

Estimating Long Term Equity Implied Volatility

Danielle Ana Crawford

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.

June 23, 2019

*MPhil in Mathematical Finance,
University of Cape Town.*



The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

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.

Signed by candidate

June 23, 2019

Abstract

Estimating and extrapolating long term equity implied volatilities is of importance in the investment and insurance industry, where 'long term' refers to periods of ten to thirty years. Market-consistent calibration is difficult to perform in the South African market due to lack of long term liquid tradable derivatives. In this case, practitioners have to estimate the implied volatility surface across a range of expiries and moneyness levels. A detailed evaluation is performed for different estimation techniques to assess the strengths and weaknesses of each of the models. The estimation techniques considered include statistical and time-series techniques, non-parametric techniques and three potential methods which use the local volatility model.

Acknowledgements

I would like to express my deep gratitude to Obeid Mahomed, my research supervisor. I am grateful for his patience, guidance, useful critiques and immense knowledge he provided throughout my dissertation. Furthermore, I would like to thank the lectures at AIFMRM for the valuable courses that provided me with the skills to be able to complete this dissertation. Lastly, I would like to thank the Banking Sector Education and Training Authority (BANKSETA) for the financial support they provided.

Contents

1. Introduction	1
2. Traditional Estimation Techniques	3
2.1 Historical Volatility	3
2.2 EWMA Volatility	4
2.3 GARCH Volatility	5
2.4 Overlapping vs. Non-Overlapping Returns	6
2.5 Data Sample Size	10
2.6 IVHV Ratios	11
3. Non-Parametric Techniques	14
3.1 Canonical Valuation	14
3.2 Break-Even Volatility	16
3.3 Data Sample Size	16
3.3.1 Canonical Valuation Implementation	17
3.3.2 Break-Even Volatility Implementation	19
3.4 I.I.D. Returns	22
4. Potential Solutions Using a Local Volatility Model	30
4.1 Local Volatility	30
4.2 Forecasting Using the Local Volatility Model	32
4.2.1 Method 1: Assume Constant Gradient Ratio	32
4.2.2 Method 2: Random Selection	34
4.2.3 Method 3: Random Path-Wise Selection	35
5. Conclusion	39
Bibliography	41
A. Supporting Figures	43
B. Alternative Models	48
B.1 Heston Model	48
B.2 Merton Model	49

C. Local Volatility	50
C.1 Derivation of the Local Volatility Formula	50
C.2 Finite Difference Techniques	51

List of Figures

2.1	Historical terminal realised volatility using τ -period returns.	7
2.2	Historical path-wise realised volatility using τ -period returns.	8
2.3	EWMA monthly terminal volatility using overlapping and non-overlapping monthly returns.	9
2.4	EWMA monthly path-wise volatility using overlapping and non-overlapping monthly returns.	9
2.5	Volatility estimates using non-overlapping daily returns with varying τ	10
2.6	GJR-GARCH(1,1) volatility forecasts.	11
2.7	Comparison between historical and implied volatilities for a range of maturities.	12
2.8	Average IVHV Ratios for the JSE Top40 Index from 2010-2018.	13
3.1	Volatility estimates using Canonical Valuation with increasing GBM sample size while applying different sampling techniques.	18
3.2	Volatility skew using Canonical Valuation while applying different sampling techniques using GBM data.	19
3.3	Profit and loss profile of a three month at-the-money call option using one thousand years of simulated GBM price data with $\sigma = 30\%$	20
3.4	Volatility estimates using Break-Even volatility with increasing GBM sample size while applying different sampling techniques.	21
3.5	Volatility skew using Break-Even volatility while applying different sampling techniques using GBM data.	22
3.6	Volatility estimates using Canonical Valuation with increasing Heston sample size while applying different sampling techniques.	23
3.7	Volatility skew using Canonical Valuation while applying different sampling techniques using Heston data.	24
3.8	Volatility estimates using Break-Even volatility with increasing Heston sample size while applying different sampling techniques.	25
3.9	Volatility skew using Break-Even volatility while applying different sampling techniques using Heston data.	26
3.10	Volatility estimates using Canonical Valuation with increasing Merton sample size while applying different sampling techniques.	27
3.11	Volatility skew using Canonical Valuation while applying different sampling techniques using Merton data.	28
3.12	Volatility estimates using Break-Even volatility with increasing Merton sample size while applying different sampling techniques.	28

3.13	Volatility skew using Break-Even volatility while applying different sampling techniques using Merton data.	29
4.1	JSE Top40 volatility surfaces on the 20/11/17.	33
4.2	Implied volatility estimates using Method 1 for JSE Top40 at-the-money European call.	34
4.3	Implied volatility estimates using Method 2 for JSE Top40 at-the-money European call.	36
4.4	Implied volatility estimates using Method 3 for JSE Top40 at-the-money European call.	38
A.1	Volatility estimates using Canonical Valuation with GBM data while applying different sampling techniques with the same Monte Carlo estimates.	43
A.2	Volatility estimates using Break-Even volatility with GBM data while applying different sampling techniques with the same Monte Carlo estimates.	44
A.3	Volatility estimates using Canonical Valuation with Heston data while applying different sampling techniques with the same Monte Carlo estimates.	45
A.4	Volatility estimates using Break-Even with Heston data while applying different sampling techniques with the same Monte Carlo estimates.	45
A.5	Volatility estimates using Canonical Valuation with Merton data while applying different sampling techniques with the same Monte Carlo estimates.	46
A.6	Volatility estimates using Break-Even volatility with Merton data while applying different sampling techniques with the same Monte Carlo estimates.	47

List of Tables

2.1	Average IVHV Ratios for JSE Top40 Index from 1996-2018.	13
B.1	Heston parameters used in Chapter 3.	48
B.2	Merton parameters used in Chapter 3.	49

Chapter 1

Introduction

The ability to estimate and extrapolate implied volatilities is critical to the insurance industry. This is due to the complex nature of long-term savings, retirement and investment products, offered by insurance companies, which incorporate embedded financial derivatives. In the insurance and investment industry 'long term' refers to periods of ten to thirty years. There is a lack of liquid tradable derivatives with long enough expiries to facilitate market-consistent calibration for the purposes of pricing and hedging of insurance investment products. This is particularly true in the South African market, where there are few long term traded instruments (Flint *et al.*, 2014). In this case, practitioners have to resort to extrapolating implied volatility surfaces which are defined across a range of strikes and expiries. It is recommended by current legislative and advisory practice notes¹ to use market-consistent models; these are models that can reproduce market prices as closely as possible (Flint *et al.*, 2014).

Flint *et al.* (2014) extensively covers commonly used implied volatility models, applying South African market data to these models. The objective of Flint *et al.* (2014)'s study is to accurately estimate long term implied volatility while using market consistent models. This dissertation reviews the ability of these commonly used techniques to estimate or extrapolate long term implied volatility using statistical and time-series techniques in Chapter 2 as well as non-parametric techniques in Chapter 3, both of which are covered in Flint *et al.* (2014). To extend on the research done by Flint *et al.* (2014), potential solutions using a local volatility model are proposed and reviewed.

¹ Actuarial Society of South Africa. APN 110: Allowance for Embedded Investment Derivatives, Version 4. Advisory Practice Note, 2012. Actuarial Society of South Africa. Market Consistent Calibration in South Africa. APN 110 sub-committee presentation, 2010

This study will be done in three sections, each describing different estimation techniques:

1. **Statistical and Time-Series Techniques:** Methods which only require historical time-series data to estimate volatility. The three methods considered are historical volatility, GJR-GARCH and exponentially weighted moving average. These methods do not directly estimate implied volatility.
2. **Non-Parametric Techniques:** Methods which only require historical time-series data to estimate fair volatilities. The two non-parametric methods considered are Canonical Valuation and Break-Even volatility. These methods directly estimate implied volatility.
3. **Potential Solutions Using a Local Volatility Model:** These methods require cross sectional option data. Three models are considered which make use of the Local Volatility model. These models directly estimate implied volatility.

This dissertation will include a detailed evaluation of each model's ability to estimate the entire implied volatility surface. Furthermore, it will investigate the ability of each model to extrapolate the implied volatilities that are not readily observable in the market (i.e for non-existing long dated options). Various features regarding each model are evaluated. These features include the data requirements, type of volatility estimated, if term structure estimation is possible, if skew estimation is possible across a range of moneyness levels, if statistical confidence intervals for the volatility estimates are possible to construct and the sensitivity of the volatility estimate to new data. The analysis is performed using both simulated and real-world data. Data is simulated using geometric Brownian motion (GBM), Heston (Drgulescu and Yakovenko, 2002) and Merton models (Linghao, n.d.). The FTSE/JSE Top40 Index is the most liquid in the South African equity market thus this data is used in the analysis. The JSE Top40 Index data set used in this paper includes daily spot index levels spanning a 22.75 year period from 30 August 1995 to 29 April 2018 and option implied volatilities for a range of expiries.

In the analysis which follows, the underlying asset price at time t will be denoted by the variable S_t . The τ -period log-return of the underlying asset will be denoted as follows:

$$r_{t,\tau} = \ln \left(\frac{S_t}{S_{(t-\tau)}} \right). \quad (1.1)$$

The holding period of the asset is from $(t - \tau)$ to t resulting in one log-return over the τ -period. It is assumed throughout this dissertation that τ is expressed in terms of business days, with 21 business days in a month and 252 business days in a year (Flint *et al.*, 2012).

Chapter 2

Traditional Estimation Techniques

Market implied volatility data can be sparse. In this case, techniques which use only historical time-series data of the underlying asset can be used. These estimation techniques result in historical volatility estimates which need to be scaled to produce an appropriate implied volatility estimate. A commonly used scaling factor is referred to as the implied volatility historical volatility (IVHV) ratio (Flint *et al.*, 2014). The estimation techniques that will be considered can be categorised as follows:

- (i) Statistical Techniques
- (ii) Time-Series Techniques

These techniques use different assumptions to characterize the underlying asset's data generating process.

2.1 Historical Volatility

Historical volatility, also referred to as statistical volatility, is often defined as the standard deviation of the log-returns of an asset. This method of calculating volatility is more commonly known as sample standard deviation and assumes that future volatility will follow the same path as historical volatility. The estimate is usually annualised by multiplying the sample variance of return data by $\frac{252}{\tau}$, where 252 is the assumed number of business days in a year and τ is the number of business days in the period under consideration. The important variables to consider in this calculation are the length of the sampling period (τ) and the number of sampling periods (n). It is important to ensure that the asset returns are non-overlapping, and independently and identically distributed (i.i.d.). Given the latter, classical historical volatility is defined as:

$$\hat{\sigma}(t, \tau, n) = \sqrt{\left(\frac{252}{\tau}\right) \left(\frac{1}{n-1}\right) \sum_{i=1}^n [r_{(t_i, \tau)} - \bar{r}(t, \tau, n)]^2}, \quad (2.1)$$

where $\bar{r}(t, \tau, n) = \frac{1}{n} \sum_{i=1}^n r_{(t_i, \tau)}$ is the average τ -period asset return and t is the current time t_n (Fildes, 1994). The resultant volatility estimates the deviation of the τ -period returns from the average τ -period return. An alternative to equation

(2.1) is known as historical realised volatility where the average asset return is not subtracted from the return values or rather, is assumed to be zero (Flint *et al.*, 2014):

$$\hat{\sigma}(t, \tau, n) = \sqrt{\left(\frac{252}{\tau}\right) \left(\frac{1}{n-1}\right) \sum_{i=1}^n (r_{(t_i, \tau)})^2}. \quad (2.2)$$

Equation (2.2) is more commonly used to calculate historical volatility. This method assumes that the mean τ -period return is zero. Thus, realised volatility is generally used when the average returns are close to zero. Assuming the mean return is zero will avoid errors caused by estimating the mean, this is especially prevalent for short periods. Historical realised volatility estimates may be more reliable than classical historical volatility (Ederington and Guan, 2006).

It is common in academic literature to consider daily returns over a τ -period when estimating the volatility for derivative pricing (Flint *et al.*, 2012). Given daily asset returns the sample historical volatility is given by:

$$\hat{\sigma}(t, 1, \tau) = \sqrt{(252) \left(\frac{1}{\tau-1}\right) \sum_{i=1}^{\tau} [r_{(t_i, 1)} - \bar{r}(t, 1, \tau)]^2}. \quad (2.3)$$

where $\bar{r}(t, 1, \tau)$ is the one day average asset return for the τ -period under consideration. The resultant volatility estimate using equation (2.3) can be described as a path-wise estimate. This is because the volatility estimate for the τ -period will depend on the daily returns in the τ -period. Equation (2.1) results in a terminal volatility estimate where only τ -period returns are considered. It should also be considered if each volatility estimate will use overlapping (rolling) or non-overlapping (blocked) price data. Overlapping data estimates are advantageous when there is a lack of price data but results in biased volatility estimates due to auto-correlation in the overlapping τ -period return data.

Historical volatility is a simple method for forecasting the volatility of stock price returns. Implied volatility is not calculated directly when using historical volatility methods and therefore needs to be scaled with an estimated IVHV ratio. This estimating technique has been criticised for assuming that future volatility will follow the same path as historical volatility. To forecast thirty year implied volatility using equation (2.1), thirty years of historical stock price data are needed for just one estimate. This is also true for the path-wise volatility estimate where thirty years of daily returns are needed for the thirty year volatility estimate. These amounts of data are not always available therefore this forecasting technique may not be viable. This estimation technique has also been criticised for weighting all historical returns equally irrespective of how far the data extends into the past (Ederington and Guan, 2006).

