_m, t_n) &= \frac{1}{\tau}\left(\frac{\exp\{r(t, t_n)(t_n - t)\}}{\exp\{r(t, t_m)(t_m - t)\}} - 1\right) \\
Z(t, t_n) &= \exp\{-r(t, t_n)(t_n - t)\} \\
d_1 &= \frac{\ln(F(t; t_m, t_n)/K)}{\sigma\sqrt{t_m - t}} + \frac{1}{2}\sigma\sqrt{t_m - t} \\
d_2 &= \frac{\ln(F(t; t_m, t_n)/K)}{\sigma\sqrt{t_m - t}} - \frac{1}{2}\sigma\sqrt{t_m - t}
\end{aligned}
$$

where η = 1 for caplets, -1 for floorlets. $r(t, \cdot)$ are Nominal Annual Compounded Continuously (NACC) spot rates. $\Phi(x)$ is the standard Gaussian cumulative distribution function at x. Caps are simple adjacent series of caplets and so the price of a cap would be calculated accordingly. In South Africa, caplets are settled in advance (like forward rate agreements) and caps are settled in arrears (like swaps).

Traditionally, the delta of Black's model is given by:

$$
\delta : \frac{\partial V(t)}{\partial F(t; t_m, t_n)} = P\tau Z(t, t_n)\eta\Phi(\eta d_1)
$$

Black's model incorrectly assumes that the forward rate is an asset independent of the other interest rate factors involved. In an attempt to rectify this problem for hedging purposes, it has become standard industry practice within South Africa to express the $Z(t, t_n)$ discount factor term in equation (2.1) as a function of said forward rate as:

$$
Z(t, t_n) = \frac{\exp\{-r(t, t_m)(t_m - t)\}}{1 + \tau F(t; t_m, t_n)}
$$

with the resultant modified delta given as:

$$
\delta : \frac{\partial V(t)}{\partial F(t; t_m, t_n)} = P\tau Z(t, t_n)\eta\Phi(\eta d_1) - \frac{\tau}{1 + \tau F(t; t_m, t_n)}V(t)
$$

where the derivation is given explicitly in Appendix A.1. The replication strategy for caps is simply a linear sum of replicating strategies for caplets. Thus, in evaluating hedging techniques, it is sufficient to focus on a particular caplet. The dynamics involved in hedging a floorlet naturally corresponds to that of a caplet and as such this dissertation will consider only the former. Musiela and Rutkowski (1997) describe how caplets can be replicated through a combination of appropriate positions in at-the-money (for simplicity) forward rate agreements (FRA's) and an initial investment of the caplet premium. Alternatively, caplets can be replicated solely in the cash market. This dissertation will consider daily delta hedging using the former approach.

A swaption is an option to enter into a swap at expiry and as such the underlying asset of a swaption is a forward-starting swap. The Black 1976 formula for standard European swaptions, initiated at time t with nominal value P, strike K, expiring at time $t_m > t$, written on the future spot swap rate $R(t_m, t_n)$ with swap maturity t_n, with is given by:

$$V(t) = P(\eta[R(t;t_m,t_n)\Phi(\eta d_1) - K\Phi(\eta d_2)])\left(\sum_{i=m+1}^{n} \tau_i Z(t,t_i)\right)$$

with:

$$\begin{aligned}\tau_i &= t_i - t_{i-1}\\ R(t;t_m,t_n) &= \frac{Z(t,t_m) - Z(t,t_n)}{\sum_{i=m+1}^{n} \tau_i Z(t,t_i)}\\ Z(t,t_i) &= \exp\{-r(t,t_i)(t_i - t)\}\\ d_1 &= \frac{\ln(R(t;t_m,t_n)/K)}{\sigma\sqrt{t_m - t}} + \frac{1}{2}\sigma\sqrt{t_m - t}\\ d_2 &= \frac{\ln(R(t;t_m,t_n)/K)}{\sigma\sqrt{t_m - t}} - \frac{1}{2}\sigma\sqrt{t_m - t}\end{aligned}$$

where η = 1 for payer swaptions and $\eta = -1$ for receiver swaptions. Again, $r(t, \cdot)$ are Nominal Annual Compounded Continuously (NACC) spot rates and $\Phi(x)$ is the standard Gaussian cumulative distribution function at x.

The delta is given by:

$$\delta : \frac{\partial V(t)}{\partial R(t;t_m,t_n)} = P\eta\Phi(\eta d_1)\left(\sum_{i=m+1}^{n} \tau_i Z(t,t_i)\right)$$

where the derivation is akin to that of Black's model for caplets as seen in Appendix A.1.

Swaptions are delta hedged in a similar manner to caplets through a combination of appropriate positions in at-the-money (for simplicity) swaps and an initial investment of the swaption premium.

2.2 Rho Hedging

It is useful to have a means of comparison for the above hedging technique. Another possible rectification of the flawed assumption that the forward rate is independent of the other interest rate factors involved is to consider the forward rate as it is defined. Thus, the Black 1976 formula becomes a function of two interest rate factors which may be hedged using two zero coupon bonds of appropriate tenor. This essentially hedges the first-order sensitivity to the interest rate risk associated with the particular caplet. This process becomes impractical for swaptions so only delta hedging will be considered.

It is useful to use discount factors as a proxy for interest rate risk as this clarifies the hedging process. These discount factors, corresponding to the price of zero coupon bonds of the same tenor, are defined as:

$$Z(t, t_i) = \exp\{-r(t, t_i)(t_i - t)\}$$

The first-order derivative of the Black 1976 formula for standard European caplets and floorlets with respect to the relevant discount factors is given by:

$$\frac{\partial V(t)}{\partial Z(t, t_m)} = P\eta\Phi(\eta d_1)$$

$$\frac{\partial V(t)}{\partial Z(t, t_n)} = \frac{V(t) - P\eta\Phi(\eta d_1)Z(t, t_m)}{Z(t, t_n)}$$

where the derivation is given explicitly in Appendix A.2, all notation remaining as in equation (2.1). Daily rho hedging operates in much the same way as delta hedging, but instead neutralises the interest rate risk by taking appropriate positions in t_n and t_m tenor zero coupon bonds.

2.3 Short-Rate Models for Simulation

To compare the hedging performance of the two aforementioned techniques, market data is simulated using Vasicek and Two-Additive-Factor Gaussian short-rate models. Closed form solutions for the price of put options on zero coupon bonds exist within the Vasicek and G2++ frameworks and thus through some mathematical manipulation, caplet premiums can be calculated.

2.3.1 Vasicek Model

Vasicek (1977) assumes the short-rate process $r(t)$ follows an Ornstein-Uhlenbeck process with constant coefficients under the risk-neutral measure:

$$dr(t) = k[\theta - r(t)]dt + \sigma dW(t), \quad r(0) = r_0,$$

where r_0, θ, k and σ are positive constants.

Jamshidian (1989) gives the closed form price for a European option written on a zero coupon bond under Vasicek's short-rate model. Brigo and Mercurio (2007) use this to show that the price of a standard European caplet, initiated at time t with nominal value P, strike K, expiring at time $t_m > t$, written on the future spot simple floating rate $F(t_m, t_n)$, with $t_m < t_n$ is given by:

$$\mathrm{Cpl}(t, t_m, t_n, P, K) = P'\mathrm{ZBP}(t, t_m, t_n, K') \tag{2.2}$$

where

$$\begin{aligned} K' &= \frac{1}{1 + K(t_n - t_m)} \\ P' &= P(1 + K(t_n - t_m)) \end{aligned}$$

with $\mathrm{ZBP}(t, t_m, t_n, K)$ denoting the time t price of a standard European put option with strike K, expiring at time $t_m > t$, written on a zero coupon bond maturing at t_n. This is given explicitly in appendix A.3.

2.3.2 Two-Additive-Factor Gaussian model

As detailed explicitly in Brigo and Mercurio (2007), the Two-Additive-Factor Gaussian (G2++) model assumes that the dynamics of the short rate process $r(t)$ under the risk-neutral measure is given by:

$$r(t) = x(t) + y(t) + \varphi(t), \quad r(0) = r_0,$$

where the processes $\{x(t) : t \geq 0\}$ and $\{y(t) : t \geq 0\}$ satisfy:

$$\begin{aligned} dx(t) &= -ax(t)dt + \sigma dW_1(t), \quad x(0) = 0, \\ dy(t) &= -by(t)dt + \eta dW_2(t), \quad y(0) = 0, \end{aligned}$$

where (W_1, W_2) is a two-dimensional Brownian motion with instantaneous correlation ρ as from:

$$d\left\langle W_1, W_2 \right\rangle_t = \rho \,\mathrm{d}t$$

with $-1 \leq \rho \leq 1$, where r_0, a, b, σ, η are positive constants. The function φ is deterministic and well defined in the time interval under consideration, with $\varphi(0) = r_0$. Caplet premiums are determined as in equation 2.2, with $\mathrm{ZBP}(t, t_m, t_n, X)$ in the G2++ model given by Brigo and Mercurio (2007) in Appendix A.4.

Under certain parameterisations, this model produces yield curves are not monotonic. The differing underlying dynamics of the Two-Additive-Factor Gaussian model and the Vasicek model ensures that the hedging performance is assessed under various market conditions.

2.4 Caplet Hedging with Simulated Market Data

To assess the performance of a particular hedging technique, it is important to examine the behaviour of the resultant P&L at expiry. A hedging technique is optimal if it neutralises the resultant P&L with minimal variation over numerous simulations.

Monte Carlo realisations of short-rate paths (and hence NACC yield curves) in the Vasicek and G2++ models are used as market data with which to conduct hedging simulations. Closed form caplet premiums allow the conversion of various input parameters to Black implied volatilities as described in Brigo and Mercurio (2007) which are used in the hedging process.

The Vasicek model considered is parameterised with $r_0 = 0.07$, $k = 0.15$, $\theta = 0.09$ and $\sigma = 0.02$. The standard European caplet to be hedged is initiated at time 0 with nominal value $P = 1000000$, strike $K = 0.070568$ (at-the-money), expiring in 9-months, written on the future spot simple 3-month floating rate $F(t_m, t_n)$.

