Mixed Monte Carlo in the Foreign Exchange Market

Christopher Baker

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.

March 7, 2017

MPhil in Mathematical Finance,
University of Cape Town.

Declaration

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

Signed by candidate

Signature Removed

March 7, 2017

Abstract

The stochastic differential equation (SDE) describing the spot FX rate is of central importance to modelling FX derivatives. A Monte Carlo estimate of the discounted individual payoffs of FX derivatives is taken to arrive at the price, provided there does not exist a closed form solution for the price. One propagates the FX spot rate through time under risk-neutral dynamics to realise the before-mentioned payoffs. A drawback to Monte Carlo becomes evident when the model dynamics become more complicated, such as when more dimensions are added to the dynamics of the model. These additional dimensions can be stochastic volatility and/or stochastic domestic and foreign short rates. This dissertation describes the calibration of such a model using mixed Monte Carlo, as described in Cozma and Reisinger (2015), to both model-generated and market data. Profit and loss analysis of hedging FX derivatives using the mixed Monte Carlo method is conducted when hedging against both model-generated and market data .

Acknowledgements

I would like to thank Assoc. Prof. Thomas McWalter, Searle Silverman and Sheldon Maze for their supervision and guidance during this dissertation.

Additional thanks must be extended to my parents, Frank and Trish Baker, for their constant support as well as to Brontë Westcott for her fresh perspective on Monte Carlo matters.

Contents

1. Introduction . . . 1

2. Summary of Mixed Monte Carlo Method . . . 3
2.1 Application of Mixed Monte Carlo to FX Options . . . 3
2.2 Mixed Monte Carlo and Swaption Pricing . . . 6

3. Calibration . . . 8
3.1 Two-Part Calibration . . . 8
3.1.1 Outline of the Calibration Process . . . 8
3.1.2 Pricing the Cap Surface . . . 11
3.2 Calibration to Model-Generated Data . . . 15
3.2.1 Part 1: Calibration to the Cap Surface . . . 15
3.2.2 Part 2: Calibration to the Option Surface . . . 17
3.3 Calibration to Market Data . . . 20
3.3.1 Market Data . . . 20
3.3.2 Results of Calibration to Market Data . . . 22

4. Hedging . . . 26
4.1 Hedging Against Model-Generated Data . . . 26
4.1.1 Hedging with Full Knowledge of Parameters . . . 29
4.1.2 Hedging with Knowledge of Correlation Matrix . . . 30
4.1.3 Hedging with No Knowledge of Parameters . . . 32
4.2 Hedging Against Market Data . . . 33
4.2.1 Results of Calibrations . . . 33
4.2.2 Results of Hedging . . . 39

5. Conclusion . . . 43

Bibliography . . . 45

A. Cozma Reisinger Parameters . . . 46

B. Market Data . . . 47
B.1 Discount Curves . . . 47
B.2 USD Cap Surface . . . 48
B.3 EUR Cap Surface . . . 50
B.4 EUR/USD Option Surface . . . 52

List of Figures

3.1 Market and calibrated model surfaces for caps 16
3.2 Difference between the market and calibrated model cap surfaces . . 17
3.3 The market and calibrated model option surfaces 18
3.4 Difference between input and calibrated model option surface 20
3.5 USD market cap smile vs. calibrated model smile 24
3.6 EUR market cap smile vs. calibrated model smile 24
3.7 Market call option smile vs. calibrated model smile 25

4.1 PnL from daily delta-hedging with known parameters 29
4.2 Time series of hedge portfolio vs. option value 30
4.3 Time series of difference between hedge portfolio and option value . 31
4.4 Monte Carlo calibrated parameter values when the correlation matrix is known 32
4.5 Time series of Monte Carlo calibrated parameters 33
4.6 Time series of SSD for each instrument calibrated 34
4.7 Box plot of SSD for option smile calibration 35
4.8 Time series of calibrated USD short rate parameter values 36
4.9 Time series of calibrated EUR short rate parameter values 37
4.10 Time series of calibrated volatility parameter values 37
4.11 Time series of correlation matrix values 38
4.12 Time series of SSD for each instrument calibration 38
4.13 Hedge portfolios of options sold on 3 January 2011 40
4.14 Hedge portfolios of options sold on 3 January 2011 without outliers . 41
4.15 Difference between hedge portfolios and value of options sold on 3 January 2011 41
4.16 Evolution of the EURUSD spot FX rate from 3 January 2011 42
4.17 Evolution of the delta of options sold on 3 January 2011 42

List of Tables

3.1 Results of calibration to the cap surface 17
3.2 Results of calibration to the input option surface 20
3.3 Parameter values resulting from market calibration 25

A.1 Parameters used in Cozma and Reisinger (2015) 46

B.1 OIS Discount Curve tickers 47
B.2 Black Volatilities used to create USD Cap Surface 48
B.3 Black Volatilities used to create USD Cap Surface cont. 49
B.4 Black Volatilities used to create EUR Cap Surface 50
B.5 Black Volatilities used to create EUR Cap Surface cont. 51
B.6 Black Volatilities used to create EUR/USD Option Surface 52

Chapter 1

Introduction

In the foreign exchange (FX) market, the stochastic differential equation (SDE) describing the spot FX rate (the number of units of domestic currency required to buy one unit of foreign currency) is of central importance to modelling FX derivatives. In order to price derivatives, a Monte Carlo estimate of the discounted individual derivative payoffs is taken to arrive at the price, provided there is no closed form solution such as in the case of the Black Scholes model. This involves discretising time and propagating the FX spot rate using the Euler-Maruyama or Milstein methods under risk-neutral dynamics.