2.2 EWMA Volatility

The idea behind the exponentially weighted moving average (EWMA) model is that data that is more recent influences the volatility forecasts more than data that

are further in the past. This is enforced in the model by weighting more recent events higher than those in the past (Korkmaz and Aydin, 2002). This method improves on the estimates of equation (2.3), where daily return data equally contributes to the volatility estimate. It is assumed that the mean return of the data are zero and the returns are i.i.d. The EWMA volatility forecast at time t is:

$$\sigma_t = \sqrt{(1 - \lambda) \sum_{i=0}^{n-1} \lambda^{i-1} r_{(t-i,1)}^2}, \quad (2.4)$$

where λ , referred to as the decay parameter, is a constant between 0 and 1, $r_{(t-i,1)}$ is the log-return at time $(t - i)$ and n is the number of historical observations under consideration. The larger the λ parameter the more influence the recent returns will have on the volatility estimate. The optimal λ value can be found using an optimization method described by Ladokhin (2009). Equation (2.4) can be reduced to the following recursive form:

$$\sigma_t^2 = \lambda \sigma_{(t-1)}^2 + (1 - \lambda) r_{(t,1)}^2, \quad (2.5)$$

where the next period's variance σ_t^2 depends on the current period's variance $\sigma_{(t-1)}^2$ and return $r_{(t,1)}$. It is recommended by Korkmaz and Aydin (2002) to use $\lambda = 0.94$ for daily volatility forecasting and 0.97 for monthly volatility forecasting (Vieira *et al.*, 2012). Again, the resultant volatility estimate needs to be scaled by an appropriate IVHV ratio to obtain the implied volatility estimate. The disadvantage of the EWMA method is that the method does not capture asymmetry¹ in volatility. Another model shortcoming is that it is assumed that returns are i.i.d., which has been empirically proven to be incorrect (Ederington and Guan, 2006).

2.3 GARCH Volatility

The generalised autoregressive conditional heteroscedasticity, GARCH(p,q), model is a standard model used to capture market features. These features are statistical properties commonly seen in market data which include heavy-tailed distributions, skewed distributions, volatility clustering and volatility feedback effects² (Flint *et al.*, 2014). These stylised features oppose the fundamental assumption that asset returns are i.i.d, therefore the GARCH model is a suitable model to capture these market features. These models are autoregressive models in squared returns of the data, meaning that the current return is dependent on past returns. The conditional term describes that the volatility of the following period is conditional on the features of the current period. Heteroscedasticity describes the varying volatility of each period in the time-series data (Reider, 2009).

¹ Volatility tends to be higher in periods where returns are decreasing than in periods where returns are increasing.

² Feedback effects describes the negative correlation between asset returns and volatility.

The model assumes the daily return on an asset or portfolio is:

$$r_{t,1} = \mathbb{E}[r_{t,1} | \mathcal{F}_{(t-1)}] + \epsilon_t \quad (2.6)$$

$$\epsilon_t = \sigma_t z_t, \quad (2.7)$$

where the mean return is conditional on all past information up to $(t - 1)$, denoted by $\mathcal{F}_{(t-1)}$, and $\{z_t\}$ are i.i.d. $N(0, 1)$ random variables for the standard GARCH model. The variable ϵ_t is the residual return (AL-Najjar, 2016). The conditional variance is defined as:

$$\sigma_t^2 = \omega + \sum_{j=1}^p \alpha_j \epsilon_{(t-j)}^2 + \sum_{k=1}^q \beta_k \sigma_{(t-k)}^2. \quad (2.8)$$

There are many forms of the GARCH model in which the model features are specified differently. These features include the functional form of the conditional variance equation, the distribution of $\{z(t)\}$ and the number of lags or choice of parameters p and q . The GARCH(1,1) model is the most popular volatility model used for financial time-series data with parameters: $p = 1$, $q = 1$, $\omega > 0$, $\alpha_1 > 0$, $\beta_1 > 0$ and $\alpha_1 + \beta_1 < 1$. It is recommended by Flint *et al.* (2014) to use the GJR-GARCH(1,1) model with student-t innovations for the daily JSE Top40 dataset:

$$\sigma_t^2 = \omega + [\alpha + \gamma I_{\epsilon_{(t-1)} < 0}] \epsilon_{(t-1)}^2 + \beta \sigma_{(t-1)}^2, \quad (2.9)$$

where $I_{\epsilon_{(t-1)} < 0}$ is the indicator function which is one when $\epsilon_{(t-1)} < 0$ (Hansen and Lunde, 2005).

This model accounts for asymmetric volatility due to positive and negative return shocks, which the GARCH(1,1) model does not. The γ parameter controls the effects of the asymmetric features in the data such that negative shocks result in the addition of $\alpha + \gamma$ whereas positive shocks result in the addition of α . If $\gamma > 0$, negative shocks will have a larger influence on the volatility estimate and if $\gamma < 0$ positive shocks will have a larger effect (Liu and Hung, 2010). The parameters ω , α and β must all be non-negative. It is also required that $\alpha + \beta + 0.5\gamma < 1$ and $\alpha + 0.5\gamma$ must be positive (Liu and Hung, 2010).

2.4 Overlapping vs. Non-Overlapping Returns

The first data consideration is if overlapping or non-overlapping returns should be used when calculating the volatility estimate. Using overlapping returns is advantageous when there is a lack of stock price data as it artificially creates more return data. In this case, the return data used for each τ -period would overlap with the prior period by $(N - 1)$, such that there are N returns in each period. For example, if there are two weeks of stock price data available and weekly volatility is estimated then non-overlapping returns would result in two stock price returns for the two week period whereas overlapping returns would result in six stock price returns.

Using overlapping returns may be beneficial but it may not be suitable for the estimation technique in use. Historical volatility and EWMA methods assume the




Fig. 2.1: Historical terminal realised volatility using τ -period returns.

return data are i.i.d. By using overlapping returns this will induce auto-correlation between returns therefore breaking the assumption that returns are independent. To illustrate this, consider ten years of simulated GBM data. Equation (2.2) is implemented to estimate historical realised volatility for τ -periods increasing from a week to a month, as seen in Figure 2.1. This figure illustrates the effect of using overlapping and non-overlapping τ -period returns on the volatility estimate. Using overlapping returns results in more volatile estimates that deviate more from the analytical solution in comparison to using non-overlapping returns, consequentially leading to less accurate estimates. Figure 2.2(a) illustrates the effect of using overlapping or non-overlapping returns on the volatility estimate using the historical volatility path-wise estimation technique, as seen in equation (2.3). Monthly volatility is estimated iteratively through the ten years of simulated GBM data. Since, in the overlapping case, each volatility estimate uses $(\tau - 1)$ returns from the prior period, the volatility estimates are correlated. The non-overlapping technique results in volatility estimates which coincide with the overlapping estimates when the periods coincide. Figure 2.2(b) illustrates the average non-overlapping path-wise estimates for varying τ -periods. The estimates are dependent on the stock price path. There is little variation between the overlapping and non-overlapping estimates³. The path-wise historical estimation technique using daily returns results in volatility estimates that are not very sensitive to the use of overlapping or non-overlapping data.

To illustrate the same effect, now using the EWMA estimation technique, consider ten years of simulated GBM data. Equation (2.5) is used to estimate EWMA monthly volatility using overlapping and non-overlapping return data for the monthly-periods under consideration. Monthly volatility is calculated throughout the ten years of simulated data using consecutive monthly periods. Since monthly volatility is being estimated $\lambda = 0.97$ is used. Figure 2.3 illustrates the effect of using overlapping returns in the volatility calculation in comparison to using non-overlapping returns. Using overlapping returns results in more volatile volatility

³ The values of the parameters in this example has resulted in the underestimation of the volatility. It cannot be concluded that the pathwise volatility estimate will always underestimate the volatility regardless of the size of parameters as this is not always the case.

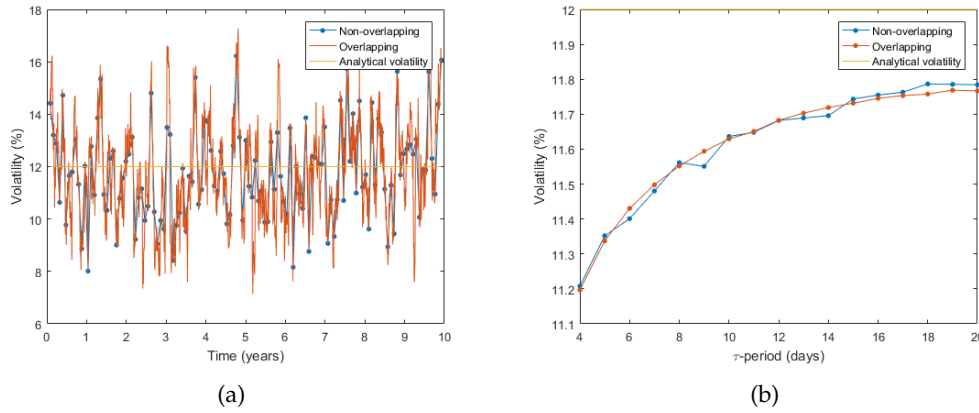


Fig. 2.2: Historical realised volatility using τ -period returns. (a) Path-wise volatility estimates with τ equal to one month. (b) Average path-wise volatility estimates for varying τ -periods.

estimates which consequently leads to less accurate estimates. This is confirmed by comparing the analytical volatility to the estimated volatility by measuring the Root Mean Squared Error (RMSE)⁴ of the estimates (Chai and Draxler, 2014). The RMSE of the volatility estimates using overlapping returns is higher than that of the non-overlapping returns. Figure 2.4(a) illustrates the effect of using overlapping or non-overlapping τ daily returns on the path-wise volatility estimates using equation (2.4). Using overlapping returns results in more volatile estimates in comparison to using non-overlapping daily returns. The non-overlapping estimates coincide with the overlapping estimates when the τ -period coincides. Path-wise volatility estimates tend to be smaller than the analytical solution. This is due to the sample size τ . Figure 2.4(b) illustrates the average EWMA volatility estimate using varying τ daily returns. The non-overlapping and overlapping techniques result in volatility estimates which do not differ greatly. The volatility estimate is dependent on the path of the return data.

It is important to ensure that the return data used, when estimating volatility with techniques that assume the return data is i.i.d., is not overlapping. Using overlapping returns will result in an autocorrelated series of returns which contradicts the assumptions of the model, leading to less accurate estimates. This is of particular importance when a terminal estimation technique is used. Since stock price returns are known to not be i.i.d. the GARCH model may be a more suitable alternative to the historical and EWMA models.

⁴ This method is used to measure the accuracy of the estimates by calculating the square-root of the average squared difference between the estimates and the analytical solution.




Fig. 2.3: EWMA monthly terminal volatility using overlapping and non-overlapping monthly returns.

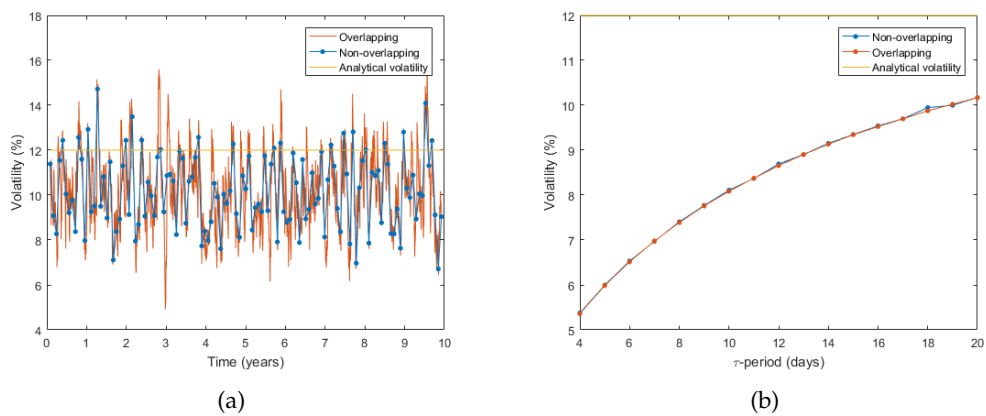


Fig. 2.4: EWMA monthly path-wise volatility using overlapping and non-overlapping monthly returns. (a) Path-wise volatility estimates with τ equal to one month. (b) Average path-wise volatility estimates with varying τ .

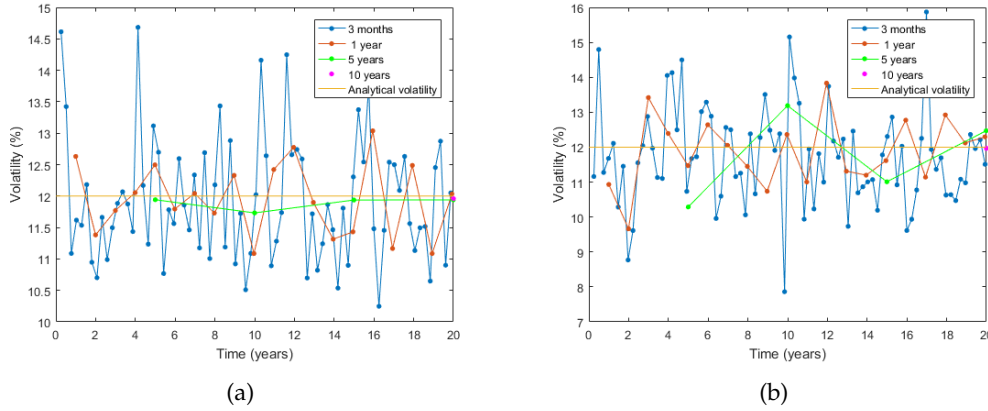


Fig. 2.5: Volatility estimates using non-overlapping daily returns with varying τ . (a) Historical realised volatility. (b) EWMA volatility.