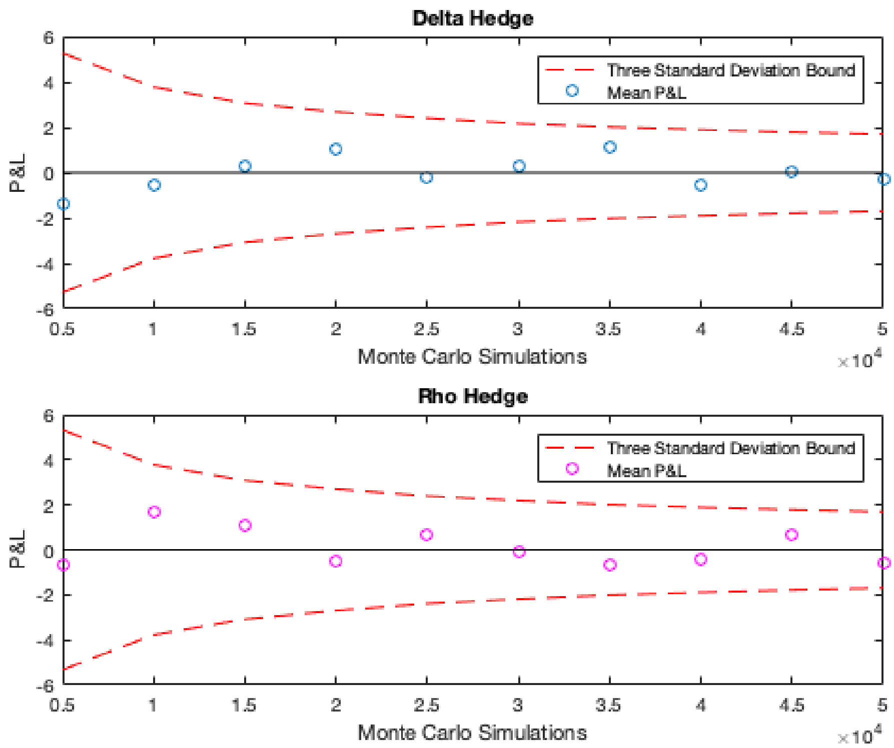

Fig. 2.1: Mean hedge portfolio P&L with Vasicek simulated market data.

As seen in Figure 2.1, there is little discernible difference in hedging performance between the two techniques. Both perform well with Vasicek generated market data as expected.

Although the differences in mean P&L are minor, the modified delta produces P&L values with a slightly lower variation than the traditional delta. Table 2.1 below shows the standard deviation of these values at the different sample sizes considered.

Sample Size	25000	30000	35000	40000	45000	50000
Modified Delta	126.67	125.28	125.88	126.34	126.02	126.30
Traditional Delta	129.82	128.40	128.89	129.44	129.19	129.48

Tab. 2.1: Standard deviation of P&L values with Vasicek generated data

.

The G2++ model considered is parameterised with $r_0 = 0.07$, $a = 0.5$, $\sigma = 0.005$, $b = 0.07$, $\eta = 0.01$ and $\rho = -0.001$. The standard European caplet to be hedged is initiated at time 0 with nominal value $P = 1000000$, strike $K = 070572$ (at-the-money), expiring in 9-months), written on the future spot simple 3-month floating rate $F(t_m, t_n)$.

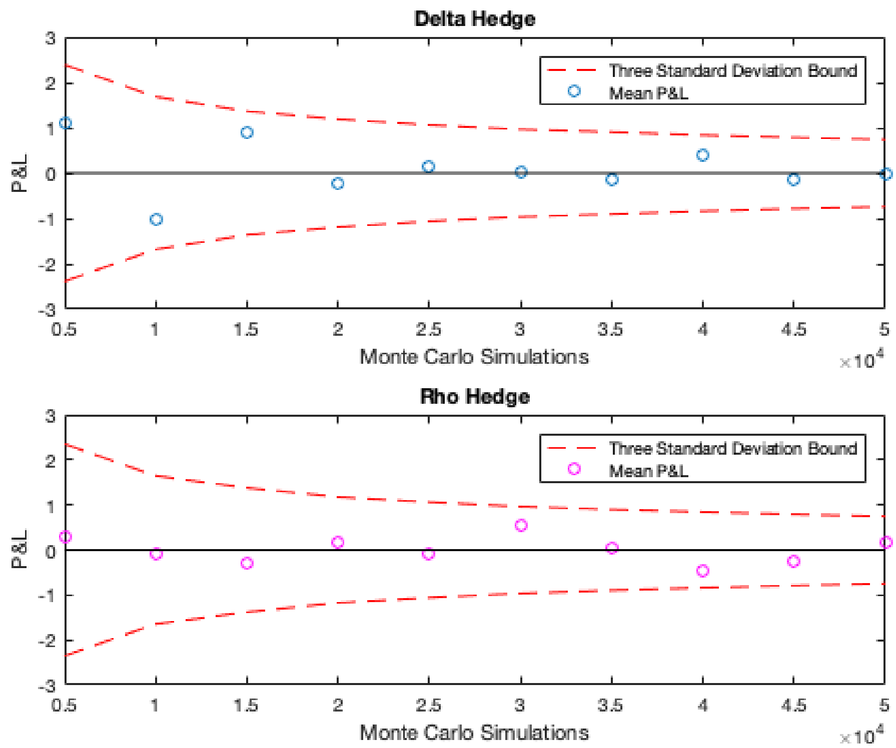

Fig. 2.2: Mean hedge portfolio P&L with G2++ simulated market data.

As seen in Figure 2.2, it is again the case that both techniques perform similarly. Thus, when calculating BEV, the choice of hedging technique between those considered is a matter of preference. As such, only Delta hedging will be considered. The hedging performance does not deteriorate with G2++ generated market data. Interestingly, the variance of the resultant P&L is less than that when using Vasicek simulated data.

As with the Vasicek generated data, the traditional delta technique produces P&L values that have a greater variance than that of the modified delta. This difference is more pronounced with G2++ simulated data, possibly due to the nature of the resultant yield curves. Table 2.2 below shows the standard deviation of these values at the different sample sizes considered.

Sample Size	25000	30000	35000	40000	45000	50000
Modified Delta	55.94	55.67	56.40	55.91	55.62	55.23
Traditional Delta	188.40	182.31	179.50	188.87	183.35	182.44

Tab. 2.2: Standard deviation of P&L values with G2++ generated data.

2.5 Swaption Hedging with Simulated Market Data

The standard European at-the-money (varies per model) Swaption to be hedged is initiated at time 0 with nominal value $P = 1000000$, expiring in 6 months, written on a 1 year interest rate swap.

Figure 2.3 depicts the delta hedging performance with data simulated with the G2++ and Vasicek models, parameterised as before. Interestingly, not unlike the with the caplet, the resultant variance of the resultant P&L is less for the G2++ simulated data than that from the Vasicek model.

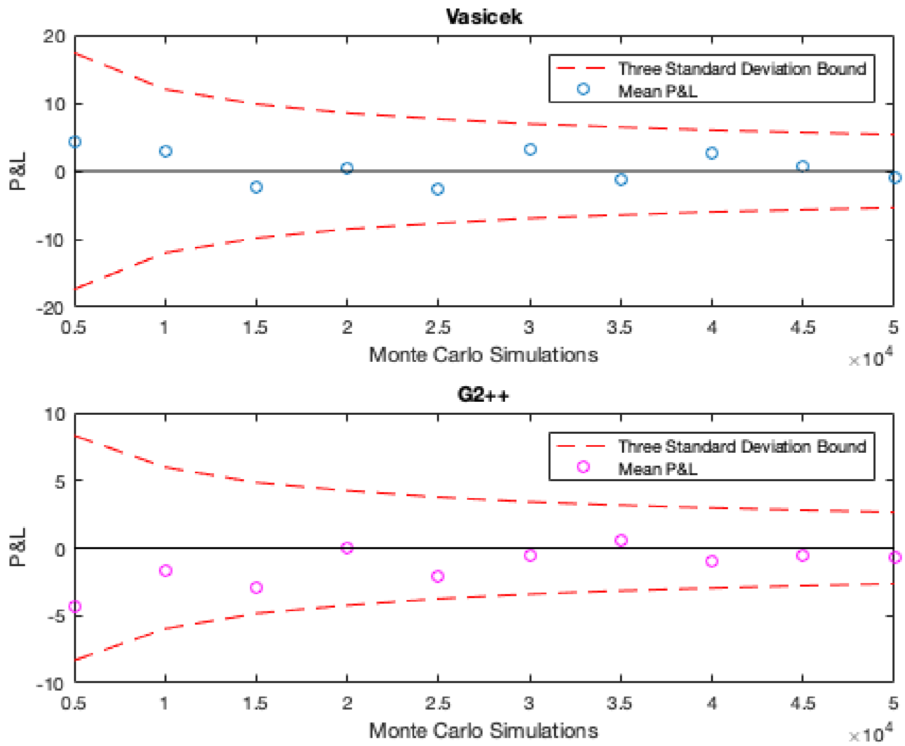

Fig. 2.3: Mean hedge portfolio P&L with G2++ and Vasicek simulated market data.

Chapter 3

Break-Even Volatility with Simulated Market Data

This dissertation will consider BEV whereby time windows are unweighted and the effects of market frictions, trading costs and profit margins are neglected. Furthermore, instead of only aggregating the time windows by means of averaging as discussed in section 1.1, BEV percentiles will be superimposed as a useful indication of the spread of the underlying distribution. Overlapping increments, non-overlapping increments as well as a statistical bootstrapping technique are investigated as methods of segmenting the data available to calculate BEV.