Monte Carlo is popular because of its flexibility and since it scales linearly with dimension. The convergence rate of Monte Carlo methods is $\mathcal{O}(N^{-\frac{1}{2}})$, where N is the sample size, i.e., the number of simulated paths of the process described by the SDE. The drawback to Monte Carlo is evident when the model dynamics become more complicated, such as when more dimensions are added to the dynamics of the model factors. Allowing the volatility of the FX rate to vary stochastically is an example of adding a dimension to the FX rate dynamics (Heston, 1993). Monte Carlo methods can become computationally expensive in high dimensions because multiple paths of each factor must be simulated. Monte Carlo variance-reduction techniques, as presented and discussed in Glasserman (2003), were developed to circumvent this drawback. Variance-reduction techniques result in Monte Carlo estimates converging to a solution at smaller values of N.

Spot FX rate dynamics incorporate the domestic short rate r^d, foreign short rate r^f and volatility of the spot FX rate. While it is convenient to assume that interest rates are constant, it is shown in Van Haastrecht *et al.* (2009) that volatility in interest rates influences FX option pricing to a greater extent than FX rate volatility. Interest rate dynamics proposed in Cox *et al.* (1985) prove popular in practice because the dynamics preserve the non-negativity and mean reversion properties of

interest rates, both of which are desirable. Heston (1993) dynamics for stochastic volatility of asset prices and FX rates are also widely used.

Combining the above-mentioned factors results in a four-factor Heston-CIR model for the spot FX rate that allows for stochastic volatility, stochastic domestic and foreign short rate dynamics. This model can be used in conjunction with Monte Carlo methods to price FX derivatives. Monte Carlo methods involve simulating multiple paths of the four stochastic processes, therefore there is interest in developing and testing variance reduction techniques for four-factor spot FX rate models.

The variance reduction technique tested in this paper is an extension of the mixed Monte Carlo/PDE method that is developed in Loeper and Pironneau (2009). Mixed Monte Carlo was initially developed for two-factor models but is extended to four-factor models in Cozma and Reisinger (2015) and Ahlip and Rutkowski (2013). The mixed Monte Carlo method invokes the Law of Total Expectation (Billingsley, 1995) to condition upon complete knowledge of the paths of the volatility of the FX rate (v_t), domestic short rate (r_t^d) and foreign short rate (r_t^f) processes, after which the expectation (a Monte Carlo estimate) of the resulting, more tractable, expression can be evaluated. Mixed Monte Carlo reduces the variance of the overall estimate by reducing the number of factors for which paths are simulated, i.e., only the paths that are conditioned upon are simulated (v_t, r_t^f and r_t^d). This allows Monte Carlo estimates with lower variance for a particular N, which means that pricing or hedging simulations consume less time for the same level of accuracy.

This dissertation describes the calibration of this model and then applies the mixed Monte Carlo method developed in Cozma and Reisinger (2015) to market data and includes a profit and loss analysis of hedging FX derivatives.

The remainder of this dissertation is organized as follows: Chapter 2 provides a summary of the mixed Monte Carlo method and its application to European-style FX options; Chapter 3 describes the calibration method for this model and evaluates the calibration to both model and market data; Chapter 4 presents a delta-hedging exercise of European-style FX options to both model and market data; and Chapter 5 concludes.

Chapter 2

Summary of Mixed Monte Carlo Method

In this chapter the method of Cozma and Reisinger (2015) is summarised and the application to European-style FX options is reviewed. There is also a short section on the application of mixed Monte Carlo to swaption pricing.

2.1 Application of Mixed Monte Carlo to FX Options

The four-factor model for the FX spot rate under investigation (as described in Chapter 1) can be described using the SDEs that follow in (2.1) to (2.4) . Under the risk-neutral, $\mathbb{Q}$, dynamics

$$dS_t = S_t(r_t^d - r_t^f)\,dt + \sqrt{v_t}S_t\,dW_t^1, \tag{2.1}$$

$$dv_t = \kappa(\theta - v_t)\,dt + \xi\sqrt{v_t}\,dW_t^2, \tag{2.2}$$

$$dr_t^d = \kappa_d(\theta_d - r_t^d)\,dt + \xi_d\sqrt{r_t^d}\,dW_t^3, \tag{2.3}$$

$$dr_t^f = \left(\kappa_f\theta_f - \kappa_f r_t^f - \rho_{sf}\xi_f\sqrt{v_t r_t^f}\right)dt + \xi_f\sqrt{r_t^f}\,dW_t^4, \tag{2.4}$$

where S is the spot FX rate, v is the volatility of S, r^f is the foreign short rate and r^d is the domestic short rate. Further, κ, κ_d and κ_f are the speed of mean-reversion for the volatility, domestic short rate and foreign short rate processes, respectively. In addition, θ, θ_d and θ_f are the level of mean-reversion for the volatility, for the domestic short rate and foreign short rate processes, respectively. Finally, ξ, ξ_d and ξ_f are the (constant) volatilities for the volatility, domestic short rate and foreign short rate processes, respectively.