2.5 Data Sample Size

The second data consideration is the size of the data sample used for the volatility estimate. Different estimation techniques require different amounts of data to ensure accurate estimates. Thus, it is important to assess how much data is needed for the estimation technique under consideration.

Term structure estimation is possible using historical or EWMA methods by varying the sample size or by obtaining the volatility estimate and then using the square-root of time rule. It is possible to use the square-root of time rule because these volatility estimation techniques assume the asset returns are i.i.d. The path-wise estimation techniques using daily non-overlapping returns are considered in this analysis.

Realised volatility estimates with varying τ -periods are shown in Figure 2.5(a) using twenty years of simulated GBM data. The smaller the τ -period the greater the variability in the volatility estimates. Therefore, calculating volatility using larger periods leads to more accurate results when forecasting long-term volatility. Smaller τ -periods are more sensitive to outliers in the data whereas larger τ -periods are less effected by sampling period (Ederington and Guan, 2006). The larger the τ -period the more normal the distribution of the return data, due to the central limit theorem.

The decay parameter λ used in the EWMA method is dependent on the size of τ . Greater τ -periods require larger decay parameter values. Figure 2.5(b) illustrates the term structure estimation for twenty years of simulated GBM daily stock price data with varying τ -periods. Shorter τ -periods are more sensitive to return shocks than longer τ -periods.

Term structure estimation is possible using GARCH models. Once a GARCH model is fitted to the return data, term structure specification is possible for volatility forecasts. To illustrate this, five years of daily GBM stock price data is simulated to calculate daily GJR-GARCH(1,1) volatility using student-t innovations. The Mat-




Fig. 2.6: GJR-GARCH(1,1) volatility forecasts.

lab function *estimate* is used to estimate the parameters in equation (2.9). These parameters are used as an input to the Matlab function *forecast* which forecasts daily volatility and *infer* which infers conditional variances. Figure 2.6 displays the daily GJR-GARCH(1,1) inferred volatility and volatility forecasts for the simulated GBM stock prices. It is important to note that the volatility forecasts do not change significantly after four months as seen in Figure 2.6. The estimates after four months are very close to the given volatility of 12%. In this case since the data is i.i.d. the GJR-GARCH(1,1) model results in a good estimate of the long term volatility. Typically the GJR-GARCH(1,1) model is used to forecast volatility from one month to one year. The fitted model may not be suitable for forecasting long term volatility (Flint *et al.*, 2014).

2.6 IVHV Ratios

Historical volatility forecasts future volatility based on past values, whereas implied volatility is the markets view on what future volatility will be; these two volatility estimates differ. This difference is called the volatility premium. Implied volatility is known to generally be higher than historical volatility (Eraker, 2008). Daily log-returns of the JSE Top40 Index levels are used to estimate historical volatility for one, three, six, nine and twelve month periods; these values are compared to market implied volatilities, for at-the-money options, to calculate the IVHV ratio.

Figure 2.7 illustrates the implied and historical volatilities for a range of ex-

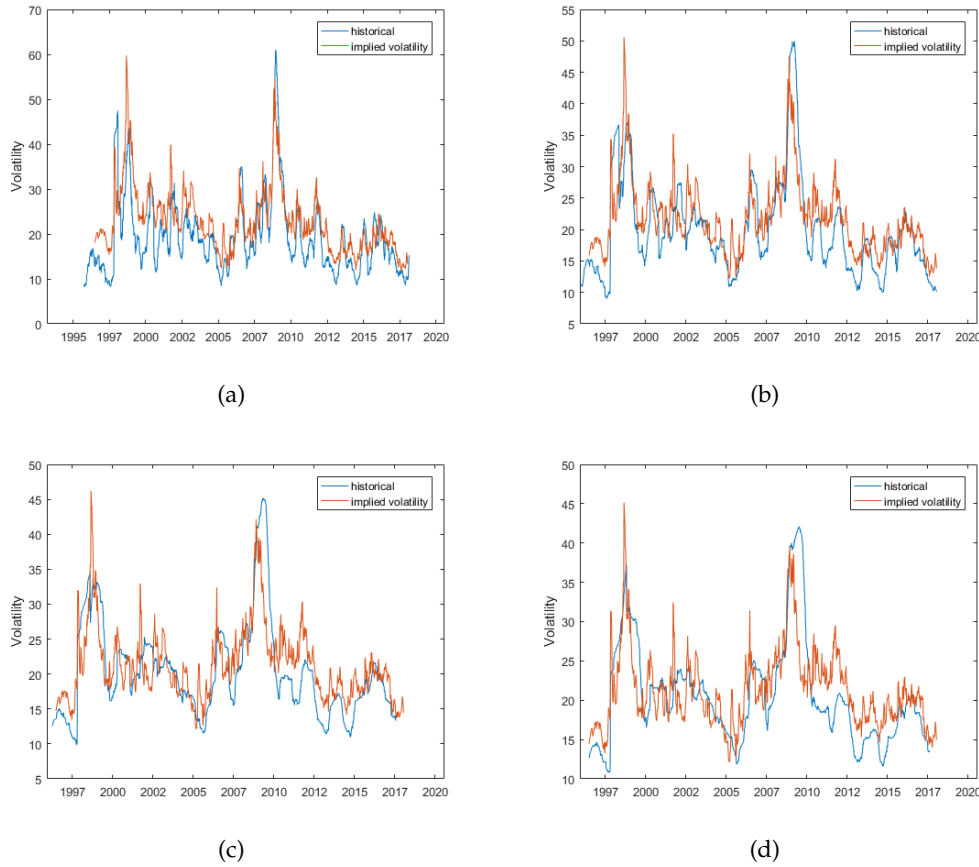


Fig. 2.7: Comparison between historical and implied volatilities for a range of maturities. (a) 3 months. (b) 6 months. (c) 9 months. (d) 1 year.

piries from 1996 to 2018. The implied volatilities for at-the-money options on the JSE Top40 Index are predominantly higher than historical volatility for each of the expiries. The average IVHV ratio for each term is above one, as seen in Table 2.1. The IVHV ratio fluctuates over time as seen in Figure 2.7. In high volatility periods i.e. 2008 financial crisis, the implied volatilities are smaller than historical volatilities for each maturity. There is a decrease in IVHV ratio with an increase in maturity with the exception of twelve months, due to the lack of data. A power trend is fitted to the mean IVHV ratios for the JSE Top40 Index from 2010-2018, as seen in Figure 2.8(b) (Flint *et al.*, 2014). The 5% and 95% confidence intervals for the IVHV ratio is used to scale resultant volatility estimates for different expiries. This results in an implied volatility range for a given maturity, rather than a point estimate.

It is not possible to directly estimate implied volatility for options that are in- or out-the-money when using the traditional estimation techniques described. Thus, volatility skew or smile can only be estimated by scaling historical volatility estimates by in- or out-the-money IVHV ratios to obtain implied volatility skew per maturity.

	1 month	3 months	6 months	9 months	12 months
Mean	1.3401	1.2443	1.1488	1.0989	1.0882
Median	1.2928	1.2328	1.1358	1.0961	1.0972

Tab. 2.1: Average IVHV Ratios for JSE Top40 Index from 1996-2018.

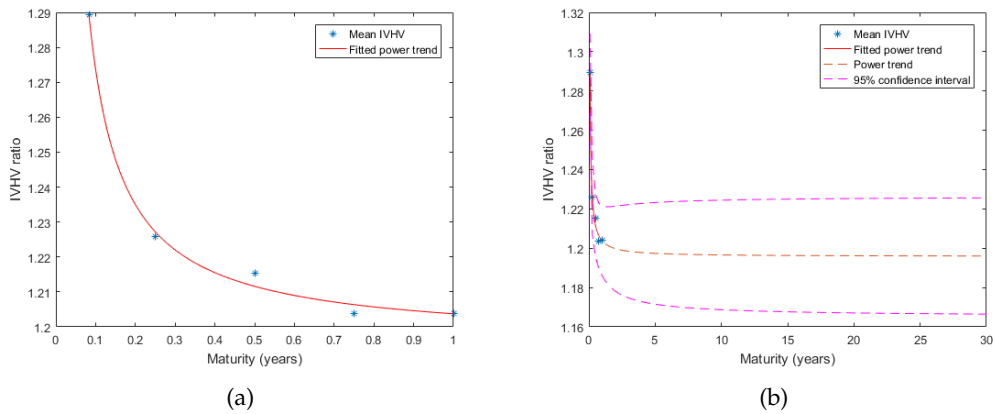


Fig. 2.8: Average IVHV Ratios for the JSE Top40 Index from 2010-2018. (a) Average IVHV Ratios for the JSE Top40 Index fitted to a power trend. (b) Extrapolated IVHV Ratios for the JSE Top40 Index.

Chapter 3

Non-Parametric Techniques

Non-parametric methods are used to obtain market-consistent fair volatilities using only historical market data. The two non-parametric methods used by [Flint et al. \(2014\)](#) to estimate volatilities are Canonical Valuation (CV) and Break-Even volatility (BEV). This chapter explores these two methods and assesses their ability to estimate the implied volatility surface.

3.1 Canonical Valuation

Canonical Valuation is a non-parametric method which values derivatives under a risk-neutral framework. This pricing method only requires a series of historical stock prices. The underlying distribution of the stock price data over some τ -period is not assumed to be log-normal, as in the Black-Scholes framework, but rather uses an estimation of the distribution of historical returns to predict the risk-neutral distribution of asset prices at expiration ([Gray and Newman, 2005](#)). This is done using the principle of minimum relative entropy. This pricing method is named after the Gibbs canonical distribution which is used to estimate the risk-neutral distribution ([Stutzer, 1996](#)).

Once this distribution is estimated, a vanilla European option can be priced; these prices are then used to imply fair volatilities for the τ -period, such that τ is defined as $\tau = [t_0, T]$. The risk-neutral value of an option at time t_0 , over the τ -period, is given by the discounted present value of the pay-off of the option under the risk-neutral measure:

$$\begin{aligned} V_{t_0, T}(S_T) &= Z_{(t_0, T)} \mathbb{E}^{\mathbb{Q}}[V_T | \mathcal{F}_{t_0}] \\ &= Z_{(t_0, T)} \int_0^{\infty} V_T \pi^*(S_T) dS_T, \end{aligned} \quad (3.1)$$

where $Z_{(t_0, T)}$ denotes the risk-free discount factor over the period $[t_0, T]$, $\mathbb{E}^{\mathbb{Q}}[\cdot | \cdot]$ is the conditional expectation under the risk-neutral probability measure \mathbb{Q} , V_T is the value of the derivative at time T and $\pi^*(S_T)$ is the risk-neutral probability density function of the underlying asset price at expiration T ([Flint et al., 2012](#)).

In order to implement this pricing method consider a European call option expiring at time T such that the current time is t_0 . Firstly, the real world probability distribution $\hat{\pi}$, of the stock price returns at expiry, is estimated using a time-series

of historical stock price returns. This is done by calculating a sample of N τ -period asset returns $r_{(t_i, \tau)}$ with $i = 1, \dots, N$, using equation (1.1) where τ is equal to the length of the option. These N asset returns are used to calculate a series of possible future stock prices at expiration as follows:

$$S_{(T_i, \tau)} = S_0 e^{r_{(t_i, \tau)}}, \quad i = 1, \dots, N. \quad (3.2)$$

where S_0 is the current stock price and $e^{r_{(t_i, \tau)}}$ is the growth rate. Each of the $S_{(T_i, \tau)}$ values are allocated an equal real-world probability of $\hat{\pi} = \frac{1}{N}$ (Alcock and Gray, 2005). This empirical distribution can be transformed into its risk-neutral counterpart π^* on condition the following is satisfied to ensure arbitrage-free option pricing:

$$S_{t_i} = \mathbb{E}_{t_i}^{\mathbb{Q}} \left[\frac{S_{(T_i, \tau)}}{e^{r\tau}} \right] = \mathbb{E}_{t_i}^{\mathbb{P}} \left[\frac{S_{(T_i, \tau)}}{e^{r\tau}} \frac{d\pi^*}{d\hat{\pi}} \right], \quad (3.3)$$

where r is the risk-free rate and $\frac{d\pi^*}{d\hat{\pi}}$ is the Radon-Nykodym derivative of the risk-neutral measure with respect the real world measure. The superscript of the expectation indicates which probability measure is used to calculate the expectation. Equation (3.3) can be simplified to:

$$1 = \sum_{i=1}^N e^{-r\tau} X_{(t_i, \tau)} \frac{\pi_i^*}{\hat{\pi}_i} \hat{\pi}_i, \quad (3.4)$$

where $X_{(t_i, \tau)} = e^{r_{(t_i, \tau)}}$ and π_i^* is the risk-neutral probability of return $r_{(t_i, \tau)}$ (Stutzer, 1996). Equation (3.4) ensures that the martingale property is satisfied and the discounted expected return on the stock sums to one (Alcock and Gray, 2005). The risk-neutral distribution can be estimated using the method of relative entropy which minimizes the difference between the empirical distribution and risk-neutral probability distribution of the underlying. Since the pay-off of the option depends on the terminal stock price, estimating the distribution of the underlying stock price is beneficial as stylized features in the data may be captured (Gray and Newman, 2005). The following relative entropy (RE) expression should be minimized to solve for π_i^* :

$$RE(\pi_i^*, \hat{\pi}_i) = \sum_{i=1}^N \pi_i^* \log \left(\frac{\pi_i^*}{\hat{\pi}_i} \right) \quad \text{s.t. (3.4)}. \quad (3.5)$$

The solution to this convex minimization problem is given by the Gibbs canonical distribution:

$$\pi_i^* = \frac{\exp(\gamma^* e^{-r\tau} X_{(t_i, \tau)})}{\sum_{i=1}^N \exp(\gamma^* e^{-r\tau} X_{(t_i, \tau)})}, \quad (3.6)$$

such that the Lagrange multiplier γ^* satisfies the following:

$$\gamma^* = \arg \min_{\gamma} \sum_{i=1}^N \exp[\gamma(e^{-r\tau} X_{(t_i, \tau)} - 1)]. \quad (3.7)$$

The call option can be priced using the resulting risk-neutral probabilities, as follows:

$$C_\tau(K) = e^{-r\tau} \sum_{i=1}^N \max[S_{(T_i,\tau)} - K, 0] \pi_i^*. \quad (3.8)$$

The difference between equation (3.8) and the Black-Scholes price is then set to zero to solve for the implied volatility (Stutzer, 1996).