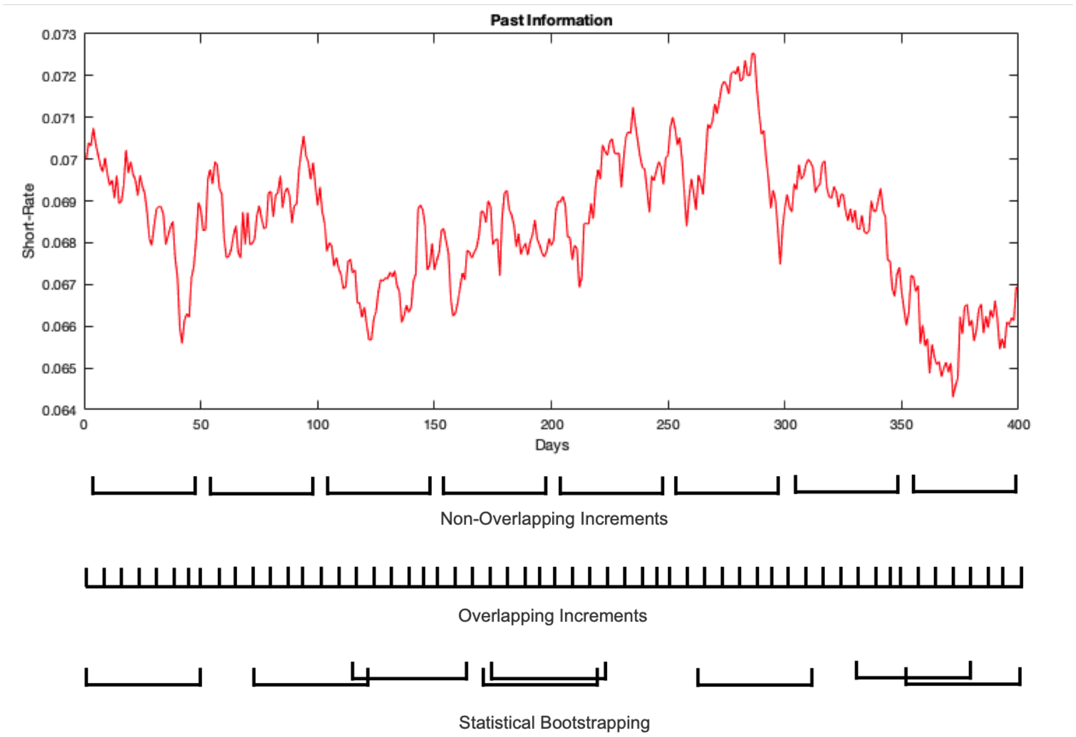

Fig. 3.1: Illustration of segmentation methods.

Figure 3.1 above visually illustrates the three segmentation methods considered.

With non-overlapping increments, the time windows used for BEV calculation are chosen such that the gap between these windows are equal. The gap is dependant on the sample size, tenor of the instrument and period of consideration. For a full non-overlapping increment BEV calculation, there will be no such gap and the time windows will be consecutive over the period of consideration. With overlapping increments, the time windows are chosen such that the overlapping section of each window is of equal size throughout the sample set. This is again dependant on sample size, tenor of instrument and period of consideration. A full overlapping BEV calculation would consist of BEV calculated on a daily basis such that each window overlaps completely with the next, not including the initial day. The statistical bootstrapping technique randomly (from a uniform distribution) selects the sample size of time windows from the period of consideration, allowing these time windows to overlap.

It must be noted that the simulation and delta hedging processes are computationally cumbersome and so it would be impractical to use a standard root finding algorithm that iterates through differing volatilities to calculate BEV. Instead, a grid containing the P&L values corresponding to particular volatilities and strikes will be generated for each time window. The volatility step sizes are chosen to minimise inaccuracy while considering computational practicality. BEV will be calculated using linear interpolation between the volatilities that bridge the P&L either side of zero for each strike.

3.1 Simulated Market Data

Two particular sets of market data are generated to investigate the various BEV methodologies, each 500 years in length. That is, a single price path is generated and fixed.

- Set A:
  - Model: Vasicek.
  - Parameters: $r_0 = 0.07$, $k = 0.15$, $\theta = 0.09$, and $\sigma = 0.02$.
- Set B:
  - Model: G2++.
  - Parameters: $r_0 = 0.07$, $a = 0.5$, $\sigma = 0.005$, $b = 0.07$, $\eta = 0.01$ and $\rho = -0.001$.

As a means of comparison, BEV is estimated using Monte Carlo simulations of market data from the current level of market rates (using the terminal short rate from the simulated price path as the initial short rate). The number of simulations correspond to the sample size used for each segmentation method.

The non-overlapping and overlapping time windows used to calculate BEV are equally spaced over the 500 year periods, regardless of sample size.

3.2 Cap Results

The standard European at-the-money caplet to be hedged is initiated at time 0 with nominal value $P = 1000000$, expiring in 9-months, written on the future spot simple 3-month floating rate $F(t_m, t_n)$. It must be noted that the Monte Carlo BEV's should not be seen as a closed form solution, but merely an indication of the current market dynamics. As the time windows are evenly spaced over the 500 year period, BEV is an average indication of the market dynamics over the entire period.

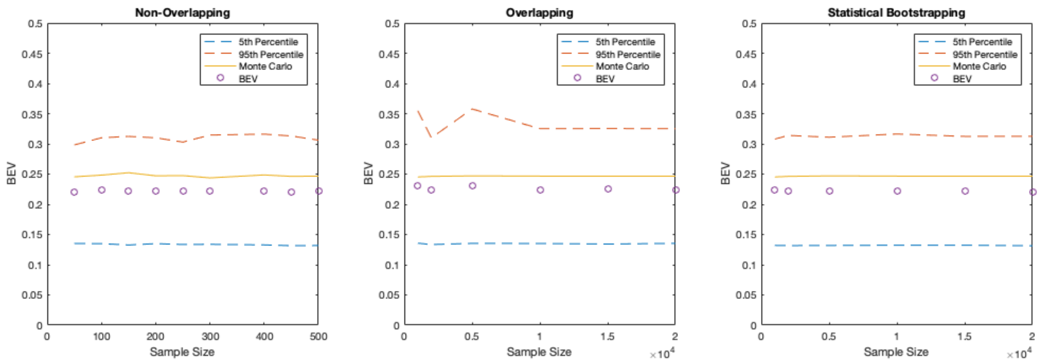

Fig. 3.2: At-The-Money Caplet BEV: Data Set A.

Figure 3.2 compares BEV using non-overlapping time windows, overlapping time windows and statistical bootstrapping over various sample sizes. All performed well with Vasicek simulated data. Calculating BEV's using past price information in this way captures the empirical dependencies between the instruments risk factors (interest rates of varying tenor) in a natural and dynamic way. However, using overlapping data will distort the tails of the resultant distribution in such a way that they become too blunt below a certain quantile. For example, this could result in the 0.1% BEV being of the same magnitude as the 1% BEV when this is not the case in reality. Although minor, this is evident in the 95$^{\text{th}}$ percentiles of Figure 3.2. Thus, when can be avoided, non-overlapping samples are preferred for their statistical properties.

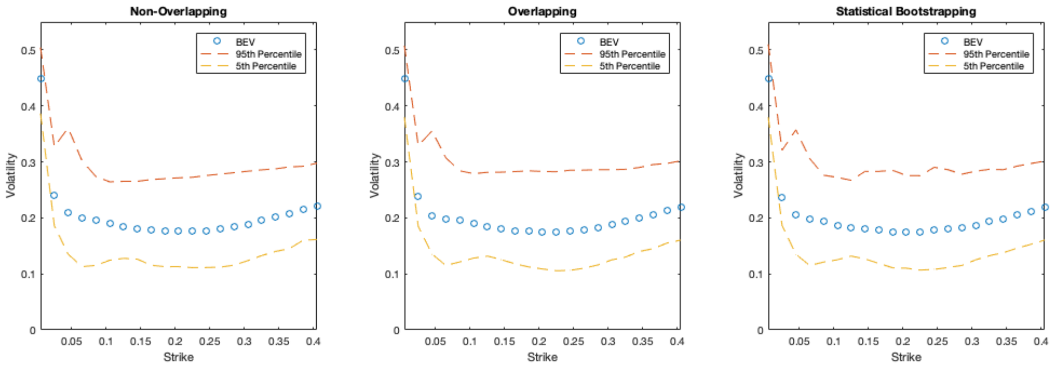

Fig. 3.3: Caplet BEV Skew: Data Set A.

Figure 3.3 conveys a powerful property of BEV in that a single price path yields an entire BEV skew. The maximum sample size from Figure 3.2 is used to generate the BEV skew. The large sample size could explain the similarity between the differing methods.

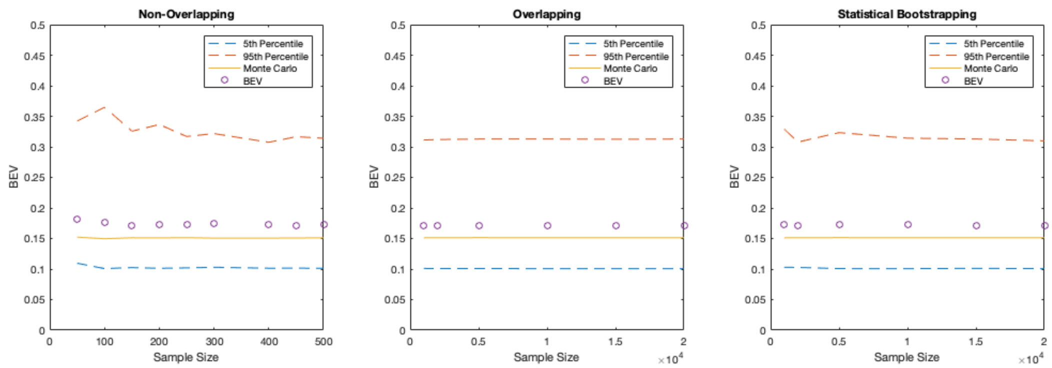

Fig. 3.4: At-The-Money Caplet BEV: Data Set B.

Figure 3.4 depicts BEV with G2++ simulated data. Again, all methodologies perform well even with the non-monotonic yield curve behaviour.

Figure 3.5 shows the BEV skew retaining its structure, performing well with G2++ simulated data.

3.3 Swaption Results

The standard European at-the-money swaption considered is initiated at time 0 with nominal value $P = 1000000$, expiring in 6 months, written on a 1 year interest rate swap.

Figures 3.6 and 3.7 convey the same story as with the caps, with BEV perform-

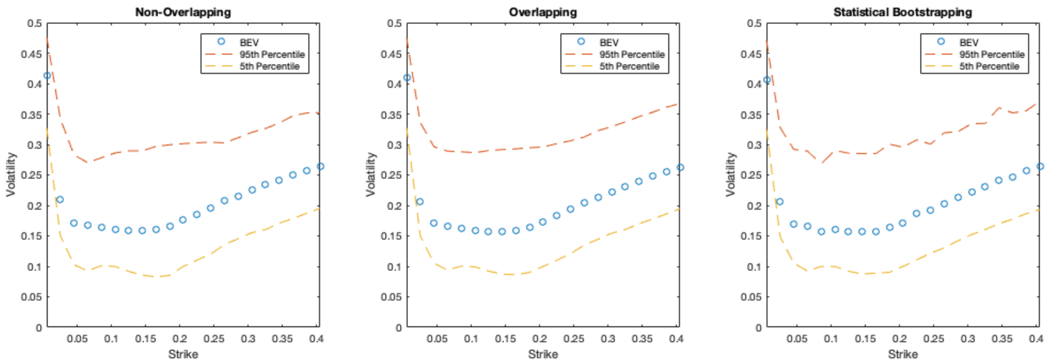

Fig. 3.5: Caplet BEV Skew: Data Set B.