The $\mathbb{Q}$ Brownian motions $W_t^i (i = 1, 2, 3, 4)$ are correlated with correlation matrix $\mathbf{\Sigma}$. Define $\mathbf{W} := [W_t^1, W_t^2, W_t^3, W_t^4]^{\mathrm{T}}$ and $\tilde{\mathbf{W}} := [\tilde{W}_t^1, \tilde{W}_t^2, \tilde{W}_t^3, \tilde{W}_t^4]^{\mathrm{T}}$, where $\tilde{W}_t^i (i = 1, 2, 3, 4)$ are uncorrelated Brownian motions. Then $\mathbf{W} = \mathbf{A}\tilde{\mathbf{W}}$ if $\mathbf{A}$ is

the upper triangular Cholesky decomposition of $\mathbf{\Sigma}$. The matrix $\mathbf{A}$ has elements a_{ij} for $i, j = 1, 2, 3, 4$ and $a_{ij} = 0$ for $i > j$.

The dynamics for the FX spot rate from (2.1) can now be written as

$$dS_t = S_t(\mu_t\, dt + a_{11}\sqrt{v_t}\, d\tilde{W}_t), \tag{2.5}$$

which implies that

$$S_t = S_0 \exp\left\{\int_0^t \mu_s - \frac{1}{2}a_{11}^2 v_s\, ds + \int_0^t a_{11}\sqrt{v_s}\, d\tilde{W}_s^1\right\}. \tag{2.6}$$

However, from (2.1), and because $W_t^1 = \sum_{j=1}^4 a_{1j}\tilde{W}_t^j$, it is the case that

$$S_t = S_0 \exp\left\{\int_0^t r_s^d - r_s^f - \frac{1}{2}v_s ds + \sum_{j=1}^4 a_{1j}\int_0^t \sqrt{v_s}\, d\tilde{W}_t^j\right\}. \tag{2.7}$$

It follows that

$$\int_0^t \mu_s ds = \int_0^t r_s^d - r_s^f - \frac{v_s}{2}(a_{11}^2 - 1)ds + \sum_{j=2}^4 a_{1j}\int_0^t \sqrt{v_s}\, d\tilde{W}_t^j. \tag{2.8}$$

When this scheme is discretised with constant time increments, for $t \in [t_i, t_{i+1}]$ and $\Delta_i := (t_{i+1} - t_i) = \Delta t$, then

$$\mu_{t_i} = r_{t_i}^d - r_{t_i}^f + \frac{v_{t_i}}{2}(a_{11}^2 - 1) + \sum_{j=2}^4 a_{1j}\sqrt{v_{t_i}}\frac{\tilde{W}_{t_{i+1}}^j - \tilde{W}_{t_i}^j}{\Delta t}. \tag{2.9}$$

If one conditions on $\mathcal{F}_t^{v,r^d,r^f}$, the natural filtration generated by v_t, r_t^d and r_t^f (and hence W_t^2, W_t^3 and W_t^4), then the FX rate process S evolves like a geometric Brownian motion with time-dependent drift and volatility (Cozma and Reisinger, 2015). As a result, the *conditional FX spot rate* can be written in a Black Scholes-type form as

$$S_t = S_0 \exp\Big\{\int_0^t (r_s^d - r_s^f - \frac{1}{2}v_s)ds + \sum_{j=2}^4 \int_0^t a_{1j}\sqrt{v_s}\frac{\delta \tilde{W}_s^j}{\delta t}ds + a_{11}\int_0^t \sqrt{v_s}d\tilde{W}_s^1\Big\}, \tag{2.10}$$

where $\Delta\tilde{W}_t^j := (\tilde{W}_{t_{i+1}}^j - \tilde{W}_{t_i}^j), \forall t \in [t_i, t_{i+1})$. Thus, conditional on $\mathcal{F}_t^{v,r^d,r^f}$, the FX spot rate can be expressed as

$$S_t = S_0 \exp\left\{(\bar{r}^d - \bar{r}^f - \frac{1}{2}\bar{\sigma}^2)t + \sum_{j=2}^4 \sum_{i=1}^N a_{1j}\sqrt{v_{t_i}}\Delta\tilde{W}_i^j + a_{11}\bar{\sigma}\sqrt{t}Z\right\}, \tag{2.11}$$

where

$$\bar{\sigma}^2 := \frac{1}{t}\int_0^t v_s^d ds = \frac{1}{N}\sum_{i=1}^{N} v_{t_i}, \tag{2.12}$$

$$\bar{r}^d := \frac{1}{t}\int_0^t r_s^d ds = \frac{1}{N}\sum_{i=1}^{N} r_{t_i}^d, \tag{2.13}$$

$$\bar{r}^f := \frac{1}{t}\int_0^t r_s^f ds = \frac{1}{N}\sum_{i=1}^{N} r_{t_i}^f \tag{2.14}$$

and Z is a standard normal random variable. The expression for S_t can be further simplified to

$$S_t = S_0 \exp\left\{(r - q - \frac{1}{2}\sigma^2)t + \sigma\sqrt{t}Z\right\}, \tag{2.15}$$

where

$$r := \bar{r}^d \tag{2.16}$$

$$\sigma^2 := a_{11}^2\bar{\sigma}^2 \tag{2.17}$$

$$q := \bar{r}^f + \frac{1}{2}(1 - a_{11}^2)\bar{\sigma}^2 - \frac{1}{t}\sum_{j=2}^{4}\sum_{i=1}^{N} a_{1j}\sqrt{v_{t_i}}\delta\tilde{W}_i^j. \tag{2.18}$$

Now S_t is a log-normal random variable and has distribution

$$S_t \sim \mathrm{LN}(\ln(S_0) + (r - q - \frac{1}{2}\sigma^2)t, \sigma^2 t). \tag{2.19}$$

Hence, one is able to derive an analytical expression for the *conditional price* of an FX call option.