3.2 Break-Even Volatility

Break-Even volatility is defined as the volatility, for a delta-hedged option, which results in zero profit and loss. This method is based on dynamic replication such that a portfolio consisting of an underlying asset and cash is used to replicate an option over a τ -period. Assuming the standard Black-Scholes framework, this delta-hedge will result in a profit and loss function which is set to zero by solving for the corresponding implied volatility. The profit and loss of a continuously delta hedged option is the gamma-weighted average of past quadratic returns:

$$\text{P\&L} = \int_0^T \frac{1}{2} \Gamma_u S_u^2 [\sigma_u^2 - \sigma_{imp}^2] du \quad (3.9)$$

where $\sigma_{imp} = \sigma_{(t_0,\tau)}$ is the fair implied volatility of the option over the τ -period, Γ is the standard Black-Scholes gamma of the option and σ_u the instantaneous realised volatility (Flint *et al.*, 2014). Equation (3.9) can be approximated by the following discrete profit and loss equation:

$$\text{P\&L}(t_0, T, \sigma_{imp}) = \frac{1}{2} \sum_{i=1}^{\tau} e^{r\tau} \Gamma_{t_i} S_{t_i}^2 [r_{(t_i,1)}^2 - \sigma_{imp}^2 \Delta t], \quad (3.10)$$

where $e^{r\tau}$ is the risk-free capitalization factor over the τ -period and $\Delta t = \frac{t_i - t_{(i-1)}}{365}$ (Flint *et al.*, 2014). Assuming the current time is t_0 with the option expiring at time T , hedging of the option occurs discretely at times $\{t_0, t_1, \dots, T\}$. It is suggested by Flint *et al.* (2014) to solve for the volatility that ensures the average P&L over the different time periods is zero; this results in a smoother volatility surface than the alternative, averaging the volatilities that results in zero P&L for each time period. Each returns period is rebased using the initial stock price, S_0 . This results in a smoother volatility surface with the consequence of considering relative strikes instead of absolute strike values (Dupire, 2006).

3.3 Data Sample Size

The first data consideration is assessing how much data is required for both CV and BEV methods to recover accurate results. Using real world stock price data is restrictive for this study as these methods are computationally intensive requiring a large amount of data to produce reasonable results. Thus, stock price data

is simulated using GBM. Using simulated GBM data is useful for testing the accuracy of the volatility estimates, as the GBM model assumes volatility is constant. Since using non-overlapping returns may result in too few τ -period returns, overlapping and statistical bootstrapping methods are considered as ways to increase the sample size.

3.3.1 Canonical Valuation Implementation

Daily stock prices are simulated and implemented to calculate an array of τ -period returns. By applying the Matlab function *fmincon*, these returns are then used to calculate the Lagrange multiplier. The risk-neutral probabilities are then estimated using equation (3.6). Call options are priced using CV across a range of expiries and moneyness levels. The Matlab function *fzero* is used to set the difference between the Black-Scholes price and the CV price to zero, to solve for the implied volatility.

To demonstrate the impact of the data sample size on the implied volatility estimates consider using CV to estimate the implied volatility of a three month at-the-money (ATM) call option using simulated GBM data, with an analytical implied volatility of 30%. The analysis is done using non-overlapping, overlapping and statistical bootstrapping methods:

- (i) **Non-Overlapping Returns:** Daily stock prices are simulated and used to calculate an array of non-overlapping τ -period returns, where τ is three months. This results in a set of i.i.d. τ -period returns.
- (ii) **Overlapping Returns:** Daily stock prices are simulated and used to calculate an array of overlapping τ -period returns, where τ is three months. Each τ -period return will overlap with the previous τ -period. This results in a set of τ -period returns that are not i.i.d.
- (iii) **Statistical Bootstrapping:** Daily stock prices are simulated and used to calculate an array of daily returns. A random sample of τ returns is selected. This is repeated, with replacement, until there is a sufficient amount of τ -periods.

Throughout this chapter if the underlying data is i.i.d., the statistical bootstrapping method will result in estimates that are more accurate than using non-overlapping returns. Since the data is i.i.d. statistical bootstrapping will only result in a larger sample size, leading to more accurate results.

Figure 3.1 illustrates the impact the size of the return data has on the accuracy of the implied volatility estimates using the CV method while using non-overlapping, overlapping and statistical bootstrapping techniques. These estimates are compared to Crude Monte Carlo (MC) estimation and three standard deviation MC bounds. The amount of simulated GBM daily stock price data starts at twenty-five years increasing in twenty-five year increments to one thousand years. Implied volatility estimates using non-overlapping returns lie within the three standard deviation MC bounds. A large amount of data is required to obtain very accurate results when using non-overlapping returns. Since using overlapping returns results in returns that are not i.i.d, the implied volatility estimates are inaccurate,

thus not lying within the three standard deviation MC bounds. Despite this sampling method generating more τ -period returns, it results in auto-correlated returns which will greatly bias the results. The statistical bootstrapping technique results in reasonable volatility estimates. This is due to the fact that the return data is i.i.d., therefore randomly selecting daily return data will not change the distribution of the returns. The CV method requires an immense amount of return data to produce accurate results. In the event that there is a lack of stock price data and the data is i.i.d., statistical bootstrapping is a good alternative to using non-overlapping returns. The latter is emphasised in figure A.1 found in Appendix A.




Fig. 3.1: Volatility estimates using Canonical Valuation with increasing GBM sample size while applying different sampling techniques.

Since simulated GBM data is used in this analysis, the volatility across a range of moneyness levels is constant. Figure 3.2 illustrates the volatility skew for a three month option using one thousand years of simulated GBM data. The estimates are compared to MC estimates and three standard deviation MC bounds. The volatility estimates for twenty-five to nine-hundred and twenty-five years of simulated data is indicated as *other*. The statistical bootstrapping method results in the most consistent volatility estimates across the range of moneyness levels. The resultant volatility skew using overlapping returns is more stable than using non-overlapping returns. Using non-overlapping return data results in unstable volatility estimates, particularly for in-the-money options.




Fig. 3.2: Volatility skew using Canonical Valuation while applying different sampling techniques using GBM data.

3.3.2 Break-Even Volatility Implementation

Daily stock prices are simulated and implemented to calculate an array of daily returns. Selected τ -periods of daily returns are used to calculate the τ -period profit and loss using equation (3.10). Each τ -period is rebased to ensure that the moneyness levels remain the same for every selected τ -period. The Matlab function *fzero* is used to find the volatility that ensures the average profit and loss for all the selected τ -periods is zero. Figure 3.3 illustrates the average profit and loss, as a result of daily delta hedging, for a three month at-the-money call option using a range of volatilities. The BEV is 30% which is an accurate estimation of the analytical volatility of the simulated data.

To demonstrate the impact of the data sample size on the implied volatility estimates, consider using BEV to estimate the implied volatility for a three month ATM call option using simulated GBM data, with an analytical implied volatility of 30%. The analysis is done using non-overlapping, overlapping and statistical bootstrapping methods as described in section 3.3.1. Less simulated GBM stock price data is used as this method is more computationally intensive. Thus, one hundred years of GBM stock price data is simulated.

Figure 3.4 illustrates the impact the size of the return data has on the accuracy of the implied volatility estimates using the BEV method while using non-overlapping, overlapping and statistical bootstrapping techniques. These estimates are compared to Crude MC estimation and three standard deviation MC bounds. Using non-overlapping returns results in BEV estimates of which the majority lie within the three standard deviation MC bounds. As the sample size is increased,




Fig. 3.3: Profit and loss profile of a three month at-the-money call option using one thousand years of simulated GBM price data with $\sigma = 30\%$.

the implied volatility estimates become increasingly more stable and accurate. Using overlapping returns results in volatility estimates with little variation between estimates but all the estimates do not lie within the three standard deviation MC bounds. The statistical bootstrapping technique results in a very similar estimate profile to that of overlapping returns, where the majority of the estimates lie within the three standard deviation MC bounds. The estimates using the overlapping and statistical bootstrapping techniques do not improve on the estimates using non-overlapping returns. This can also be seen in figure A.2 in Appendix A. If the number of years of simulated stock price data were to be increased, the increase in accuracy of the implied volatility estimates would become more apparent when using the overlapping and statistical bootstrapping methods.

Since simulated GBM data is used in this analysis, the volatility across a range of moneyness levels is constant. Figure 3.5 illustrates the volatility skew for an ATM call option using one hundred years of simulated GBM data. The estimates are compared to MC estimates and three standard deviation MC bounds. The volatility estimates for one to ninety-nine years of simulated data is indicated as *other*. The non-overlapping method results in the volatility estimates which lie within the three standard deviation bounds across the range of moneyness levels. The resultant volatility skew using overlapping and statistical bootstrapping returns also lie within the bounds. All three sampling methods result in estimates that are unstable for options that are deep in- or out-the-money. The estimates are unstable due to the short length of the option, thus there is a sparsity of paths that reach the specified moneyness level in the short time period.




Fig. 3.4: Volatility estimates using Break-Even volatility with increasing GBM sample size while applying different sampling techniques.

BEV is computationally expensive in comparison to CV. This is due to the manner in which the implied volatility is estimated. Since the BEV method requires numerical optimization per strike for each maturity i.e. searching for the implied volatility which results in zero profit and loss, this is very computationally expensive in comparison to the alternative CV. When using the CV method only one numerical optimization is required per maturity to enable implied volatility estimation for a range of strikes. When using the CV method it is important to use τ -periods that are non-overlapping to ensure volatility estimates are accurate. In the case of the CV and BEV methods, if the daily return data is i.i.d., using the statistical bootstrapping method is a good alternative to using non-overlapping returns when there is a lack of stock price data. Term structure estimation is possible using both the CV and BEV techniques. Both the CV and BEV methods are able to estimate volatility skew, on the condition that a large amount of return data is used in the estimation. As the length of the option increases the number of non-overlapping τ -period returns decreases which decreases the accuracy of the volatility estimate. Therefore the longer the length of the option the larger the amount of data is required to obtain accurate results. Both the CV and BEV techniques require a large amount of stock price data for accurate volatility estimation.

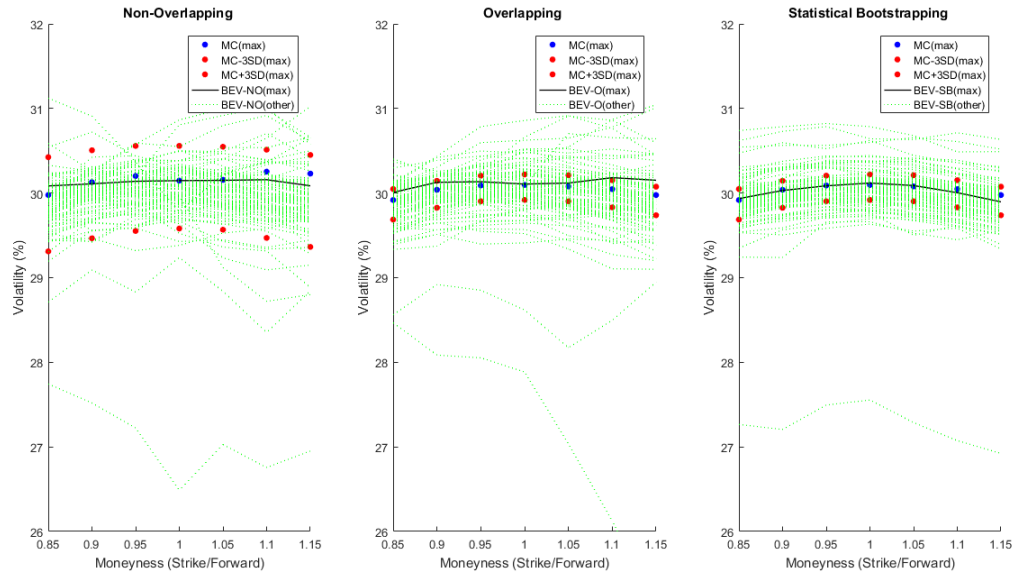


Fig. 3.5: Volatility skew using Break-Even volatility while applying different sampling techniques using GBM data.