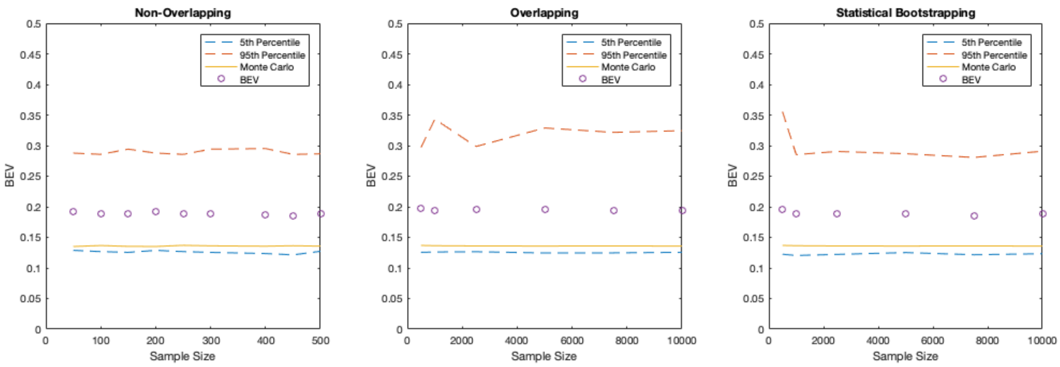

Fig. 3.6: At-The-Money Swaption BEV: Data Set A.

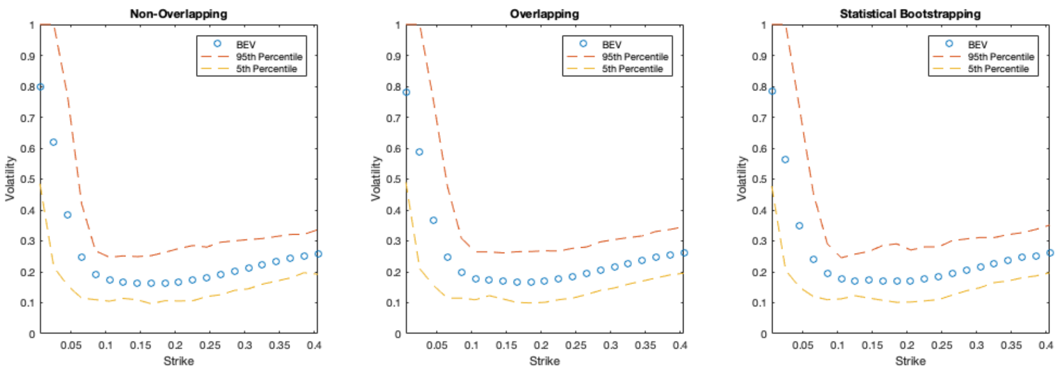

Fig. 3.7: Swaption BEV Skew: Data Set A.

ing well with Vasicek generated data. When the option is deep in-the-money, BEV seems to become nonsensical at the 95th percentile. However, other than that specific case BEV reproduces the skew.

Figure 3.8 depicts the ATM swaption BEV with G2++ simulated data. Figure 3.9 shows the BEV skew retaining its structure with swaptions, as with caps. Not

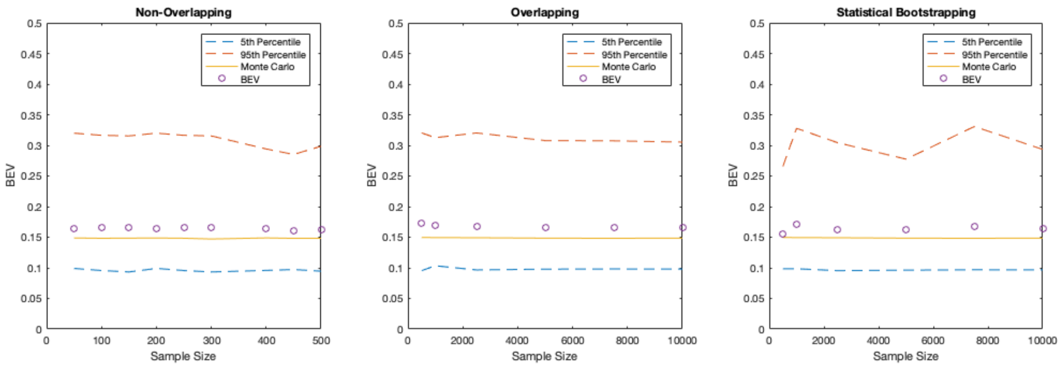

Fig. 3.8: At-The-Money Swaption BEV: Data Set B.

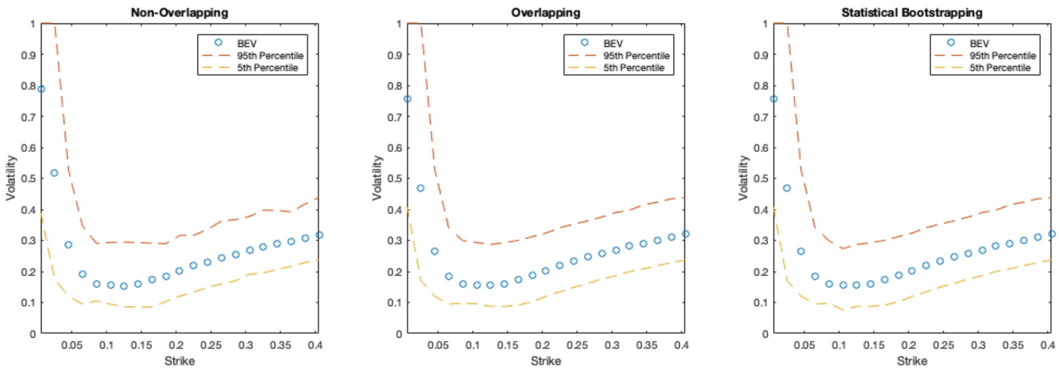

Fig. 3.9: Swaption BEV Skew: Data Set B.

unlike with the Vasicek generated data there exist anomalies when the option is deeply in-the-money, but overall BEV reproduces the skew. The differing non-monotonic yield curve behaviour does not negatively impact the swaption BEV estimates.

Chapter 4

Break-Even Volatility in the South African Interest Rate Market

Several caps, floors and swaptions trade daily on the Johannesburg Stock Exchange (JSE). The interest rate market has grown substantially over the last few years, with an increased variety of instruments available. As mentioned previously, the 3-month JIBAR is the main reference rate used in the interest rate market in South Africa.

As discussed in Chapter 3, calculating BEV with time windows that are non-overlapping is preferable to overlapping. However, sample size is absolutely integral to the quality of the BEV estimate. In the case of using limited market data, the need for a practical assumption overrides the need for technical consistency. Even though the data does overlap, including more information in the BEV calculation is worth the loss of desired statistical properties. This will result in slightly skewed results due to autocorrelation, but practically this is permissible. There are several ways to correct for this problem but they lie beyond the scope of this dissertation.

Furthermore, practitioners often desire estimates that reflect conditions in particular time frames. Thus, BEV will be calculated using overlapping time windows from the most recent 6-month and 2-year time intervals. Additionally, a BEV estimate including the maximum amount of data available for each particular instrument will be presented. Instrument schedules were calculated as per the JSE's specifications using the modified following day count convention.

4.1 Data

In order to calculate the various BEV's, daily nominal annual compounded continuously (NACC) yield curve data has been retrieved from the JSE. This data spans from the 2$^{\text{nd}}$ of January 2004 until the 30$^{\text{th}}$ of March 2018. In conjunction with the yield curves, daily volatility quotes of various caps, floors and swaptions have also been provided spanning the same dates. Caps and floors are quoted with term up to 10 years and strikes varying from 0.5% to 10%. Swaptions are quoted only at-the-money, with term and tenor up to 10 years.

4.1.1 Yield Curve Interpolation

The NACC yield curves have been interpolated as per the JSE's specifications, using an algorithm described in Du Preez (2011). A monotone preserving cubic Hermite interpolation scheme is applied to the function $r(t)t$, where $r(t)$ is the NACC rate at t. This algorithm was developed inter alia by Akima (1970), Fritsch and Carlson (1980), De Boor (1978) and Hyman (1983).

4.2 Caps

This section examines the BEV of 1 year, 2 year and 5 year caps quoted on the 30$^{\text{th}}$ March 2018.

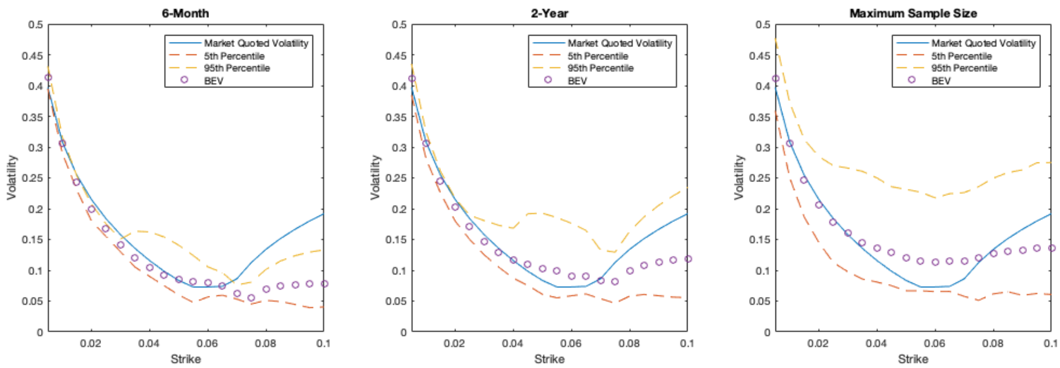

Fig. 4.1: 1-Year Cap: 30$^{\text{th}}$ March 2018.