The time 0 price of a contingent claim X_0, with maturity T, which derives value from the FX spot rate S_T, is

$$X_0 = \mathbb{E}_{\mathbb{Q}}\left[\exp\left\{-\int_0^T r_s ds\right\} f(S_T)\right] = \mathbb{E}_{\mathbb{Q}}\left[\mathbb{E}_{\mathbb{Q}}\left[\exp\left\{-\int_0^T r_s ds\right\} f(S_T)\middle|\mathcal{F}_t^{v,r^d,r^f}\right]\right]. \tag{2.20}$$

The inner, conditional, expectation can be written as

$$S_0 \exp\{-qT\}\,\Phi(d_1) - K\exp\{-rT\}\,\Phi(d_2) \tag{2.21}$$

where

$$d_1 = \frac{\ln(\frac{S_0}{K}) + (r - q + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}} \tag{2.22}$$

$$d_2 = d_1 - \frac{1}{2}\sigma\sqrt{T} \tag{2.23}$$

and Φ is the cumulative distribution function of a standard normal random variable.

The price of the FX call option is estimated by means of Monte Carlo estimates of (2.21). In order to do that, realisations of the v, r_d and r_f processes are generated. Next, one calculates r, σ^2, and q using the realisations of v, r_d and r_f and the expressions in (2.16), (2.17) and (2.18). Once r, σ^2 and q have been calculated, one can create a *conditional option price* as shown in (2.21), (2.22) and (2.23). The price of the FX call is estimated by taking a Monte Carlo estimate (the mean) of a sufficiently-large sample of independent realisations of the conditional option price, (2.21).

2.2 Mixed Monte Carlo and Swaption Pricing

One of the tasks of this dissertation is to investigate pricing and hedging cross-currency swaptions using mixed Monte Carlo. The implementation of mixed Monte Carlo pivots on the complete knowledge of three of the four processes in order to arrive at a conditional price for the derivative. Next, one uses the Law of Total Probability to calculate a mixed Monte Carlo price for the derivative by taking a Monte Carlo estimate of the conditional option prices.

In Section 2.1, by conditioning on the full knowledge of the volatility, domestic short rate and foreign short rate processes, one arrives at a Black-Scholes-type conditional option price for a European-style FX options.

This is not the case for cross-currency swaptions. There exists an analytical solution for swaptions in a one-factor model in the form of Jamshidian's trick (Jamshidian, 1989). One could condition on the full knowledge of the foreign short rate, spot FX rate and volatility processes in an attempt to obtain a conditional option price of the swaption. The problem is that in order to realise the spot FX rate, one needs realisations of the domestic short rate process. The spot FX rate process is bound to the remaining three processes in such a way that future realisations of the spot FX rate depend on realisations of the remaining three processes. Thus, the only way to derive a conditional option price is to condition on knowledge of the foreign short

rate, spot FX rate, and volatility processes. One is not able to derive a conditional price for swaptions under these circumstances.

Chapter 3

Calibration

This chapter consists of three sections, in which the process of calibrating the Heston-CIR-CIR model is covered. The first section describes the two-part calibration of the model. The second section outlines the calibration of the model to model-generated data, which demonstrates the ability to recover input parameters and to match the relevant pricing surface. The third section describes the calibration of the model to EURUSD FX market data. Throughout this section, when using model-generated data, the parameter values used to generate market prices are the same as those used in Cozma and Reisinger (2015), which are displayed in Appendix A.

3.1 Two-Part Calibration

This section describes the two-part calibration procedure followed when calibrating the model. It explores the assumptions made during calibration and outlines the valuation of the instruments used in the calibration.

3.1.1 Outline of the Calibration Process

The calibration of the model follows a two-part process. Firstly, the foreign short rate (r^f) process parameters ($r_0^f, \kappa_f, \theta_f$ and ξ_f) are calibrated to the foreign cap surface. One could, equivalently, calibrate to the floor surface, but in this dissertation the model is calibrated to the cap surface. The calibration to the cap surface is effected by minimising the sum of squared differences between the market prices and model prices. The same procedure is followed for the domestic short rate process (r^d) parameters ($r_0^d, \kappa_d, \theta_d$ and ξ_d) and the domestic cap surface. Secondly, the volatility process (v) parameters (v_0, κ, θ and ξ) and the elements of the correlation matrix ($\rho_{sd}, \rho_{sf}, \rho_{df}, \rho_{sv}, \rho_{vf}$ and ρ_{vd}) are calibrated to the European-style FX call option surface via Monte Carlo calibration. One could, equivalently, calibrate to the European-style FX put surface, but in this dissertation the model is calibrated to the call surface. Throughout this dissertation, the European-style FX call option

surface will be referred to as the *option surface*, since no other option surfaces are considered.