3.4 I.I.D. Returns

The assumption that the stock price process follows the Black-Scholes GBM dynamics with a constant volatility has been found to be unrealistic. It has been found that returns are not normally distributed and contain fat tails (Drgulescu and Yakovenko, 2002). For this reason, the Heston and Merton stochastic volatility models are considered and used to simulate stock price data. Since this simulated data may represent a more realistic stock price path, it will be used to assess the ability for the CV and BEV methods to estimate volatility and volatility skew. Overlapping, non-overlapping and bootstrapping techniques are also considered as in Section 3.3.

- (i) **Heston Model:** The Heston model assumes volatility is stochastic therefore the resultant returns series is not i.i.d. See Appendix B.1 for more information on the Heston model. See table B.1 in Appendix B.1 for the parameters used in the Heston model.

To illustrate the effect of non i.i.d. returns on the volatility estimates using the CV method, consider one thousand years of simulated Heston stock price data. Figure 3.6 illustrates the resultant volatility estimates for a three month ATM call option using the CV method. These estimates are compared to Crude MC estimation and three standard deviation MC bounds. The estimates using overlapping and statistical bootstrapping techniques result in unstable volatility estimates, irrespective of the sample size. Thus, these two

sampling techniques are not effective methods to obtain more accurate results by increasing the sampling size. Although the use of non-overlapping returns results in estimates that are within the three standard deviation MC bounds, the estimates are unstable. The latter is emphasised in figure A.3 in Appendix A.




Fig. 3.6: Volatility estimates using Canonical Valuation with increasing Heston sample size while applying different sampling techniques.

Skew estimation can be seen in Figure 3.7 for a three month option. The CV technique using all three sampling techniques does not accurately estimate volatility skew, resulting in unstable estimates. Using non-overlapping returns results in volatility estimates that lie within the three standard deviation MC bounds whereas using non-overlapping and statistical bootstrapping techniques results in estimates which do not lie within these bounds.

To illustrate the effect of non-i.i.d. returns on the volatility estimates using the BEV method, consider one hundred years of simulated Heston stock price data. Figure 3.8 illustrates the resultant volatility estimates for a three month ATM call option using the BEV method with non-overlapping, overlapping and statistical bootstrapping techniques. These estimates are compared to Crude MC estimation and three standard deviation MC bounds. The estimates using all three sampling techniques result in unstable volatility estimates, irrespective of the sample size. Thus, the BEV method is not capable of estimating volatility for stock price return data that is not i.i.d. The latter is emphasised in Figure A.4 in Appendix A, where the MC estimates are the same for all three sampling techniques.

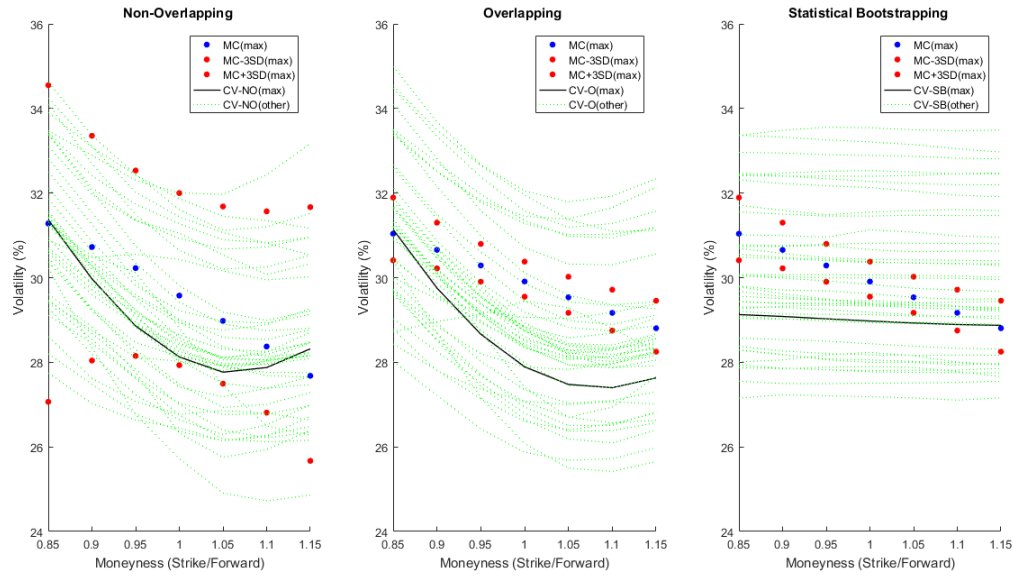


Fig. 3.7: Volatility skew using Canonical Valuation while applying different sampling techniques using Heston data.

The volatility estimates using the BEV method for a three month option at a range of moneyness levels is illustrated in Figure 3.9. The estimates using all three sampling techniques result in unstable volatility estimates for a range of moneyness levels. Thus, the BEV method is unable to recover a volatility skew for return data that is not i.i.d.

The Heston model results in a returns series that is not i.i.d. Both CV and BEV methods do not perform well when the returns series is not i.i.d., resulting in unstable volatility estimates. Thus, the CV and BEV techniques may only be employed when the series of stock price returns is in-fact i.i.d.

- (ii) **Merton Model:** The Merton model adds random jumps to the GBM dynamics to account for random stock price shocks. The resultant returns series remains i.i.d. See Appendix B.2 for more information on the Merton model. See table B.2 in Appendix B.2 for the parameters used in the Merton model.

To illustrate the effect of market data following the Merton model on the volatility estimates using CV and BEV methods, consider one thousand years of simulated Merton stock price data. Non-overlapping, overlapping and statistical bootstrapping techniques are used in this analysis. Figure 3.10 illustrates the volatility estimates using the CV method for a range of sample sizes. The volatility estimates are compared to MC simulation and three standard deviation MC bounds. Using non-overlapping returns results in stable estimates that lie within the three standard deviation bounds. As the sam-




Fig. 3.8: Volatility estimates using Break-Even volatility with increasing Heston sample size while applying different sampling techniques.

ple size is increased there is less variation in the estimates. Using overlapping returns results in unstable estimates due to the autocorrelation in the returns. Using the statistical bootstrapping technique results in stable volatility estimates. The variability in the estimates is decreased as the sample size increases thus, the statistical bootstrapping technique is a good alternative to increase the sample size. The latter is evident in Figure A.5 in Appendix A.

The volatility estimates using the CV method for a three month option at a range of moneyness levels is illustrated in Figure 3.11. The estimates using all three sampling techniques result in a volatility skew where the non-overlapping and statistical bootstrapping techniques result in estimates which lie within the three standard deviation MC bounds.

Figure 3.12 illustrates the volatility estimates using the BEV method for a three month ATM call option with a range of sample sizes using one hundred years of simulated Merton data. Using non-overlapping returns results in unstable volatility estimates of which the majority lie within the three standard deviation MC bounds. As the sample size is increased, the volatility estimates become more stable. Using statistical bootstrapping and overlapping techniques to increase the sample size does not result in an increase of the accuracy of the estimates. This is due to the immense amount of data that the BEV method requires to produce accurate results. The latter is emphasised in figure A.6 in Appendix A. If the number of years of stock price data were to increase, the statistical bootstrapping and overlapping sampling techniques would im-

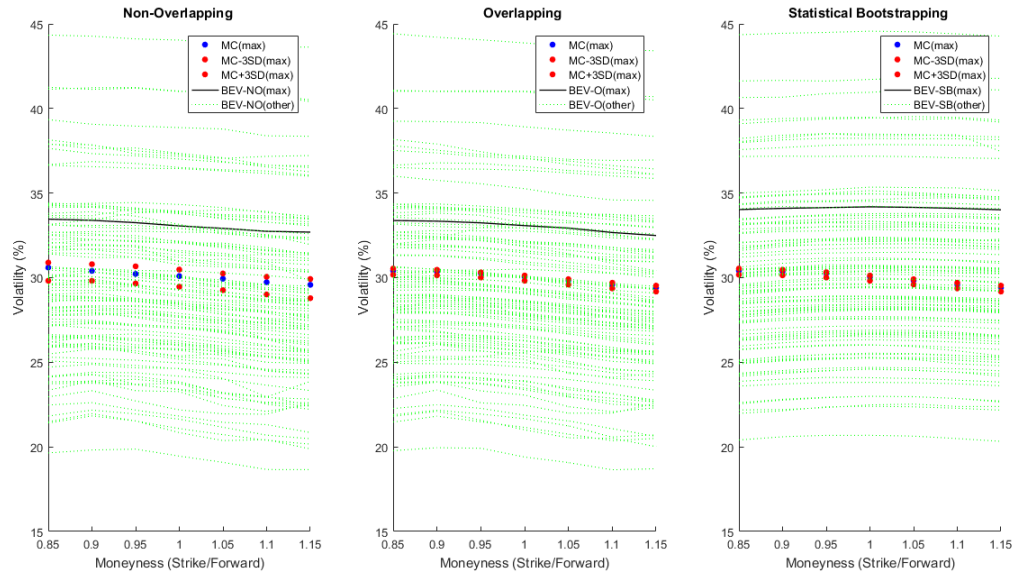


Fig. 3.9: Volatility skew using Break-Even volatility while applying different sampling techniques using Heston data.

prove on the accuracy of the non-overlapping estimates.

Figure 3.13 illustrates the volatility estimates for a range of moneyness levels for a three month option using simulated Merton stock price data. Using non-overlapping returns results in volatility estimates which lie within the three standard deviation MC bounds. Using overlapping or statistical bootstrapping techniques results in estimates which also lie within these bounds. Thus, given i.i.d. return data and a large enough sample size the BEV method is able to recover a volatility skew.

The BE and CV methods result in satisfactory estimates provided the sample size is large enough and stock price returns series is i.i.d. If the stock price returns are i.i.d., the statistical bootstrapping method can be used to increase the sample size and is a good alternative to using non-overlapping returns when using the CV and BEV methods. Term structure estimation is possible using CV and BEV methods. The disadvantage of the estimation techniques described is that the length of term structure is limited to the amount of available stock price data. Thus, longer maturities will result in unstable volatility estimates because of the lack of stock price data. Skew estimation is possible, as seen in the analysis using the Merton model to simulate stock price data, thus both the CV and BEV models are able to recover a skew when the data is i.i.d. Another drawback of these estimation techniques is that they only use equity stock price market data and no equity option market data. It is not possible to estimate statistical confidence intervals for the




Fig. 3.10: Volatility estimates using Canonical Valuation with increasing Merton sample size while applying different sampling techniques.

estimates. Lastly, these estimation techniques require the stock price returns series to be i.i.d. The i.i.d. condition may be unrealistic for market stock price returns.

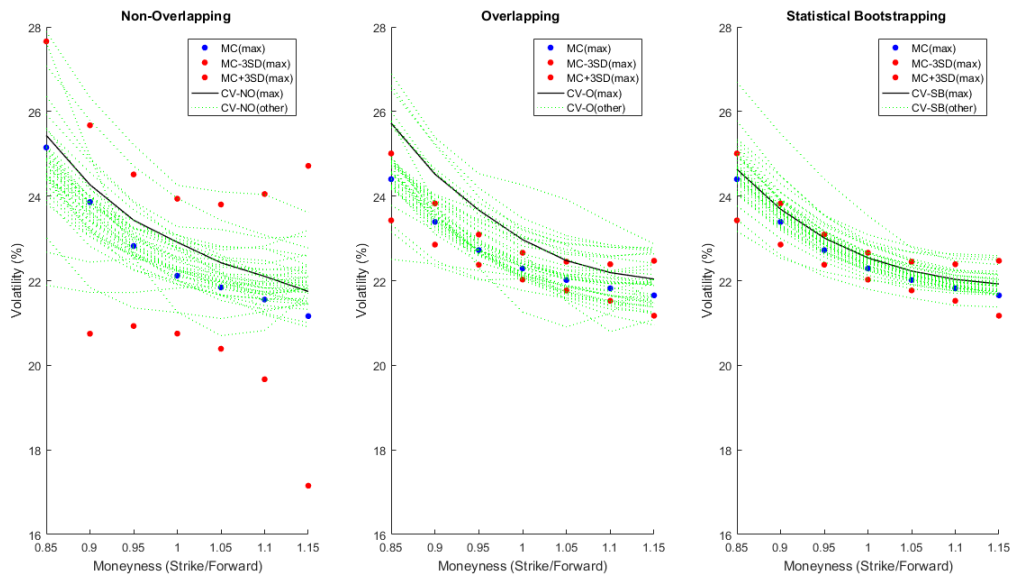


Fig. 3.11: Volatility skew using Canonical Valuation while applying different sampling techniques using Merton data.




Fig. 3.12: Volatility estimates using Break-Even volatility with increasing Merton sample size while applying different sampling techniques.