As can be seen from Figure 4.1, the BEV is relatively well aligned with the volatilities quoted in the market. This is to be expected, as 1-year caps tend to be among the most liquid and would naturally attain this relationship. However, irregularities remain, particularly in out-the-money caps whereby the 95th percentile 6-month BEV lies below the market quotes. Thus, a trade whereby a market maker sells these particular options and merely delta hedges their risk daily would've been profitable with 95% confidence over the last 6 months. Although the confidence is not of the magnitude of 95%, this property holds over the last 2 years as well as over the maximum sample size of market data.

Figure 4.2 shows that the various BEV skews for the 2-year cap is less aligned with market quoted volatilities than that of the 1-year cap. The BEV structure indi-

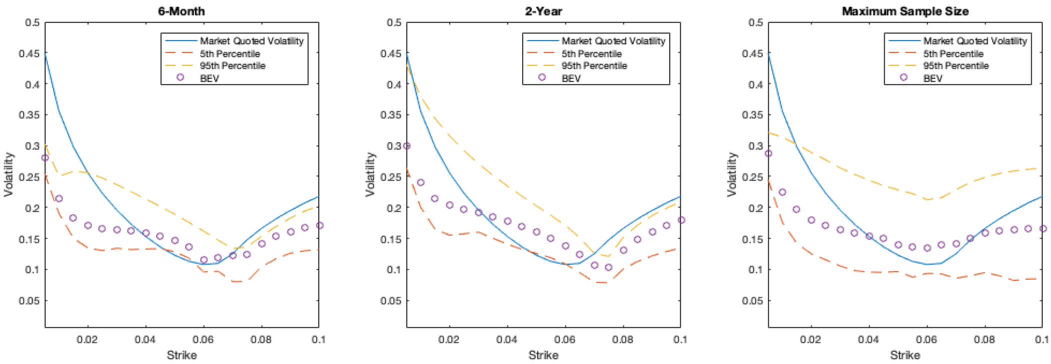

Fig. 4.2: 2-Year Cap: 30th March 2018.

cates that as a market maker who daily delta hedges their risk, one would prefer to sell deeply in-the-money and out-the money 2-year caps while purchasing at-the-money.

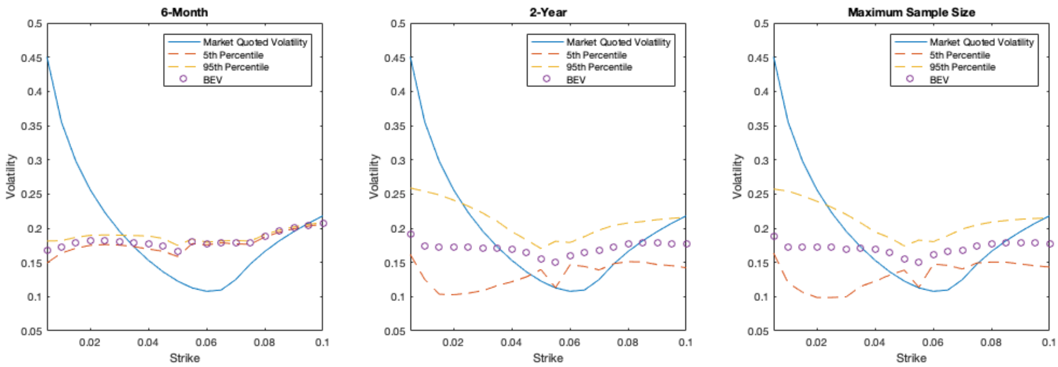

Fig. 4.3: 5-Year Cap: 30th March 2018.

Interestingly, Figure 4.3 shows a far flatter BEV skew than that of the 1-year and 2-year caps. Furthermore, the BEV estimates are less volatile. This is potentially due to the state of the interest rate market over the period used to calculate BEV and could indicate a possible need to re-evaluate the deeply in-the-money market quoted volatilities. The similarities seen between the 2-year and maximum sample size BEV can be attributed the fact that the longer term caps will naturally have a lower maximum sample size with which to calculate BEV. Thus, there is a large percentage overlap in the data used.

4.3 Swaptions

Figure 4.4 below shows the market quoted volatility being significantly greater than the BEV. This indicates that a large premium is being demanded by the market maker, especially considering the market dynamics over the last 6 months.

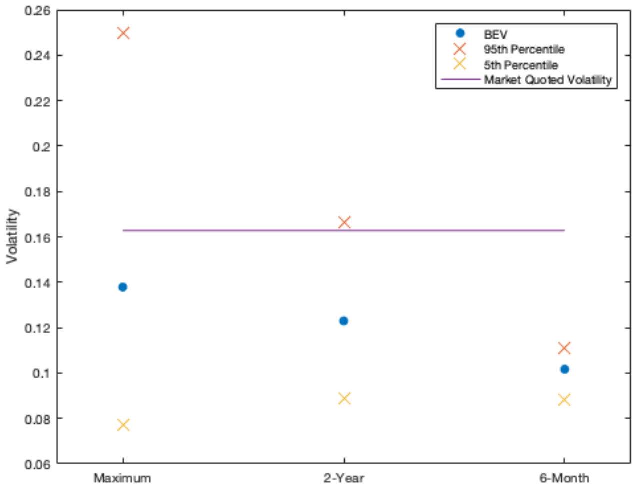

Fig. 4.4: 1Y×1Y ATM Swaption.

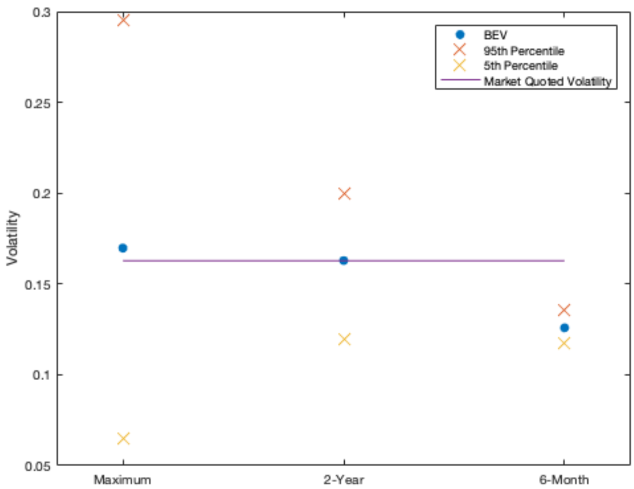

Fig. 4.5: 1Y×5Y ATM Swaption.

However, the 1Y×5Y ATM swaption shown in Figure 4.5 aligns better with the historical BEV's. There is a small discrepancy over the last 6 months, but not to the same magnitude as the 1Yx1Y ATM swaption.

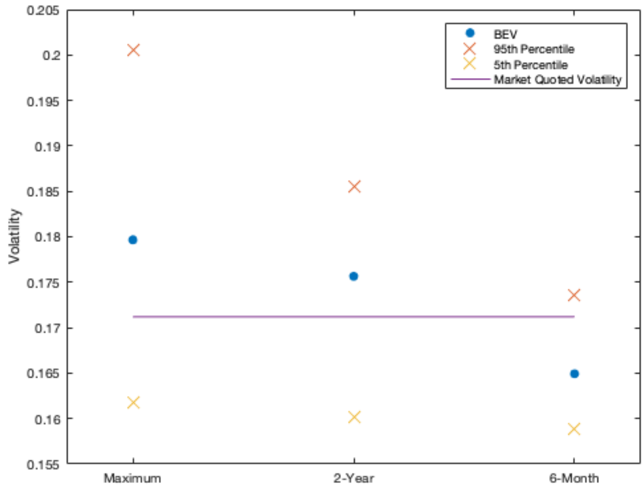

Fig. 4.6: 5Yx7Y ATM Swaption.

Interestingly, the 5Y×7Y ATM swaption BEV lies mostly above the market quoted volatility as seen in Figure 4.6. However, this could indicate that the market makers place emphasis on the most recent 6 months, where BEV lies below the market quoted volatility.

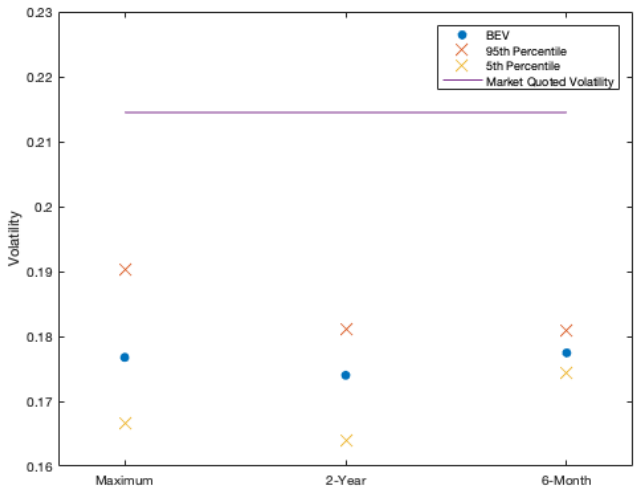

Fig. 4.7: 10Yx10Y ATM Swaption.

Figure 4.7 shows that the market quoted volatility of the largest expiry and tenor combination in the market lies significantly above the BEV estimates, regardless of the sample size. This is indicative of a mispriced long-dated swaption market, where the market maker demands a relatively large premium for the instrument. This strengthens the case for BEV for the insurance houses mentioned in Chapter 1.

Chapter 5

Non-Parametric Calibration of the LFMM with Break-Even Volatilities

As mentioned in Chapter 1, insurance houses in South Africa are interested in modelling long-dated interest rate derivatives embedded within their liabilities and are currently calibrating the Lognormal Forward-LIBOR Market Model (LFMM) to market prices.

One of the important strengths of BEV is the ability to synthesise a BEV estimate for any particular instrument given the historical market information. This allows for the optimal calibration of various models in cases where this would otherwise be impossible given the limited number of instruments traded in the relevant market. The purpose of this section is to calibrate the LFMM in a non-parametric fashion using BEV. This removes unnecessary noise as a result of various parameterisations and sheds light on the compatibility between the underlying dynamics of the model and market.