Restrictions are placed on the upper and lower bounds of the calibrated parameters. The elements of the correlation matrix are restricted to fall between -1 and 1, i.e. $(\rho_{sd}, \rho_{sf}, \rho_{df}, \rho_{sv}, \rho_{vf}, \rho_{vd}) \in [-1; 1]$. The speeds of mean reversion, κ, κ_d and κ_f, are constrained to be positive. The mean reversion level of the volatility process, θ, is constrained to be positive. The level of mean reversion for the short rate processes, θ_d and θ_f, are also restricted to be positive.

In practice, the domestic short rate (r^d) is not calibrated to the domestic cap surface in the same way as the foreign short rate. This is because the domestic short rate process models the short rate that determines the basis zero curve. The basis zero curve differs from the swap zero curve, or zero coupon bond curve (ZCB), in that it incorporates country risk into the curve. Thus, in order to accurately model the dynamics, the domestic short rate has to be calibrated using financial instruments that include country risk, such as FX options. In this dissertation, this restriction is relaxed, so the domestic short rate process is calibrated to the domestic cap surface. This assumption simplifies the calibration process, allowing for a faster calibration as well as shifting the focus of this dissertation to the implementation of mixed Monte Carlo rather than an arbitrage-free calibration. The effect of the relaxation of the above-mentioned restriction is minimised by looking at FX crosses where the counter-parties have minimal country risk, such as EURUSD. An extension of this dissertation could focus on an arbitrage-free calibration of the model, without the above assumption. The EURUSD currency cross represents the price of EUR1 in terms of USD.

Effectively, calibration to the cap surface is an exercise in solving for four parameters at a time: $r_0^d, \kappa_d, \theta_d$ and ξ_d or $r_0^f, \kappa_f, \theta_f$ and ξ_f, depending on which process is being calibrated. The number of parameters for which one is solving creates the opportunity for different model parameters to generate the same pricing surface, thus hindering the calibration process. By reducing the number of parameters for which one is solving, it is possible to stabilise the calibration process. 'Stability' in this sense refers to the ability to recover input model parameters from model-generated data through calibration.

Carr *et al.* (2009) highlight that in short rate models, the very short end of the ZCB curve is governed by the initial short rate, i.e., r_0^d and r_0^f. Further, Carr *et al.* (2009)

identify that the mean reversion level of the process, θ_d and θ_f, controls the long-end of the ZCB curve.

Thus, if the shortest tenor of the domestic ZCB curve is used as a proxy for r_0^d and the longest tenor is used as a proxy for θ_d, one is able to change the calibration from a *four parameter calibration* to a *two parameter calibration*. The same result holds for the foreign short rate process.

Since the shortest tenors on the prevailing ZCB curves (both domestic and foreign) are being used as proxies for the initial short rate, one can estimate the correlation between the Brownian motions that drive the domestic and foreign short rate processes, ρ_{df}. The reasoning behind this claim is illustrated below.

Consider the following:

$$\begin{aligned} dr_t^d &= \kappa_d(\theta_d - r_t^d)\,dt + \xi_d\sqrt{r_t^d}\,dW_t^3 \\ \implies \frac{dr_t^d}{\sqrt{r_t^d}} &= \frac{\kappa_d(\theta_d - r_t^d)d}{\sqrt{r_t^d}}\,dt + \xi_d\,dW_t^3. \end{aligned}$$

Also,

$$\begin{aligned} dr_t^f &= \left(\kappa_f\theta_f - \kappa_f r_t^f - \rho_{sf}\xi_f\sqrt{v_t r_t^f}\right)\,dt + \xi_f\sqrt{r_t^f}\,dW_t^4, \\ \implies \frac{dr_t^f}{\sqrt{r_t^f}} &= \frac{\left(\kappa_f\theta_f - \kappa_f r_t^f - \rho_{sf}\xi_f\sqrt{v_t r_t^f}\right)}{\sqrt{r_t^f}}\,dt + \xi_f\,dW_t^4. \end{aligned}$$

When working in continuous time, this means that:

$$\begin{aligned} &\text{Corr}\left[\frac{dr_t^d}{\sqrt{r_t^d}}, \frac{dr_t^f}{\sqrt{r_t^f}}\right], \\ &= \frac{\text{Cov}\left[\frac{dr_t^d}{\sqrt{r_t^d}}, \frac{dr_t^f}{\sqrt{r_t^f}}\right]}{\text{SD}\left[\frac{dr_t^d}{\sqrt{r_t^d}}\right]\text{SD}\left[\frac{dr_t^f}{\sqrt{r_t^f}}\right]}, \\ &= \frac{\text{Cov}\left[\xi_d\,dW_t^3, \xi_f\,dW_t^4\right]}{\text{SD}\left[\xi_d\,dW_t^3\right]\text{SD}\left[\xi_f\,dW_t^4\right]}. \end{aligned}$$

In this dissertation, time is discretised into time steps of equal size, i.e., time step $i = \Delta_i = \Delta \; \forall i$, so

$$\begin{aligned}
&\frac{\text{Cov}\left[\xi_d \, dW_t^3, \xi_f \, dW_t^4\right]}{\text{SD}\left[\xi_d \, dW_t^3\right] \text{SD}\left[\xi_f \, dW_t^4\right]} \\
&\approx \frac{\xi_d \xi_f \sqrt{\Delta}\sqrt{\Delta}\text{Cov}\left[Z^3, Z^4\right]}{\xi_d \xi_f \sqrt{\Delta}\sqrt{\Delta}\text{SD}\left[Z^3\right] \text{SD}\left[Z^4\right]}, \\
&= \frac{\xi_d \xi_f \, \Delta \text{Cov}\left[Z^3, Z^4\right]}{\xi_d \xi_f \, \Delta}, \\
&= \rho_{df},
\end{aligned}$$

where $Z^3, Z^4 \sim N(0,1)$ with correlation ρ_{df} .