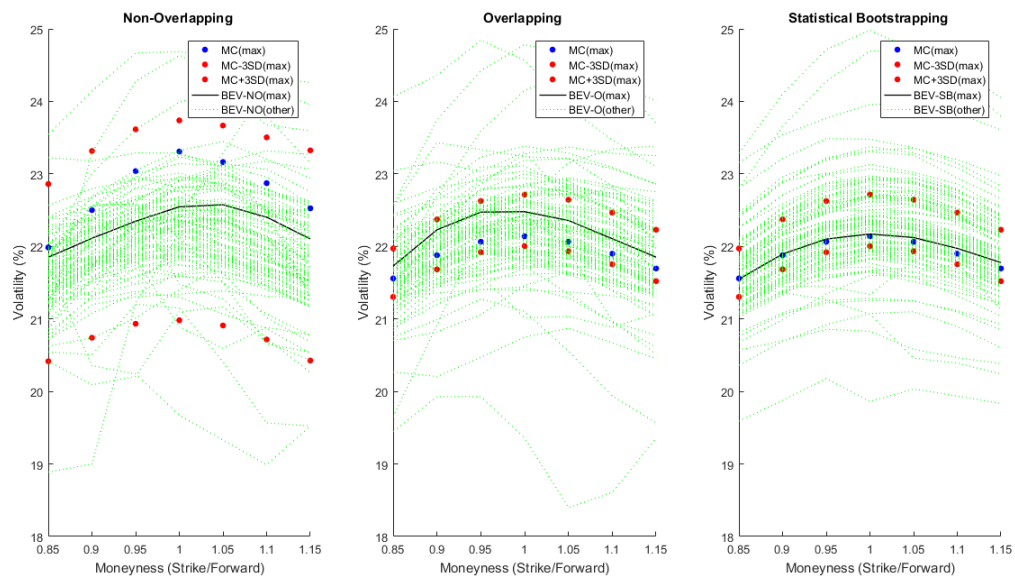


Fig. 3.13: Volatility skew using Break-Even volatility while applying different sampling techniques using Merton data.

Chapter 4

Potential Solutions Using a Local Volatility Model

The estimation procedures outlined in this chapter attempts to recover option prices at a given point in time. The models described are based on the SAFEX parametrization proposed by [Kotzé *et al.* \(2015\)](#). The estimation techniques used can be categorised as deterministic volatility models requiring a snapshot of option implied volatility data which directly models local volatility. Local volatility can be described as the instantaneous volatility at a specific point in time. Thus, local volatility is time and state dependent ([Kotzé *et al.*, 2015](#)). Moreover, given a set of European stock prices, there exists a unique risk neutral diffusion parameter, the local variance, which can be used to generate these prices ([Gatheral and Lynch, 2004](#)).

Simulation techniques are used to extrapolate local volatilities for maturities that are unobservable in the market. Thus, the dynamics of the local volatility for maturities beyond what is observable in the market become partially stochastic in nature and are based on the assumptions used in the three proposed models described in this chapter. Since local volatility is described as an instantaneous quantity, the proposed simulation extrapolation can be viewed as generating a partially stochastic volatility model for the asset price. The local volatility term structure is effectively reproduced and extended to unobservable maturities by extrapolation simulations of the existing local volatility surface. It should be emphasised that the existing local volatility surface is derived from a sample of market observable implied volatilities for a given range of maturities.

4.1 Local Volatility

Time-varying deterministic volatility models aim to create a volatility surface by fitting a deterministic model to an underlying stochastic asset-price process ([Flint *et al.*, 2014](#)). A popular non-parametric deterministic model considered is the local volatility model. This is a model in which volatility is a deterministic function of two variables, time (t) and the asset price at each time (S_t). Under the risk-neutral measure, the standard Black-Scholes framework assumes volatility is constant and asset prices follow a geometric Brownian motion:

$$dS_t = rS_t dt + \sigma S_t dW_t, \quad (4.1)$$

where r is the risk-free rate and $\{W_t\}$ denotes a standard Brownian motion under the risk neutral measure \mathbb{Q} (Linghao, n.d.). In reality the assumption of constant volatility has been found to be inaccurate. Dupire has shown that the constant volatility in equation (4.1) can rather be modelled using the deterministic function $\sigma(t, S_t)$ which consequently adds some randomness to the volatility process (Davis, 2011). The randomness is added due to the dependence of the volatility on the stock price process, which is driven by Brownian motion. This results in a stock price process that will no longer have a constant volatility. The dynamics are described by:

$$dS_t = rS_t dt + \sigma(t, S_t)S_t dW_t. \quad (4.2)$$

Since there is still only one source of randomness in the model, the local volatility model is complete ensuring unique prices (Kamp, 2009). The local volatility function $\sigma(t, S_t)$ is consistent with market prices of options on the underlying S such that $\sigma(t, S_t)$ is the markets view of future volatility at t with stock price S_t .

The resultant local volatility may not be realised. Local volatility can be viewed as the average instantaneous volatility for a specified time and stock price hence the name *local*, whereas Black-Scholes implied volatility can be viewed as the overall *global* future volatility for the lifetime of the option (Derman et al., 1996). The local volatility can be extracted analytically from European option prices with Dupire's formula for local volatility (Kotzé et al., 2015):

$$\sigma_L(T, K) = \sqrt{\frac{\frac{\partial}{\partial T} C(S_t, T, K) + rK \frac{\partial}{\partial K} C(S_t, T, K)}{\frac{1}{2} K^2 \frac{\partial^2}{\partial K^2} C(S_t, T, K)}}, \quad (4.3)$$

where $C(S_t, T, K)$ is the call price today and $\sigma_L(T, K)$ is continuous. Local volatility can be uniquely computed for a range of strike/asset prices and expirations using equation (4.3). It is assumed that the call price is twice differentiable with respect to the strike and once with respect to expiration. Since liquid option prices are only known for a discrete set of expiries and strikes, price data can be interpolated to create a smoother price surface. It should be ensured that the option price function is convex with respect to strike to guarantee the second derivative in the denominator is positive and real. The derivatives can be approximated numerically using finite difference techniques (see appendix (C.2)) (Kotze et al., 2014). It is suggested by Kotze et al. (2014) to use one to ten index points for the change in strike for the Top40 Index and one to ten basis points for the change in time. A percentage of the strike price could also be used while ensuring that it is small enough. There are a few areas which may result in errors using Dupire's equation. Firstly, small absolute errors in the approximation of the second derivative in the denominator are magnified by the squared strike price. Secondly, this derivative will result in very small values for options deep in- or out-the-money, this is especially prevalent for shorter termed options. An alternative method is to calculate local volatility in terms of market implied volatilities (Kotze and Oosthuizen, 2015):

$$\sigma_L(S_t, T, K) = \sqrt{\frac{\sigma_{imp}^2 + 2\sigma_{imp}T \left(\frac{\partial \sigma_{imp}}{\partial T} + rK \frac{\partial \sigma_{imp}}{\partial K} \right)}{1 + 2d_1 K \sqrt{T} \frac{\partial \sigma_{imp}}{\partial K} + K^2 T \left(d_1 d_2 \left(\frac{\partial \sigma_{imp}}{\partial K} \right)^2 + \sigma_{imp} \frac{\partial^2 \sigma_{imp}}{\partial K^2} \right)}}, \quad (4.4)$$

where

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + rT + \frac{\sigma_{imp}^2 T}{2}}{\sigma_{imp}\sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \sigma_{imp}\sqrt{T}. \quad (4.5)$$

The derivation of equation (4.4) is included in Appendix C.1. This method is more stable due to the elimination of the single standing dual Gamma in the denominator. Small errors in the second derivative calculation will no longer lead to large overall errors, as they are now part of a summation rather than standing alone (Kotze *et al.*, 2014).

4.2 Forecasting Using the Local Volatility Model

In this section three methods are proposed as potential techniques for obtaining long-term equity implied volatility estimates. Local volatilities can be estimated using equation (4.4) for a range of strikes and expiries where the maximum expiry is limited to what is observable in the market. The following three methods outlined are used to obtain implied volatilities for expiries longer than what is available in the market. Each method contains different assumptions about how future local volatility proceeds beyond what is observable in the market. In order to demonstrate these three methods consider an implied volatility surface with a range of strikes K and a range of expiries $[t, t + dt, t + 2dt, \dots, T]$, where dt is the interval between consecutive expiries. This implied volatility surface is applied to equation (4.4) resulting in local volatilities for the same range of strikes and expiries.

Given the local volatility surface, one of the three methods is applied to the local volatility data. Each of these methods comprises of an iterative process to forecast local volatility one step into the future. Once this local volatility is calculated, Milstein discretization is applied to equation (4.2) to obtain the stock price one time step into the future. This is performed iteratively until the desired number of stock prices are obtained. Monte Carlo simulation is performed with the resultant stock prices to obtain an option price. This option price is then used to extract the Black-Scholes implied volatility. It is desirable to obtain a distribution of forecasts, rather than a point estimate. Each method can be repeated to obtain a distribution of implied volatility forecasts.

To illustrate each method, JSE Top40 implied volatility data on the 20th of November 2017 is interpolated across market observable expiries such that δt is one business day. It is assumed that implied volatility prior to the first market expiry is flat. The implied volatility surface is illustrated in Figure 4.1(a). Equation (4.4) is used to calculate the local volatility surface, with the longest observable expiry from the JSE Top40 data being $T = 1.7$ years and the shortest three months. The local volatility surface is illustrated in Figure 4.1(b).

4.2.1 Method 1: Assume Constant Gradient Ratio

In this algorithm it is assumed that the change in local volatility corresponding to expiry $(T - dt)$ to T persists into the future.

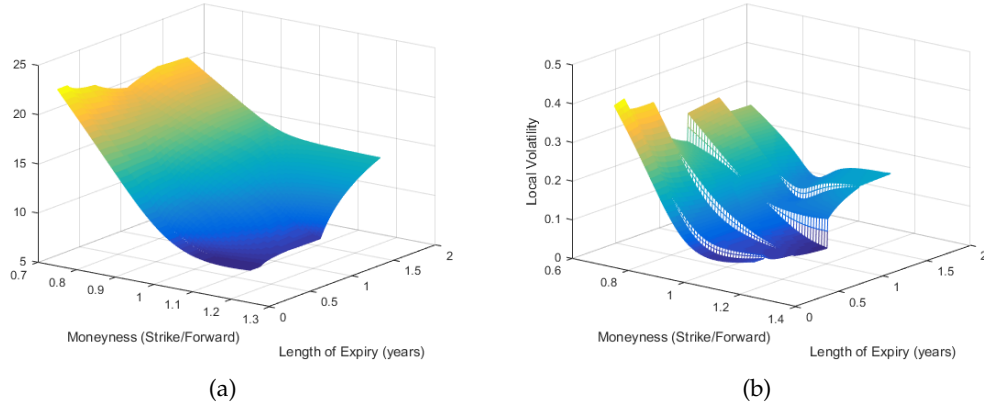


Fig. 4.1: JSE Top40 volatility surfaces on the 20/11/17. (a) JSE Top40 implied volatility surface on the 20/11/17. (b) JSE Top40 local volatility surface on the 20/11/17.

Algorithm 1: Gradient from $(T - dt)$ to T

Input : Number of stock price simulations (N)

Number of stock price paths (M)

Market local volatilities (σ_L)

Maturities within local volatility surface ($\underline{\mathbf{T}}$)

Strikes within local volatility surface ($\underline{\mathbf{K}}$)

Output: Stock price paths (\mathbf{S})

- 1 Create empty matrix of size $[N, M]$ and set this equal to $\widetilde{\sigma}_L$.
 - 2 Create empty matrix of size $[N + 1, M]$ and set this equal to \mathbf{S} .
 - 3 Set first row of \mathbf{S} equal to \mathbf{S}_0 .
 - 4 **for** $i = 2$ **to** N **do**
 - 5 **if** $\underline{\mathbf{T}}(i) \leq \max \underline{\mathbf{T}}$; **then**
 - 6 Interpolate local volatility surface at the minimum of [\mathbf{S}_{i-1} or max $\underline{\mathbf{K}}$] and $\underline{\mathbf{T}}(i)$. Set the first row of $\widetilde{\sigma}_L$ equal to result.
 - 7 **else**
 - 8 Interpolate local volatility surface at the minimum of [$\mathbf{S}(i-1)$ or max $\underline{\mathbf{K}}$] and $\underline{\mathbf{T}}_X$. Set σ_{L1} equal to result.
 - 9 Interpolate local volatility surface at the minimum of [$\mathbf{S}(i-1)$ or max $\underline{\mathbf{K}}$] and $\underline{\mathbf{T}}_{X-1}$. Set σ_{L2} equal to result.
 - 10 Set row (i) of $\widetilde{\sigma}_L$ equal to $\widetilde{\sigma}_L(i-1)\sigma_{L1}/\sigma_{L2}$.
 - 11 **end**
 - 12 Use Milstein to calculate $\mathbf{S}(i)$:
 - 13
$$\mathbf{S}(i) = S(i-1) + rS(i-1)\delta T + \widetilde{\sigma}_L(i-1)S(i-1)\sqrt{\delta T}Z(i) + \frac{1}{2}\widetilde{\sigma}_L(i-1)^2S(i-1)(Z(i)^2 - 1)\delta T.$$
 - 14 **end**
-

Figure 4.2 illustrates the average implied volatilities for one hundred simulations of Algorithm 1 with one hundred thousand stock price paths of at-the-money options for observable market expiries as well as expiries extending to thirty years. The average implied volatilities for market observable expiries lies within the 5% and 95% percentile interval and does not vary much from the market values. The variation in the forecasts are fairly narrow and remains a consistent distance from the average implied volatility as the expiry is increased, as seen by the 5% and 95% percentiles. As the option expiration is increased, the variation between 5% and 95% percentiles are increased. Estimates lying further in the future are more uncertain than estimates that are closer to the present time. Implied volatilities with expiries close to the longest observable market expiry are sensitive to the change in implied volatilities from $(T - dt)$ to T .