5.1 LFMM Structure

Let $Z(t, t_n)$ be the value at time t of a zero-coupon bond delivering $Z(t_n, t_n) = 1$ at time $t_n > t$. The simply-compounded forward-LIBOR (analogous to JIBAR) rate is defined as before as:

$$F(t, t_m, t_n) := \frac{1}{\tau}\Big(\frac{Z(t, t_m)}{Z(t, t_n)} - 1\Big), \quad 0 \leq t < t_m < t_n.$$

Consider a set $T_0 < ... < T_M < T_{M+1}$ of bond tenor dates, expressed in years from the current time t. The corresponding simple forward rate $F_i(t)$ of the LIBOR rate $L(T_i, T_{i+1})$ from T_i to T_{i+1} at some time $t \leq T_i$ is given by $F_i(t) := F(t, T_i, T_{i+1})$, with $F_i(T_i) = L(T_i, T_{i+1})$. Consider the probability measure $\mathbb{Q}^{i+1}$ associated with the numeraire $Z(t, T_{i+1})$. $\mathbb{Q}^{i+1}$ is called the forward adjusted measure for maturity T_{i+1}, or simply the T_{i+1} forward measure. Let τ_i be the year fraction of the period $(T_i; T_{i+1})$. By definition, we have $F_i(t)Z(t, T_{i+1}) := \frac{1}{\tau_i}(Z(t, T_i) - Z(t, T_{i+1}))$. It follows that $F_i(t)Z(t, T_{i+1})$ is the price of a tradable asset (difference between two zero-coupon bonds with a nominal amount of $\frac{1}{\tau_i}$). Hence, $F_i(t)$ is a martingale under $\mathbb{Q}^{i+1}$. The LFMM assumes the following driftless dynamics for the forward LIBOR rate $F_i(t)$ under $\mathbb{Q}^{i+1}$:

$$dF_i(t) = \sigma_i(t)F_i(t)\, dW_i^{i+1}(t), \quad t \leq T_k$$

where $\sigma_i(t)$ is the instantaneous volatility at time t of the forward LIBOR rate $F_i(t)$ and $W_i^{i+1}(t)$ is the i^{th} component of the M-dimensional Brownian motion $W^{i+1}(t)$ under $\mathbb{Q}^{i+1}$ with instantaneous covariance given by:

$$d\left\langle W_i^{i+1}, W_j^{i+1}\right\rangle_t = \rho_{i,j}\, \mathrm{d}t$$

where the matrix formed by elements $\rho_{i,j}$ is denoted by ρ.

Thus, the forward LIBOR rate $F_i(t)$ is lognormally distributed under $\mathbb{Q}^{i+1}$. However, the dynamics of $F_i(t)$ under a measure $\mathbb{Q}^{m+1}$ different from $\mathbb{Q}^{i+1}$ are not martingales and are given by:

$$\mathrm{d}F_i(t) = \mu_i(t)\, \mathrm{d}t + \sigma_i(t)F_i(t)\, \mathrm{d}W_i^{m+1}(t), \quad t \leq T_i$$

$$\text{where} \quad \mu_i(t) := \begin{cases} -\sigma_i(t)F_i(t) \sum\limits_{j=i+1}^{m} \frac{\rho_{i,j}\tau_j\sigma_j(t)F_j(t)}{1+\tau_jF_j(t)} & \text{for } i < m; \\ 0 & \text{for } i = m; \\ \sigma_i(t)F_i(t) \sum\limits_{j=m+1}^{i} \frac{\rho_{i,j}\tau_j\sigma_j(t)F_j(t)}{1+\tau_jF_j(t)} & \text{for } i > m, \end{cases}$$

Swaption Price Approximation

The forward swap rate $S_{\alpha,\beta}(t)$ is the rate of the fixed leg of an interest rate swap that ensures the contract is fair and thus has no costs to entry. The set of payment dates for this swap is $\mathcal{T} = \{T_{\alpha+1}, ..., T_\beta\}$. In the context of swaptions, T_α would be the expiry and $T_\beta - T_\alpha$ the tenor. With this we obtain:

$$S_{\alpha,\beta}(t) = \frac{Z(t, T_\alpha) - Z(t, T_\beta)}{\sum_{i=\alpha}^{\beta-1} \tau_i Z\big(t, T_{i+1}\big)}$$

Dividing both the numerator and denominator by $Z(t, T_\alpha)$ and noticing that $\frac{P(t,T_k)}{P(t,T_\alpha)} = \prod_{i=\alpha}^{k-1} \frac{P(t,T_{i+1})}{P(t,T_i)} = \prod_{i=\alpha}^{k-1} \frac{1}{1+\tau_iF_i(t)}$ for all $k > \alpha$, we can express the forward swap rate

$S_{\alpha,\beta}(t)$ in terms of the forward-LIBOR rates as:

$$S_{\alpha,\beta}(t) = \frac{1 - \prod_{j=\alpha}^{\beta-1} \frac{1}{1+\tau_j F_j(t)}}{\sum_{i=\alpha}^{\beta-1} \tau_i \prod_{j=\alpha}^{i} \frac{1}{1+\tau_j F_j(t)}}$$

Hence, the forward swap rate is not lognormally distributed under the LFMM and an approximation is necessary for pricing swaptions.

Rebonato (1999) notes that the swap rate can be seen as a linear combination of forward rates as:

$$S_{\alpha,\beta}(t) = \sum_{i=\alpha}^{\beta-1} w_i(t) F_i(t),$$

with weights defined by

$$w_i(t) = \frac{\tau_i \prod_{j=\alpha}^{i} \frac{1}{1+\tau_j F_j(t)}}{\sum_{k=\alpha}^{\beta-1} \tau_k \prod_{j=\alpha}^{k} \frac{1}{1+\tau_j F_j(t)}}.$$

Moreover, Rebonato (1999) makes the following simplifying assumptions:

- Each forward LIBOR rate $F_i(t)$ and its weight $w_i(t)$ are independent.
- $w_i(t) \approx w_i(0)$.
- $F_i(t) \approx F_i(0)$.

In combination, these assumptions result in the following approximation:

Approximation 5.1.1 (The Rebonato formula)**.** The squared Black swaption volatility is

$$\left(v_{S_{\alpha,\beta}}\right)^2 \approx \frac{1}{T_\alpha} \sum_{i,j=\alpha}^{\beta-1} \frac{w_i(0) w_j(0) F_i(0) F_j(0)}{S^2_{\alpha,\beta}(0)} \rho_{ij} \int_0^{T_\alpha} \sigma_i(t)\sigma_j(t)\, \mathrm{d}t.$$

Hull and White (2013) extended the Rebonato formula by using a first-order Taylor expansion to approximate the weight w_i.

Approximation 5.1.2 (The Hull-White formula)**.** The squared Black swaption volatility is

$$\left(v_{S_{\alpha,\beta}}\right)^2 \approx \frac{1}{T_\alpha} \sum_{h,j=\alpha}^{\beta-1} G_{h,j}(0) \rho_{hj} \int_0^{T_\alpha} \sigma_h(t)\sigma_j(t)\, \mathrm{d}t.$$

where

$$G_{h,j}(t) = \frac{\tilde{w}_h(t)\tilde{w}_j(t)F_h(t)F_j(t)}{S^2_{\alpha,\beta}(t)}.$$

and

$$\tilde{w}_h(t) = w_h(t) + \sum_{i=\alpha}^{\beta-1} F_i(t)\frac{\partial w_i(t)}{\partial F_h} \tag{5.1}$$

and

$$\frac{\partial w_i(t)}{\partial F_h} = \frac{w_i(t)\tau_h}{1+\tau_h F_h(t)}\left[\frac{\sum_{k=h}^{\beta-1}\tau_k \prod_{j=\alpha}^{k}\frac{1}{1+\tau_j F_j(t)}}{\sum_{k=\alpha}^{\beta-1}\tau_k \prod_{j=\alpha}^{k}\frac{1}{1+\tau_j F_j(t)}} - \mathbb{I}_{\{i\geq h\}}\right]$$

This dissertation considers approximations generated using the Hull-White formula.

Note on Implementation of Approximations

For calibration purposes it is necessary to generate, as efficiently as possible, Black swaption volatilities of swaptions with multiple expiries (α) and tenors ($\beta - \alpha$). Calculating these approximations independently for each specific α and β results in much of the same information being generated repetitively, ultimately rendering the calibration process inefficient. As discussed in Kleisinger-Yu *et al.* (2018), it is possible to separate much of the information common to all α and β combinations so as to only generate this once, which greatly improves the efficiency of generating multiple Black swaption volatilities. For example, when calculating the weights used in the Hull-White Formula given by 5.1, one would (only once, used for all combinations of α and β) create an $N \times M$ matrix with entries given by

$$\text{where} \quad d_{i,j} = \begin{cases} \tau \prod_{z=\alpha_i}^{j} \frac{1}{1+\tau_z F_z(t)} & \text{for } j > \alpha_i; \\ 0 & \text{otherwise;} \end{cases}$$

where $i = 1, ..., N$ and $j = \alpha_i, ..., M$ where N is total number of α_i's considered. From this, it is possible (through various selective summation schemes) to extract the information necessary to calculate the respective weights for each considered α and β combination.

5.2 Optimal Non-Parametric Calibration

The purpose of this section is to calibrate the LFMM to BEV's in a non-parametric fashion. This removes unnecessary noise resulting from differing parameterisations, allowing a clean analysis of the LFMM's fit to the South African interest rate market. An optimal non-parametric calibration is only possible due to the fact that, using BEV, one can synthesise any particular instrument using only past information. Rebonato (2005) provides an argument for the optimal calibration of the LFMM using co-terminal swaptions. This provides the basis of the methodology to follow. However, this dissertation does not consider the parameterisations proposed, but instead builds the correlation matrix using a novel 'hybrid bootstrapping' technique. BEV allows for the extraction of terminal correlations, which is not possible when using the historical correlations between forward LIBOR rates for calibration.

5.2.1 Methodology

Initially, the instantaneous volatility structure is fixed using 2-year caplet BEV's. As such, no cap stripping is necessary. An assumption that the instantaneous volatility of each forward rate depends only on the time until maturity is made, as suggested by Brigo and Mercurio (2007). Through this assumption, one can create Table 5.1 of instantaneous volatilities below:

Instantaneous Volatilities	Time $t \in (0; T_1]$	$(T_1; T_2]$	$(T_2; T_3]$	...	$(T_{M-1}; T_M]$
Forward Rate $F_1(t)$	η_1	Dead	Dead	...	Dead
$F_2(t)$	η_2	η_1	Dead	...	Dead
...	...	...	...	...	...
$F_M(t)$	η_M	η_{M-1}	η_{M-2}	...	η_1

Tab. 5.1: Matrix of Instantaneous Volatilities $\sigma_{i,j}(t)$.