As a result of the above, and because the shortest tenors of the EUR and USD swap curve are used as proxies for r_0^f and r_0^d (Carr *et al.*, 2009), the correlation of the Brownian motions driving the domestic and foreign short rates (ρ_{df}) is estimated using historical swap curve data. Estimating the other correlations by the same method as above is not possible because the time t value of the volatility process v is not observable, i.e., one can not identify a value for v_t.

Once the domestic and foreign short rate processes (r^d and r^f) have been calibrated to the respective cap surfaces, it remains to calibrate the volatility (v) process parameters (κ, θ and ξ) and remaining correlation coefficients ($\rho_{sf}, \rho_{sv}, \rho_{sd}, \rho_{vd}$ and ρ_{vf}). The volatility parameters and five elements of the correlation matrix are calibrated via Monte Carlo calibration to the option surface.

When performing the Monte Carlo calibration, the model price changes because of changing parameters as well as changes in the Monte Carlo sample used to create the surface. In order to isolate impact of the change in model parameters, the same set of input standard normal random variables is used to price the calibrated surface. The standard normal random variables are replaced by Sobol sequences to reduce the variance in the calibrated pricing surface further (McWalter, 2016). The transformed Sobol numbers are used in all model pricing and calibration exercises.

3.1.2 Pricing the Cap Surface

This subsection outlines the process of creating a cap surface under the Cox-Ingersoll-Ross (CIR) model for short rates. The presentation of formulae as well as the notation is the same as in Brigo and Mercurio (2001). The presentation and pricing of interest rate derivatives is as per Ouwehand (2015).

The risk-neutral ($\mathbb{Q}$) dynamics for the short rate in the CIR model are

$$dr_t = \kappa(\theta - r_t)\,dt + \xi\sqrt{r_t}\,dW_t, \tag{3.1}$$

where κ is the speed of mean reversion, θ is the mean-reversion level and ξ is the volatility of the short rate process.

In the CIR model, the time t price of a zero coupon bond that matures at time T, where t, T are year fractions from t_0 that obey the day count convention, is

$$P(t,T) = A(t,T) \cdot \exp\{-B(t,T)r_t\},$$

where

$$\begin{aligned} A(t,T) &= \Big[\frac{2h\exp\{(\kappa+h)(T-t)/2\}}{2h + (\kappa+h)(\exp\{(T-t)h\}-1)}\Big]^{2\kappa\theta/\xi^2}, \\ B(t,T) &= \frac{2(\exp\{(T-t)h\}-1)}{2h + (\kappa+h)(\exp\{(T-t)h\}-1)}, \\ h &= \sqrt{\kappa^2 + 2\xi^2}. \end{aligned}$$

A caplet can be thought of as a call on the prevailing interest rate over a future time period and a cap is a portfolio of consecutive caplets struck at the same strike.

A cap is defined by its tenor structure $\mathcal{T} = \{t_0, t_1, ..., t_n\}$ such that $t_0 < t_1 < ... < t_n$, a reference rate (such as LIBOR or EURIBOR, which both follow the ACT/360 day count convention) and the strike rate R. The reference rate that applies to the period $[t_i, t_j]$ is denoted by $L(t_i, t_j)$. The reference rate pays interest of $L(t_i, t_j)(t_j - t_i)$ at t_j for the period (t_i, t_j), where t_i, t_j are year fractions (from t_0) that agrees with the rate-specific day-count convention and $i < j$. A cap is made up of i caplets, where the i^{th} caplet, c_i, pays

$$c_i(t_i) = (L(t_{i-1}, t_i) - R)^+(t_i - t_{i-1}) \tag{3.2}$$

at time t_i per unit nominal. The notation $(x-y)^+$ means the larger of $(x-y)$ and 0, i.e., $\max(x-y, 0)$.

A payment of $(L(t_{i-1}, t_i) - R)^+(t_i - t_{i-1})$ at time t_i can be discounted to a value of

$$\frac{(L(t_{i-1}, t_i) - R)^+(t_i - t_{i-1})}{1 + L(t_i, t_j)(t_j - t_i)} \tag{3.3}$$

at time t_{i-1}.

The working that follows in (3.4) to (3.7) shows how a call option on a LIBOR rate (a caplet) can be expressed as a put option on a ZCB. The expression in (3.3) can be further manipulated to show that, at time t_{i-1}, a payment of

$$\frac{(L(t_{i-1},t_i)-R)^+(t_i-t_{i-1})}{1+L(t_i,t_j)(t_j-t_i)} \tag{3.4}$$

$$= \frac{(1+L(t_{i-1},t_i)(t_i-t_{i-1})-(1+R(t_i-t_{i-1})))^+}{1+L(t_i,t_j)(t_j-t_i)} \tag{3.5}$$

$$= (1+R(t_i-t_{i-1}))\left(\frac{1}{1+R(t_i-t_{i-1})}-\frac{1}{1+L(t_i,t_j)(t_j-t_i)}\right)^+ \tag{3.6}$$

$$= (1+R(t_i-t_{i-1}))\left(\frac{1}{1+R(t_i-t_{i-1})}-P(t_{i-1},t_i)\right)^+, \tag{3.7}$$

which is the payoff of $(1+R(t_i-t_{i-1}))$ put options, expiring at t_{i-1}, on a ZCB that matures at t_i with a strike price of $\frac{1}{1+R(t_i-t_{i-1})}$ per unit nominal.