Fig. 4.2: Implied volatility estimates using Method 1 for JSE Top40 at-the-money European call.

4.2.2 Method 2: Random Selection

In this algorithm a random expiry t_R is selected. It is assumed that the change in local volatility is equal to that from t_{R-1} to t_R .

Figure 4.3 illustrates the average implied volatilities for one hundred simulations of Algorithm 2 with one hundred thousand stock price paths of at-the-money options for observable market expiries as well as expiries extending to thirty years. The average implied volatilities for market observable expiries lies within the 5% and 95% percentile interval and does not vary much from the market values. The variation in the forecasts is fairly wide, in comparison to Method 1. As the expiry

Algorithm 2: Random selection

Input : Number of stock price simulations (N)
Number of stock price paths (M)
Market local volatilities (σ_L)
Maturities within local volatility surface ($\underline{\mathbf{T}}$)
Strikes within local volatility surface ($\underline{\mathbf{K}}$)

Output: Stock price paths (\mathbf{S})

- 1 Create empty matrix of size $[N, M]$ and set this equal to $\widetilde{\sigma}_L$.
- 2 Create empty matrix of size $[N + 1, M]$ and set this equal to \mathbf{S} .
- 3 Set first row of \mathbf{S} equal to \mathbf{S}_0 .
- 4 **for** $i = 2$ **to** N **do**
- 5 **if** $\underline{\mathbf{T}}(i) \leq \max \underline{\mathbf{T}}$; **then**
- 6 Interpolate local volatility surface at the minimum of $[\mathbf{S}_{i-1}$ or $\max \underline{\mathbf{K}}$] and $\underline{\mathbf{T}}(i)$. Set the first row of $\widetilde{\sigma}_L$ equal to result.
- 7 **else**
- 8 Select a random expiry $\underline{\mathbf{T}}_R$.
- 9 Interpolate local volatility surface at the minimum of $[\mathbf{S}(i-1)$ or $\max \underline{\mathbf{K}}$] and $\underline{\mathbf{T}}_R$. Set σ_{L_1} equal to result.
- 10 Interpolate local volatility surface at the minimum of $[\mathbf{S}(i-1)$ or $\max \underline{\mathbf{K}}$] and $\underline{\mathbf{T}}_{R-1}$. Set σ_{L_2} equal to result.
- 11 Set row (i) of $\widetilde{\sigma}_L$ equal to $\widetilde{\sigma}_L(i-1)\sigma_{L_1}/\sigma_{L_2}$.
- 12 **end**
- 13 Use Milstein to calculate $\mathbf{S}(i)$:
- 14
$$\mathbf{S}(i) = S(i-1) + rS(i-1)\delta T + \widetilde{\sigma}_L(i-1)S(i-1)\sqrt{\delta T}Z(i) + \frac{1}{2}\widetilde{\sigma}_L(i-1)^2S(i-1)(Z(i)^2 - 1)\delta T.$$
- 15 **end**

increases the variation in the forecasts increases, as seen by the 5% and 95% percentiles. This is expected due to the assumption that the change in local volatility is equal to that from a random observable expiry t_{R-1} to t_R . Because random volatilities are selected in this algorithm, the implied volatility estimates are not sensitive to the volatilities at $(T - dt)$ and T .

4.2.3 Method 3: Random Path-Wise Selection

In this algorithm a random expiry t_R is selected. It is assumed that the change in local volatility at $(t + 1)$ is equal to that from t_{R-1} to t_R . The proceeding change in local volatility is equivalent to that from t_R to t_{R+1} and so forth until T is reached. Once the latter is reached another random expiry is selected and the process is continued.

Figure 4.4 illustrates the average implied volatilities for one hundred simulations of Algorithm 3 with one hundred thousand stock price paths for at-the-money options for observable market expiries as well as expiries extending to thirty years. The average implied volatilities for market observable expiries lies within the 5%

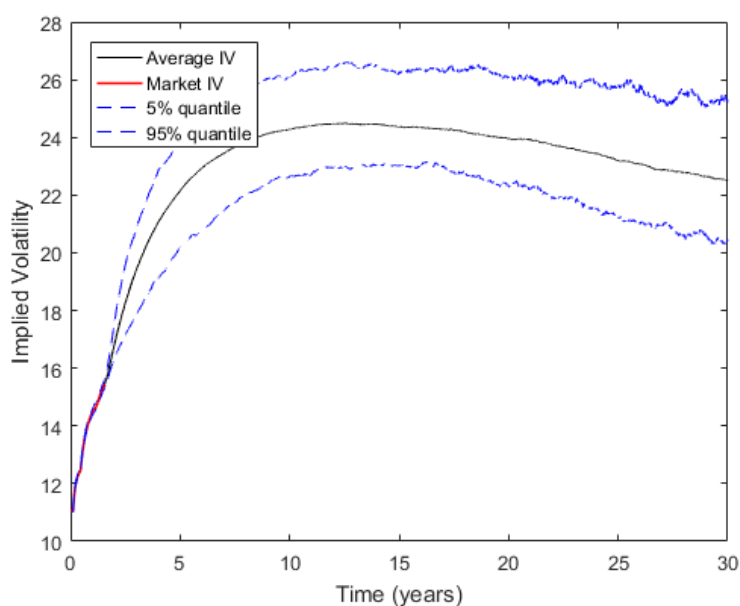


Fig. 4.3: Implied volatility estimates using Method 2 for JSE Top40 at-the-money European call.

and 95% percentile interval and does not vary much from the market values. The variation in the forecasts is fairly narrow.

The three methods outlined produce market-consistent estimates as well as reasonable long-term implied volatility estimates. Method 2 has a larger variation between the 5% and 95% confidence interval, which is expected, in comparison to method 1 and 3. This is expected as a random expiry is used to estimate the change in local volatility for each iteration, which will result in more variability between consecutive estimates. It is possible to estimate implied volatility for a range of moneyness levels. The advantages of using these simulation-based extrapolation techniques are firstly, that it is possible to estimate implied volatilities for maturities which are not available in the market. This was not possible for both the traditional estimation techniques and non-parametric estimation techniques which were outlined. Secondly, it is possible to estimate statistical confidence intervals for the resultant implied volatility estimates. This is important because there is variability in the future outcomes for option implied volatility, thus a point estimate is very unreliable. And thirdly, market implied volatility is used to extrapolate implied volatilities that are not available in the market, rather than using stock price data. The disadvantage of using these extrapolation techniques is that the change in the proceeding local volatility should be assumed to follow a chosen methodology for each time step.

Algorithm 3: Path-wise selection

Input : Number of stock price simulations (N)
Number of stock price paths (M)
Market local volatilities (σ_L)
Maturities within local volatility surface (\underline{T})
Strikes within local volatility surface (\underline{K})

Output: Stock price paths (S)

- 1 Create empty matrix of size $[N, M]$ and set this equal to $\widetilde{\sigma}_L$.
- 2 Create empty matrix of size $[N + 1, M]$ and set this equal to S .
- 3 Set first row of S equal to S_0 .
- 4 Set t equal to T
- 5 **for** $i = 2$ **to** N **do**
- 6 **if** $\underline{T}(i) \leq \max \underline{T}$; **then**
- 7 Interpolate local volatility surface at the minimum of $[S_{i-1}$ or \max
 $\underline{K}]$ and $\underline{T}(i)$. Set the first row of $\widetilde{\sigma}_L$ equal to result.
- 8 **else**
- 9 Set $t = t + \delta \underline{T}$
- 10 **if** $t > T$ **then**
- 11 Select a random expiry $t = \underline{T}_R$.
- 12 **end**
- 13 Interpolate local volatility surface at the minimum of $[S(i-1)$ or
 $\max \underline{K}]$ and t . Set σ_{L_1} equal to result.
- 14 Interpolate local volatility surface at the minimum of $[S(i-1)$ or
 $\max \underline{K}]$ and $(t - \delta \underline{T})$. Set σ_{L_2} equal to result.
- 15 Set row (i) of $\widetilde{\sigma}_L$ equal to $\widetilde{\sigma}_L(i-1)\sigma_{L_1}/\sigma_{L_2}$.
- 16 **end**
- 17 Use Milstein to calculate $S(i)$:
- 18
$$S(i) = S(i-1) + rS(i-1)\delta T + \widetilde{\sigma}_L(i-1)S(i-1)\sqrt{\delta T}Z(i) + \frac{1}{2}\widetilde{\sigma}_L(i-1)^2S(i-1)(Z(i)^2 - 1)\delta T.$$
- 19 **end**




Fig. 4.4: Implied volatility estimates using Method 3 for JSE Top40 at-the-money European call.

Chapter 5

Conclusion

Estimating long-term implied volatility for equity derivatives is difficult in South Africa, where 'long term' refers to expiries from ten to thirty years. This due to the lack of liquid tradable long-term derivative instruments. Thus, it is important to be able to estimate and extrapolate fair, market-consistent volatility surfaces. This dissertation evaluates different techniques used to estimate volatilities for a range of moneyness levels and expiries. The methods evaluated include statistical and time-series techniques, non-parametric techniques and simulation-based extrapolation techniques which forecast implied volatilities making use of a local volatility model.

The statistical and time-series techniques include historical volatility, EWMA and GJR-GARCH(1,1) which only need a time-series of historical stock price data. These three methods cannot estimate volatility skew or statistical confidence intervals for the volatility estimates. These methods result in realised volatility estimates which need to be scaled by an estimated IVHV ratio to obtain implied volatility estimates. Term structure estimation is possible for these methods but maximum possible expiry for historical and EWMA techniques is limited to the amount of stock price data available. GJR-GARCH(1,1) is able to forecast long-term volatilities but the volatility estimates for longer termed maturities are flat.

The non-parametric techniques include Canonical Valuation and Break-Even volatility which only use a time-series of historical stock price data. Both of these methods estimate implied volatility with a maximum expiry limited to the amount of stock price data available. These methods require an immense amount of historical data to produce reasonable results and are computationally expensive. These methods are able to estimate the volatility surface across a range of expiries and moneyness levels with the ability to estimate a reasonable volatility skew. Two sampling techniques, using overlapping returns and statistical bootstrapping, are considered to increase the size of the return data. The use of statistical bootstrapping results in implied volatility estimates which improve on the estimates using non-overlapping returns. Both CV and BEV methods require the stock price return data to be independently and identically distributed.

Three potential models are proposed, using Dupire's local volatility model, to estimate implied volatilities for expiries beyond what is available in the market. Each method contains different assumptions about how future volatility proceeds beyond what is observable in the market. These methods require a snapshot of option implied volatilities which are used for the estimation of local volatilities.

These methods directly estimate implied volatilities for a range of expiries and moneyness levels. It is possible to calculate statistical confidence intervals for the resultant volatility estimates. All three methods are able to recover market implied volatilities accurately.

Further research should be done on each of these methods to ensure that the long-term implied volatility is market-consistent and fair. A further study should be done on how to estimate long-term implied volatilities using only option market data.

Bibliography

- AL-Najjar, D. M. (2016). Modelling and estimation of volatility using arch/garch models in Jordan's stock market, *Asian Journal of Finance & Accounting* **8**(1): 152–167.
- Alcock, J. and Gray, P. (2005). Dynamic, nonparametric hedging of European style contingent claims using canonical valuation, *Finance Research Letters* **2**(1): 41–50.
- Chai, T. and Draxler, R. R. (2014). Root mean square error (rmse) or mean absolute error (mae)?—arguments against avoiding rmse in the literature, *Geoscientific model development* **7**(3): 1247–1250.
- Davis, M. H. (2011). The Dupire formula, *Imperial College London, Finite Difference Methods Course material*.
- Derman, E., Kani, I. and Zou, J. Z. (1996). The local volatility surface: Unlocking the information in index option prices, *Financial Analysts Journal* **52**(4): 25–36.
- Drgulescu, A. A. and Yakovenko, V. M. (2002). Probability distribution of returns in the Heston model with stochastic volatility, *Quantitative Finance* **2**(6): 443–453.
- Dupire, B. (2006). Fair skew: Break-even volatility surface, *Bloomberg LP*.
- Eberly, D. (2008). Derivative approximation by finite differences, *Magic Software, Inc.*
- Ederington, L. and Guan, W. (2006). Measuring historical volatility.
- Eraker, B. (2008). The volatility premium, *Manuscript*.
- Figlewski, S. (1994). Forecasting volatility using historical data.
- Flint, E. J., Chikurunhe, F. and Seymour, A. J. (2012). (un) modelling the volatility surface: Valuing South African volatility surfaces via risk-neutral historic return distributions.
- Flint, E., Ochse, E. and Polakow, D. (2014). Estimating long-term volatility parameters for market-consistent models, *South African Actuarial Journal* **14**(1): 19–72.
- Gatheral, J. and Lynch, M. (2004). Lecture 1: Stochastic volatility and local volatility, *Case Studies in Financial Modeling Notes, Courant Institute of Mathematical Sciences*.