The above η values allow for an exact fit of the instantaneous volatilities to the caplet volatilities. Given the evaluation method used to find the caplet volatilities, this instantaneous volatility structure constitutes a piece-wise constant instantaneous volatility assumption. Specifically, each $T_i - T_{i-1}$ is taken as 3 months and 40 forward LIBOR rates are considered.

Insurance houses require prudence when reporting the value of their embedded long-dated derivatives. Thus, in order to not understate their liabilities, the LFMM is also calibrated to 95^{th} percentile BEV's.

The instantaneous correlation matrix is built, element by element, using a 'hybrid bootstrapping' technique with co-terminal swaptions. The optimisation seeks to minimise each particular absolute relative error $\epsilon_{\alpha,\beta} := \frac{|\sigma^{\text{BEV}}_{\alpha,\beta} - \sigma^{\text{model}}_{\alpha,\beta}|}{\sigma^{\text{BEV}}_{\alpha,\beta}}$ where $\sigma^{\text{model}}_{\alpha,\beta}$ denotes the corresponding model-based swaption volatility. Firstly a swaption with $\alpha = 1$ and $\beta = 3$ is used to calibrate the $\rho_{1,2}$ and $\rho_{2,1}$ elements of the instantaneous correlation matrix ρ. Then, with the aforementioned elements remaining

α	β	Matrix Elements
1	3	$\rho_{1,2}$, $\rho_{2,1}$
2	4	$\rho_{2,3}$, $\rho_{3,2}$
1	4	$\rho_{1,3}$, $\rho_{3,1}$
3	5	$\rho_{3,4}$, $\rho_{4,3}$
2	5	$\rho_{2,4}$, $\rho_{4,2}$
1	5	$\rho_{1,4}$, $\rho_{4,1}$
...	...	...
1	41	$\rho_{1,40}$, $\rho_{40,1}$

Tab. 5.2: Co-Terminal Swaptions used for Calibration.

fixed, another swaption with $\alpha = 2$ and $\beta = 4$ is used to calibrate the elements $\rho_{2,3}$ and $\rho_{3,2}$. Finally, a co-terminal swaption with $\alpha = 1$ and $\beta = 4$ is used to calibrate the remaining elements $\rho_{1,3}$ and $\rho_{3,1}$ of the now 3×3 matrix ρ. The process continues in this fashion using consecutive co-terminal swaptions until the full 40x40 instantaneous correlation matrix is populated. The exact swaptions used (in order) are shown in Table 5.2.

It is necessary at each step to constrain the optimisation to ensure that the entire correlation matrix remains valid. This is done using the algorithm developed by Higham (2002), with which a MATLAB function is provided. For a matrix to be considered valid, the following properties must hold:

- $|\rho_{i,j}| \leq 1$ for all i, j,
- ρ positive semidefinite,
- $\rho_{i,i} = 1$ for all i.

Furthermore, correlation of forward LIBOR rates should exhibit the following desired properties:

- $i \to \rho_{i,j}$ increasing for all $i \geq j$ "decreasing along row,"
- $i \to \rho_{i+p,i}$ increasing for fixed p "increasing along sub-diagonals."

Certain parameterisations force the matrix to exhibit the desired properties. The above methodology does not, allowing for maximum flexibility in order to fit the market (BEV) forward LIBOR rates. Kleisinger-Yu *et al.* (2018) found that when calibrating the LFMM to market quoted volatilities in South Africa, using several desirable parameterisations of the correlation matrix all yielded average absolute relative errors in excess of 30%.

5.2.2 Results

2-year caplet BEV's used to fix the LFMM's instantaneous volatility structure have a term structure as shown below in figure 5.1:

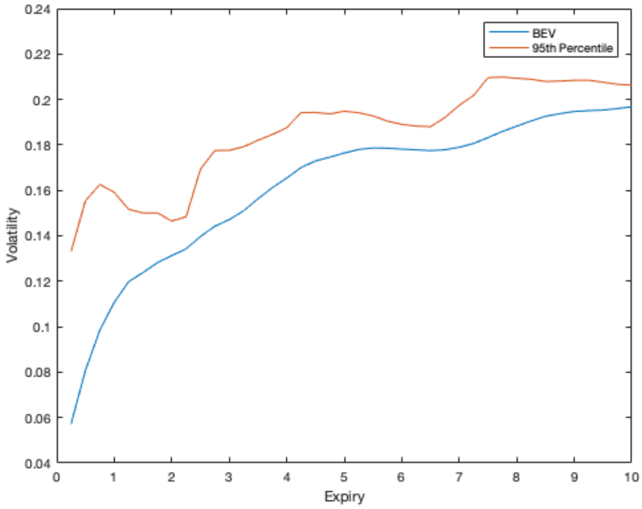

Fig. 5.1: Term Structure of Caplet BEVs.

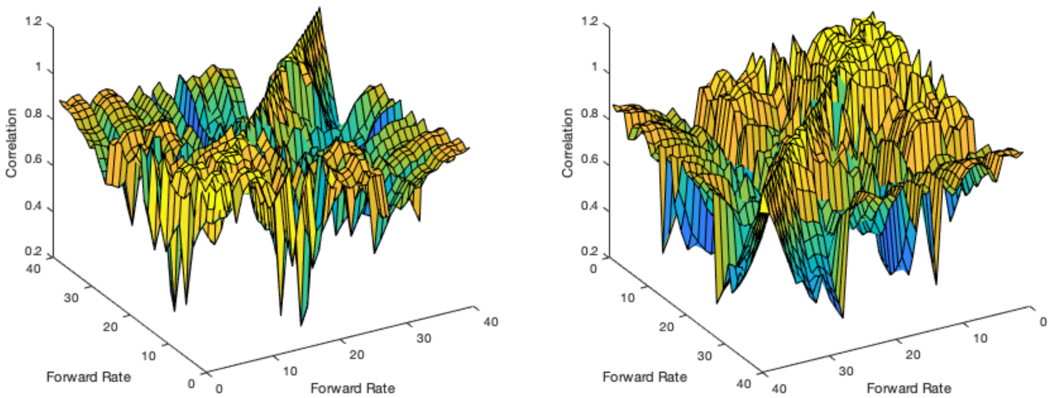

Fig. 5.2: Correlation Matrix: Calibration to BEV.

Figures 5.2 and 5.3 show the resultant instantaneous correlation matrices. The calibrated LFMM is used to price the full 40×40 grid of swaption volatilities with $\alpha = 1, 2, ..., 40$ and $\beta = 2, 3, ..., 41$ and this is compared to the swaption BEVs.

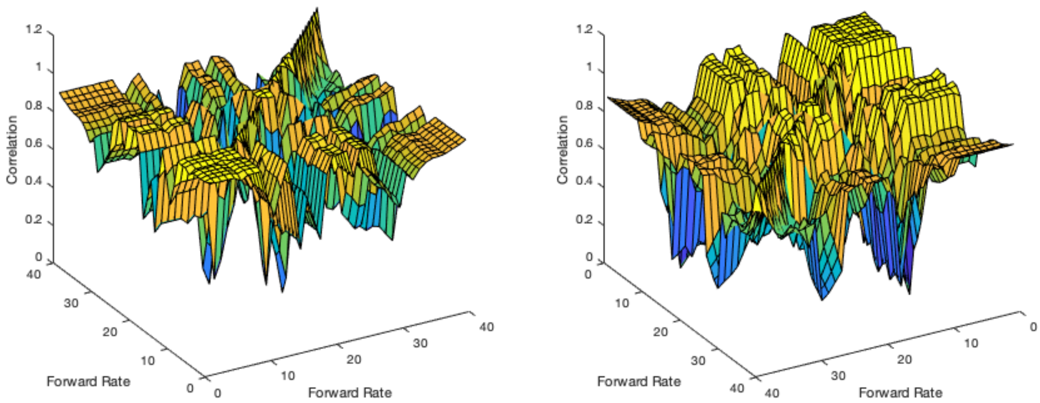

Fig. 5.3: Correlation Matrix: Calibration to 95^{th} Percentile BEV.

	BEV	95th Percentile BEV
Average Absolute Relative Error	3.8036%	9.6329%

Tab. 5.3: Goodness of Fit.

As can be seen in table 5.3, the calibrated model recovers the swaption BEV's excellently. However, this is only achieved due to the large degree of freedom in the correlation matrix. The matrix exhibits none of the desired properties and is nonsensical. The fact the LFMM in this context is unable to recover BEV swaption prices generated solely from past interest rate information and not market volatility quotes, without a nonsensical correlation matrix, is an indication that the underlying market dynamics are not compatible with those of the LFMM. This corroborates the findings of Kleisinger-Yu *et al.* (2018).

This is likely due to the deficiency of the LFMM in that it cannot reproduce the volatility skew as seen with BEV. Therefore, alternative approaches that such as the Shifted-Lognormal Model and other extensions to the LFMM including various Stochastic Volatility models and Uncertain-Parameters models as unpacked in Brigo and Mercurio (2007) that incorporate the volatility skew should be considered for this purpose. However, this analysis lies beyond the scope of this dissertation.

Chapter 6

Conclusion

This dissertation ultimately assessed the use of the Lognormal Forward-LIBOR Market Model in the South African interest rate market using BEV.

Chapter 2 analysed various hedging techniques for caps, floors and swaptions using simulated data from the Vasicek and G2++ short-rate models. The modified delta was found to perform better than the traditional and was used as the technique of choice throughout the rest of the dissertation.

Chapter 3 investigated BEV with simulated market data. Although the results differ minimally between segmentation methods, non-overlapping segments are preferred for their statistical properties. A single historical price path yielded a well structured BEV skew with both Vasicek and G2++ simulated data.

Chapter 4 assessed BEV within the South African interest rate market. Although Chapter 3 argued that non-overlapping segments are preferred, the need for a practical assumption with the limited sample size with South African interest rate market data overrides the need for technical consistency and thus overlapping segments were used. It was found that several caps and swaptions are trading at volatilities that differ significantly from their BEV estimates.

Lastly, Chapter 5 calibrated the Lognormal Forward-LIBOR Market Model with South African BEVs using co-terminal swaptions. An adequate model fit was attained, but only when the instantaneous correlation matrix was allowed to become nonsensical. Using this, it was argued that the underlying dynamics of the South African interest rate market are incompatible with those underlying the Lognormal Forward-LIBOR Market Model.

Bibliography

Akima, H. (1970). A new method of interpolation and smooth curve fitting based on local procedures, *Journal of the Association for Computing Machinery* **17**(4): 589–602.