Cox *et al.* (1985) show that the time t price of a European-style call option, with maturity $T>t$, on a ZCB, with maturity $S>T$, strike price X and the prevailing short rate r_t is

$$\begin{aligned}\text{ZBC}(t,T,S,X) = {} & P(t,S)\chi^2\left(2\bar{r}[\rho+\psi+B(T,S)];\frac{4\kappa\theta}{\xi^2},\frac{2\rho^2 r_t\exp\{h(T-t)\}}{\rho+\psi+B(T,S)}\right)\\ & - XP(t,T)\chi^2\left(2\bar{r}[\rho+\psi];\frac{4\kappa\theta}{\xi^2},\frac{2\rho^2 r_t\exp\{h(T-t)\}}{\rho+\psi}\right),\end{aligned} \tag{3.8}$$

where

$$\begin{aligned}\rho &= \frac{2h}{\xi^2(\exp\{h(T-t)\}-1)}\\ \psi &= \frac{\kappa+h}{\xi^2}\\ \bar{r} &= \frac{\ln(A(T,S)/X)}{B(T,S)}\end{aligned}$$

and $\chi^2(x,v,\lambda)$ is the non-central chi-squared distribution function evaluated at x with v degrees of freedom and non-centrality parameter λ (Brigo and Mercurio, 2001).

A European put at time t, with maturity T and strike price X, written on a ZCB with maturity S, denoted by $\text{ZBP}(t,T,S,X)$, can be priced using put call parity as

$$\begin{aligned}&\text{ZBC}(t,T,S,X)+XP(t,T) = \text{ZBP}(t,T,S,X)+P(t,S)\\ &\implies \text{ZBP}(t,T,S,X) = \text{ZBC}(t,T,S,X)-P(t,S)+XP(t,T).\end{aligned} \tag{3.9}$$

Combining (3.8) and (3.9), one can price a cap at time t, with tenor structure $\mathcal{T}$, reference rates $L(t_i, t_{i+1})$ and strike R. The value of a cap with the before-mentioned specification is represented by $\text{Cap}(t, \mathcal{T}, R)$. The expression for the value of cap is

$$\text{Cap}(t, \mathcal{T}, R) = \sum_{i=1}^{n} (1 + X(t_i - t_{i-1})\text{ZBP}\left(t, t_{i-1}, t_i, \frac{1}{1 + X(t_i - t_{i-1})}\right). \tag{3.10}$$

Fair caps are struck at the prevailing fair swap rate. A swap is a contract in which two interest rates are exchanged at tenor dates that form a tenor structure $\mathcal{T}$. A fixed-for floating swap exchanges a fixed interest payment for a floating interest payment at the tenor dates. Thus, for tenor structure $\mathcal{T} = \{t_0, t_1, ..., t_n\}$, a payer swap with swap rate K has a payoff of $(L(t_{i-1}, t_i) - K)(t_{i-1} - t_i)$ per unit nominal at time t_i.

A problem that arises is that the above cashflow is only known at time t_{i-1}. At this point it is important to define the fair forward rate observed at time t for the period (t_i, t_{i+1}), where $t_i, t_{i+1} > t$. This fair forward rate is denoted by $f(t, t_{i-1}, t_i)$. The fair forward rate is defined in such a way that

$$(1 + L(t, t_{i-1})(t_{i-1} - t))\,(1 + f(t, t_{i-1}, t_i)(t_i - t_{i-1})) = (1 + L(t, t_i)(t_i - t))\,.$$

So,

$$\begin{aligned} f(t, t_{i-1}, t_i) &:= \left(\frac{1 + L(t, t_i)(t_i - t)}{1 + L(t, t_{i-1})(t_{i-1} - t)} - 1\right) \cdot \frac{1}{t_i - t_{i-1}} \\ &= \left(\frac{P(t, t_{i-1})}{P(t, t_i)} - 1\right) \cdot \frac{1}{t_i - t_{i-1}} \end{aligned}$$

The fair forward rate $f(t, t_{i-1}, t_i)$ is known at time t and can be used to value the swap. Trivially, $f(t, t, t_i) = L(t, t_i)$. Thus, if one represents the time $t = t_0$ value of a unit nominal, payer swap with swap rate K, tenor structure $\mathcal{T}$ reference rate $L(t_{i-1}, t_i)$ nominal by $\text{Swap}(t, \mathcal{T}, K)$. Then the time $t = t_0$ value of the swap is

$$\begin{aligned} \text{Swap}(t, \mathcal{T}, K) &= \sum_{i=1}^{n} (L(t_{i-1}, t_i) - K)(t_i - t_{i-1}) \frac{1}{1 + L(t, t_i)(t_i - t)} \\ &= \sum_{i=1}^{n} (f(t, t_{i-1}, t_i) - K)(t_i - t_{i-1}) P(t, t_i) \\ &= \sum_{i=1}^{n} \left(\left(\frac{P(t, t_{i-1})}{P(t, t_i)} - 1\right) \cdot \frac{1}{t_i - t_{i-1}} - K\right)(t_i - t_{i-1}) P(t, t_i) \\ &= \sum_{i=1}^{n} (P(t, t_{i-1}) - P(t, t_i) - K(t_i - t_{i-1}) P(t, t_i)) \\ &= 1 - P(t, t_n) - K \sum_{i=1}^{n} (t_i - t_{i-1}) P(t, t_i). \end{aligned}$$

A fair swap has zero value at inception, so it is required that K is defined as

$$K := \frac{1 - P(t, t_n)}{\sum_{i=1}^{n}(t_i - t_{i-1})P(t, t_i)}. \tag{3.11}$$

By combining (3.11), (3.10) and (3.9), one can price the cap surface using the observable ZCB curve.