- Gray, P. and Newman, S. (2005). Canonical valuation of options in the presence of stochastic volatility, *Journal of Futures Markets* **25**(1): 1–19.
- Hansen, P. R. and Lunde, A. (2005). A forecast comparison of volatility models: does anything beat a garch (1, 1)?, *Journal of applied econometrics* **20**(7): 873–889.
- Kamp, R. (2009). *Local volatility modelling*, Master's thesis, University of Twente.
- Korkmaz, T. and Aydin, K. (2002). Using evma and garch methods in var calculations: Application on ise-30 index, *ERC/METU 6. International Conference in Economics*, pp. 11–14.
- Kotze, A. and Oosthuizen, R. (2015). Pricing jse exotic can-do options: Monte carlo simulation.
- Kotze, A., Oosthuizen, R. and Pindza, E. (2014). Local volatility modeling of jse exotic can-do options.
- Kotzé, A., Oosthuizen, R. and Pindza, E. (2015). Implied and local volatility surfaces for south african index and foreign exchange options, *Journal of Risk and Financial Management* **8**(1): 43–82.
- Ladokhin, S. (2009). Volatility modeling in financial markets, *VU University Amsterdam. Faculty of Sciences, Business Mathematics and Informatics. De Boelelaan 1081a* **1**: 081.
- Linghao, Y. (n.d.). *The pricing of options with jump diffusion and stochastic volatility*, PhD thesis.
- Liu, H.-C. and Hung, J.-C. (2010). Forecasting s&p-100 stock index volatility: The role of volatility asymmetry and distributional assumption in garch models, *Expert Systems with Applications* **37**(7): 4928–4934.
- Reider, R. (2009). Volatility forecasting i: Garch models, *New York* .
- Stutzer, M. (1996). A simple nonparametric approach to derivative security valuation, *The Journal of Finance* **51**(5): 1633–1652.
- Vieira, J. B. *et al.* (2012). Modeling volatility: an assessment of the value at risk approach, *PhD thesis*.
- Wang, D.-M. (2010). Monte Carlo Simulations for Complex Option Pricing, *PhD thesis, The University of Manchester (United Kingdom)*.

Appendix A

Supporting Figures

The following figures illustrates the implied volatility estimates using the same specified simulated data with MC estimates and three standard deviation MC bounds that are the same for all the three sampling techniques.




Fig. A.1: Volatility estimates using Canonical Valuation with GBM data while applying different sampling techniques with the same Monte Carlo estimates.




Fig. A.2: Volatility estimates using Break-Even volatility with GBM data while applying different sampling techniques with the same Monte Carlo estimates.




Fig. A.3: Volatility estimates using Canonical Valuation with Heston data while applying different sampling techniques with the same Monte Carlo estimates.




Fig. A.4: Volatility estimates using Break-Even with Heston data while applying different sampling techniques with the same Monte Carlo estimates.




Fig. A.5: Volatility estimates using Canonical Valuation with Merton data while applying different sampling techniques with the same Monte Carlo estimates.




Fig. A.6: Volatility estimates using Break-Even volatility with Merton data while applying different sampling techniques with the same Monte Carlo estimates.

Appendix B

Alternative Models

B.1 Heston Model

The Heston model is a stochastic volatility model. This model assumes the stock price volatility follows the Cox-Ingersoll-Ross process, which is random process. The stock price process follows the following underlying process:

$$dS_t = \mu S_t dt + \sqrt{v_t} S_t dW_t^1, \quad (\text{B.1})$$

where dv_t follows the following underlying process:

$$dv_t = \kappa(\theta - v_t)dt + \sigma\sqrt{v_t} dW_t^2, \quad (\text{B.2})$$

with

$$dW_t^2 = \rho dW_t^1 + \sqrt{1 - \rho^2} dZ_t. \quad (\text{B.3})$$

where $0 \leq \rho \leq 1$ is the correlation coefficient for the two Weiner processes W_t^1 and W_t^2 , and W_t^1 is independent of the Weiner process Z_t . The instantaneous variance is denoted by v_t and the long term mean of v_t is denoted by θ . The variable σ is the volatility of the volatility and κ is the rate of mean reversion to θ ([Drgulescu and Yakovenko, 2002](#)).

Parameter	Value
S_0	100
r	0.1
σ	0.1
dt	1/252
v_0	0.09
κ	0.15
θ	0.09
ρ	-0.9

Tab. B.1: Heston parameters used in Chapter 3.

Parameter	Value
S_0	100
r	0.1
σ	0.2
dt	1/252
σ_p	0.1
μ_p	-0.05
λ	1

Tab. B.2: Merton parameters used in Chapter 3.

B.2 Merton Model

The Merton model models the stock price dynamics as a jump diffusion process. The stock price returns composes of the sum of three components, the liner drift, Brownian motion and the Poisson process which adds random jumps to the stock price process to account for stock price shocks:

$$dS_t = (r - \lambda\kappa)S_t dt + \sigma S_t dW_t + Jdq, \quad (\text{B.4})$$

where dq is the independent Poisson process with intensity λ such that the probability that dq is equal to one is λdt or zero. The variable κ is the magnitude of the jump size over the time period dt . The jumps are assumed to be log-normally distributed and i.i.d. (Linghao, n.d.).

The following expression is the Euler discretization of equation (B.4) and can be used to generate M independent stock price paths where $S_j(t_i)$ for $j = 1, \dots, M$ and $x(t)$ is the log price (Wang, 2010):

$$x_j(t_{i+1}) = x_j(t_i) + (r - \lambda\kappa - \frac{1}{2}\sigma^2)dt + \sigma Z_j(t_i)\sqrt{dt} + \mu_p \rho_j(t_i) + \sigma_p \sqrt{\rho_j(t_i)} \bar{Z}_j(t_i) \quad (\text{B.5})$$

where

$$\kappa = \exp(\mu_p + \frac{1}{2}\sigma_p^2) - 1 \quad (\text{B.6})$$

$$\rho \sim \text{Poisson}(\lambda dt) \quad (\text{B.7})$$

Appendix C

Local Volatility

C.1 Derivation of the Local Volatility Formula

Following [Kamp \(2009\)](#) the derivation of the local volatility function using implied volatilities is as follows.

Firstly a change of variables is made to Dupire's equation such that the call option price is a function of other variables:

$$y = \ln \left(\frac{K}{S_0} \right) + \int_{t_0}^T (q_s - r_s) ds \quad (\text{C.1})$$

$$w = \Sigma^2 \tau, \quad (\text{C.2})$$

such that $\Sigma = \sigma_{imp}(K, T)$ and $\tau = T - t_0$. The resultant option price is expressed as follows:

$$C(S_0, t_0, K, T, \Sigma) = S_0 e^{-\int_{t_0}^T q_s ds} [N(d_1) - e^y N(d_2)], \quad (\text{C.3})$$

where

$$d_1 = -\frac{y}{\sqrt{w}} + \frac{\sqrt{w}}{2} \quad (\text{C.4})$$

$$d_2 = -\frac{y}{\sqrt{w}} - \frac{\sqrt{w}}{2}. \quad (\text{C.5})$$

The partial derivatives of the call option with respect to K and T in Dupire's equation are as follows:

$$\begin{aligned} \frac{\partial C}{\partial K} &= \frac{\partial C}{\partial y} \frac{\partial y}{\partial K} + \frac{\partial C}{\partial w} \frac{\partial w}{\partial K} \\ &= \frac{1}{K} \frac{\partial C}{\partial y} + \frac{\partial w}{\partial K} \frac{\partial C}{\partial w}, \end{aligned} \quad (\text{C.6})$$

$$\begin{aligned} \frac{\partial^2 C}{\partial K^2} &= -\frac{1}{K^2} \frac{\partial C}{\partial y} + \frac{1}{K} \frac{\partial}{\partial K} \left(\frac{\partial C}{\partial y} \right) + \frac{\partial^2 w}{\partial K^2} \frac{\partial C}{\partial w} + \frac{\partial w}{\partial K} \frac{\partial}{\partial K} \left(\frac{\partial C}{\partial w} \right) \\ &= \frac{1}{K^2} \left(\frac{\partial^2 C}{\partial y^2} - \frac{\partial C}{\partial y} \right) + \frac{2}{K} \frac{\partial w}{\partial K} \frac{\partial^2 C}{\partial w \partial y} + \frac{\partial^2 w}{\partial K^2} \frac{\partial C}{\partial w} + \left(\frac{\partial w}{\partial K} \right)^2 \frac{\partial^2 C}{\partial w^2}, \end{aligned} \quad (\text{C.7})$$

$$\frac{\partial C}{\partial T} = -q_T C + (q_T - r_T) \frac{\partial C}{\partial y} + \frac{\partial w}{\partial T} \frac{\partial C}{\partial w}. \quad (\text{C.8})$$

Substituting equation (C.6), (C.7) and (C.8) into Dupire's local volatility equation (4.3) results in the following equation:

$$\sigma_L^2 = 2 \frac{-q_T C + (q_T - r_T) \frac{\partial w}{\partial y} + \frac{\partial w}{\partial T} \frac{\partial C}{\partial w} + (r_T - q_T) \frac{\partial C}{\partial y} + (r_T - q_T) K \frac{\partial w}{\partial K} \frac{\partial C}{\partial w} + q_T C}{\left(\frac{\partial^2 C}{\partial y^2} - \frac{\partial C}{\partial y} \right) + 2K \frac{\partial w}{\partial K} \frac{\partial^2 C}{\partial w \partial y} + K^2 \frac{\partial^2 w}{\partial K^2} \frac{\partial C}{\partial w} + K^2 \frac{\partial^2 C}{\partial w^2}}. \quad (\text{C.9})$$

The latter can be simplified by substituting the following identities into equation (C.9):

$$\frac{\partial^2 C}{\partial w^2} = \left(-\frac{1}{8} - \frac{1}{2w} + \frac{y^2}{2w^2} \right) \frac{\partial C}{\partial w} \quad (\text{C.10})$$

$$\frac{\partial^2 C}{\partial w \partial y} = \left(\frac{1}{2} - \frac{y}{w} \right) \frac{\partial C}{\partial w} \quad (\text{C.11})$$

$$\frac{\partial^2 C}{\partial y^2} - \frac{\partial C}{\partial y} = 2 \frac{\partial C}{\partial w}, \quad (\text{C.12})$$

resulting in a local volatility formula with respect to w and y :

$$\sigma_L^2 = \frac{\frac{\partial w}{\partial T} + (r_T - q_T) K \frac{\partial w}{\partial K}}{1 + K \frac{\partial w}{\partial K} \left(\frac{1}{2} - \frac{y}{w} \right) + \frac{1}{2} K^2 \frac{\partial^2 w}{\partial K^2} - \frac{1}{4} K^2 \left(\frac{\partial w}{\partial K} \right)^2 \left(\frac{1}{4} + \frac{1}{w} - \frac{y^2}{w^2} \right)}. \quad (\text{C.13})$$

To remove the variable w from equation (C.9) the partial derivatives of w with respect to K and T are given by:

$$\frac{\partial w}{\partial K} = 2\Sigma\tau \frac{\partial \Sigma}{\partial K} \quad (\text{C.14})$$

$$\frac{\partial^2 w}{\partial K^2} = 2\tau \left(\frac{\partial \Sigma}{\partial K} \right)^2 + 2\Sigma\tau \frac{\partial^2 \Sigma}{\partial K^2} \quad (\text{C.15})$$

$$\frac{\partial w}{\partial T} = \Sigma^2 + 2\Sigma\tau \frac{\partial \Sigma}{\partial T}. \quad (\text{C.16})$$

Substituting the above into equation (C.13) results in the local volatility equation which depends on implied volatilities:

$$\sigma_L^2 = \frac{\Sigma^2 + 2\Sigma\tau \left(\frac{\partial \Sigma}{\partial T} + (r_T - q_T) K \frac{\partial \Sigma}{\partial K} \right)}{1 + K\Sigma\tau \frac{\partial \Sigma}{\partial K} - 2\frac{Ky}{\Sigma} \frac{\partial \Sigma}{\partial K} + K^2 \Sigma \tau \frac{\partial^2 \Sigma}{\partial K^2} - \frac{1}{4} K^2 \Sigma^2 \tau^2 \left(\frac{\partial \Sigma}{\partial K} \right)^2 + \frac{K^2 y^2}{\Sigma^2} \left(\frac{\partial \Sigma}{\partial K} \right)^2}. \quad (\text{C.17})$$

C.2 Finite Difference Techniques

The following central difference approximations can be used to calculate the first and second derivative of a function $F(x)$ (Eberly, 2008):

$$F'(x) = \frac{F(x+h) - F(x-h)}{2h} + O(h^2) \quad (\text{C.18})$$

$$F''(x) = \frac{F(x+h) - 2F(x) + F(x-h)}{h^2} + O(h^2), \quad (\text{C.19})$$

where $O(h^2)$ is the truncation error. The partial derivatives in equation (4.3) can now be approximated as follows:

$$\frac{\partial C}{\partial T} = \frac{C(S_t, T + \delta T, K) - C(S_t, T - \delta T, K)}{2\delta T} \quad (\text{C.20})$$

$$\frac{\partial C}{\partial K} = \frac{C(S_t, T, K + \delta K) - C(S_t, T, K - \delta K)}{2\delta K} \quad (\text{C.21})$$

$$\frac{\partial^2 C}{\partial K^2} = \frac{C(S_t, T, K + \delta K) - 2C(S_t, T, K) + C(S_t, T, K - \delta K)}{\delta K^2}. \quad (\text{C.22})$$