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

Brigo, D. and Mercurio, F. (2007). *Interest rate models-theory and practice: with smile, inflation and credit*, Springer Science and Business Media.

De Boor, C. (1978). *A practical guide to splines*, Vol. 27, Springer-Verlag New York.

Du Preez, P. F. (2011). *An investigation into popular methods for constructing yield curves*, Msc dissertation, university of pretoria.

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

Fritsch, F. N. and Carlson, R. E. (1980). Monotone piecewise cubic interpolation, *SIAM Journal on Numerical Analysis* **17**(2): 238–246.

Higham, N. (2002). Computing the nearest correlation matrix—a problem from finance, *Journal of Numerical Analysis* **22**(3): 329–343.

Hull, J. and White, A. (2013). Forward rate volatilities, swap rate volatilities, and implementation of the LIBOR market model, *Journal of fixed income* **10**: 46–62.

Hyman, J. M. (1983). Accurate monotonicity preserving cubic interpolation, *SIAM Journal on Scientific Statistical Computing* **4**(4): 645–654.

Jamshidian, F. (1989). An exact bond option formula, *The journal of Finance* **44**(1): 205–209.

Kleisinger-Yu, X., Cresswell, W., Sterley, C. and De Savigny, A. (2018). An assessment on the appropriateness of the use of the LFMM in South Africa, *Report*, African Collaboration for Quantitative Finance and Risk Research.

Mahomed, O. (2015). *South African Financial Markets*, Unpublished.

Musiela, M. and Rutkowski, M. (1997). *Martingale methods in financial modelling*, Vol. 36, Springer Finance, New York.

Rebonato, R. (1999). On the pricing implications of the joint lognormal assumption for the swaption and cap markets, *Journal of computational finance* **2**(3): 57–76.

Rebonato, R. (2005). *Volatility and correlation: the perfect hedger and the fox*, John Wiley and Sons.

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

Appendix A

Hedging of Caps, Floors and Swaptions

A.1 Derivation of Modified Black 1976 Model Delta

$$\delta : \frac{\partial V(t)}{\partial F(t;t_m,t_n)} = P\tau Z(t,t_n)\frac{\partial (F(t;t_m,t_n)\Phi(\eta d_1) - K\Phi(\eta d_2))}{\partial F(t;t_m,t_n)} + \frac{V(t)}{Z(t,t_n)}\left(\frac{\partial Z(t,t_n)}{\partial F(t;t_m,t_n)}\right)$$

where

$$\frac{\partial (F(t;t_m,t_n)\Phi(\eta d_1) - K\Phi(\eta d_2))}{\partial F(t;t_m,t_n)} = \Phi(\eta d_1) + F(t;t_m,t_n)\frac{\partial \Phi(\eta d_1)}{\partial F(t;t_m,t_n)} - K\frac{\partial \Phi(\eta d_2)}{\partial F(t;t_m,t_n)}$$
$$= \Phi(\eta d_1) + F(t;t_m,t_n)\Phi'(\eta d_1)\eta\frac{\partial d_1}{\partial F(t;t_m,t_n)} - K\Phi'(\eta d_2)\eta\frac{\partial d_2}{\partial F(t;t_m,t_n)}$$

with

$$\Phi'(\eta d_2) = \frac{1}{\sqrt{2\pi}}\exp\{-\frac{1}{2}(\eta d_1)^2\}\left(\frac{F(t;t_m,t_n)}{K}\right) = \Phi'(\eta d_1)\left(\frac{F(t;t_m,t_n)}{K}\right)$$

$$\frac{\partial d_1}{\partial F(t;t_m,t_n)} = \frac{\partial d_2}{\partial F(t;t_m,t_n)}$$

$$\Rightarrow \frac{\partial (F(t;t_m,t_n)\Phi(\eta d_1) - K\Phi(\eta d_2))}{\partial F(t;t_m,t_n)} = \Phi(\eta d_1).$$

Additionally,

$$\frac{\partial Z(t,t_n)}{\partial F(t;t_m,t_n)} = \frac{-Z(t,t_m)\tau}{(1+\tau F(t;t_m,t_n))^2}$$
$$= -\tau Z(t,t_n)\left(\frac{1}{1+\tau F(t;t_m,t_n)}\right)$$

therefore

$$\delta : \frac{\partial V(t)}{\partial F(t;t_m,t_n)} = P\tau Z(t,t_n)\eta\Phi(\eta d_1) - \frac{\tau}{1+\tau F(t;t_m,t_n)}V(t). \quad \square$$

A.2 Derivation of Black 1976 Model First-Order Discount Factor Sensitivities

1)

$$\frac{\partial V(t)}{\partial Z(t,t_m)} = P\tau Z(t,t_n)\eta\Bigg(\Phi(\eta d_1)\frac{\partial F(t;t_m,t_n)}{\partial Z(t,t_m)} + \eta F(t;t_m,t_n)\Phi'(\eta d_1)\frac{\partial d_1}{\partial Z(t,t_m)} - \eta K\Phi'(\eta d_2)\frac{\partial d_2}{\partial Z(t,t_m)}\Bigg)$$

where

$$\Phi'(\eta d_2) = \frac{1}{\sqrt{2\pi}}\exp\{-\frac{1}{2}(\eta d_1)^2\}\left(\frac{F(t;t_m,t_n)}{K}\right) = \Phi'(\eta d_1)\left(\frac{F(t;t_m,t_n)}{K}\right)$$

$$\frac{\partial d_1}{\partial Z(t,t_m)} = \frac{\partial d_2}{\partial Z(t,t_m)}$$

and

$$\frac{\partial F(t;t_m,t_n)}{\partial Z(t,t_m)} = \frac{1}{\tau Z(t,t_n)}.$$

Therefore,

$$\frac{\partial V(t)}{\partial Z(t,t_m)} = P\eta\Phi(\eta d_1). \quad \square$$

2)

$$\frac{\partial V(t)}{\partial Z(t,t_n)} = \frac{V(t)}{Z(t,t_n)} + P\tau Z(t,t_n)\eta\Bigg(\Phi(\eta d_1)\frac{\partial F(t;t_m,t_n)}{\partial Z(t,t_n)} + \eta F(t;t_m,t_n)\Phi'(\eta d_1)\frac{\partial d_1}{\partial Z(t,t_n)} - \eta K\Phi'(\eta d_2)\frac{\partial d_2}{\partial Z(t,t_n)}\Bigg)$$

where

$$\frac{\partial d_1}{\partial Z(t,t_n)} = \frac{\partial d_2}{\partial Z(t,t_n)}$$

$$\frac{\partial F(t;t_m,t_n)}{\partial Z(t,t_n)} = -\frac{Z(t,t_m)}{\tau Z(t,t_n)^2}.$$

Therefore,

$$\frac{\partial V(t)}{\partial Z(t,t_n)} = \frac{V(t) - P\eta\Phi(\eta d_1)Z(t,t_m)}{Z(t,t_n)}. \quad \square$$

A.3 Vasicek Model Zero Coupon Bond Price Formula

$$\text{ZBP}(t,t_m,t_n,X) = XP(t,t_m)\Phi(\sigma_p - h) - P(t,t_n)\Phi(-h)$$

where

$$P(t,t_n) = A(t,t_n)\exp\{-B(t,t_n)r(t)\}$$

with

$$A(t,t_n) = \exp\left\{(\theta - \frac{\sigma^2}{2k^2})(B(t,t_n) - (t_n - t)) - \frac{\sigma^2}{4k}B(t,t_n)^2\right\}$$
$$B(t,t_n) = \frac{1}{k}(1 - \exp\{-k(t_n - t)\})$$

and

$$\sigma_p = \sigma\sqrt{\frac{1 - \exp\{-2k(t_m - t)\}}{2k}}B(t_m,t_n)$$
$$h = \frac{1}{\sigma_p}\ln\left(\frac{P(t,t_n)}{P(t,t_m)X}\right) + \frac{\sigma_p}{2}.$$

A.4 G2++ Model Zero Coupon Bond Price Formula

$$\mathrm{ZBP}(t,t_m,t_n,X) = -P(t,t_n)N\left(\frac{\ln\left\{\frac{XP(t,t_m)}{P(t,t_n)}\right\}}{\Sigma(t,t_m,t_n)} - \frac{1}{2}\Sigma(t,t_m,t_n)\right) + XP(t,t_m)N\left(\frac{\ln\left\{\frac{XP(t,t_m)}{P(t,t_n)}\right\}}{\Sigma(t,t_m,t_n)} + \frac{1}{2}\Sigma(t,t_m,t_n)\right)$$

where

$$\Sigma(t,t_m,t_n)^2 = \frac{\sigma^2}{2a^3}\Big(1 - \exp\{-a(t_n - t_m)\}\Big)^2\Big(1 - \exp\{-2a(t_m - t)\}\Big) + \frac{\eta^2}{2b^3}\Big(1 - \exp\{-b(t_n - t_m)\}\Big)^2\Big(1 - \exp\{-2b(t_m - t)\}\Big) + 2\rho\frac{\eta\sigma}{ab(a+b)}\Big(1 - \exp\{-a(t_n - t_m)\}\Big)\Big(1 - \exp\{-b(t_n - t_m)\}\Big)\Big(1 - \exp\{-(a+b)(t_m - t)\}\Big)$$

and

$$P(t,t_n) = \exp\left\{-\int_t^{t_m}\varphi(u)du - \frac{1 - \exp\{-a(t_n - t)\}}{a}x(t) - \frac{1 - \exp\{-b(t_n - t)\}}{a}y(t) + \frac{1}{2}V(t,t_n)\right\}$$

with

$$V(t,t_n) = \frac{\sigma^2}{a^2}\left(t_n - t + \frac{2}{a}\exp\{-a(t_n - t)\} - \frac{1}{2a}\exp\{-2a(t_n - t)\} - \frac{3}{2a}\right) + \frac{\eta^2}{b^2}\left(t_n - t + \frac{2}{b}\exp\{-b(t_n - t)\} - \frac{1}{2b}\exp\{-2b(t_n - t)\} - \frac{3}{2b}\right) + 2\rho\frac{\sigma\eta}{ab}\left(t_n - t + \frac{\exp\{-a(t_n - t)\} - 1}{a} + \frac{\exp\{-b(t_n - t)\} - 1}{b} + \frac{\exp\{-(a+b)(t_n - t)\} - 1}{a+b}\right).$$