3.2 Calibration to Model-Generated Data

This section details the process of calibrating the model to model-generated data. The first subsection outlines the procedure and the results of calibrating the short rate processes to cap price surfaces. The second subsection illustrates the procedure and results of calibrating the remaining parameters to the European-style FX call option surface, hereafter referred to as the *option surface*, via Monte Carlo calibration. In both instances, the sum of squared differences (SSD) between the input model price surface and calibrated model price surfaces are minimised using the MATLAB function `fmincon`. If one denotes the input model price of an instrument at time t by X_t and the calibrated model price by X_t^*, then the SSD is defined as

$$\text{SSD} := \left(X_t - X_t^*\right)^2.$$

3.2.1 Part 1: Calibration to the Cap Surface

An input cap price surface, hereafter referred to as the *cap surface*, is created using the closed form solution (3.10), input market parameters $r_0 = 0.0291, \kappa = 0.32, \theta = 0.0248$ and $\xi = 0.0317$, strikes equal to the fair swap rate plus a premium/discount and multiple tenors. One could also minimise the difference between the Black volatility implied by the input cap surface and the calibrated model cap surface, but in this dissertation the SSD is calculated in reference to the cap price surface. The premium/discount ranges in steps from -0.5% to 0.5% in steps of 0.1%. The tenors range from one to ten years in steps of a year.

Initial and Mean Reversion Levels Unknown

If one assumes no knowledge of the initial and mean reversion levels of the short rate process, the calibration involves searching for four parameters, $r_0^f, \kappa_f, \theta_f$ and ξ_f, and minimising the SSD between the calibrated model cap surface and the input cap surface. The minimisation stops with the SSD at 1.56×10^{-5}. This particular minimisation (initial vector $x_0 = [0.01; 1; 0.01; 0.1]^T$) takes 12 minutes. The accuracy of the calibration (magnitude of SSD) depends on the initial guess vector. The

resulting calibrated parameters are shown in Table 3.1. All of the calibrated parameters are different from the input market parameters. This is potentially because the number of parameters being searched for in the calibration allows local minima in the pricing surface.

In Figure 3.1 one is able to see the slight difference in the market and calibrated model cap surfaces. The difference between the two surfaces is shown in Figure 3.2.

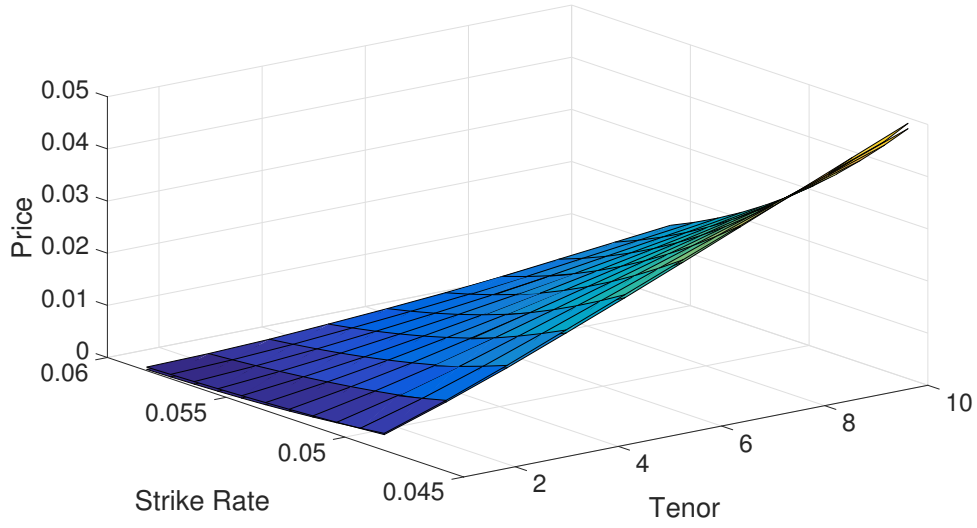

Fig. 3.1: Market and calibrated model surfaces for caps

Initial and Mean Reversion Levels Fixed

If one fixes the initial short rate and mean reversion level, r_0 and θ respectively, as mentioned earlier (Carr *et al.*, 2009), then the calibration involves minimising the SSD by changing only κ and ξ. The minimisation stops when the SSD reaches 7.68×10^{-9} . The minimisation (initial vector $x_0 = [1; 0.1]^T$) takes one minute. The resulting calibrated parameters are shown in Table 3.1. The calibrated parameters are the same as the input parameters to at least four decimal places.

It is clear that by implying r_0^f and θ_f from the ZCB curve, as suggested by Carr *et al.* (2009), one increases the stability of the first part of the calibration. Furthermore, the calibration is accelerated considerably.

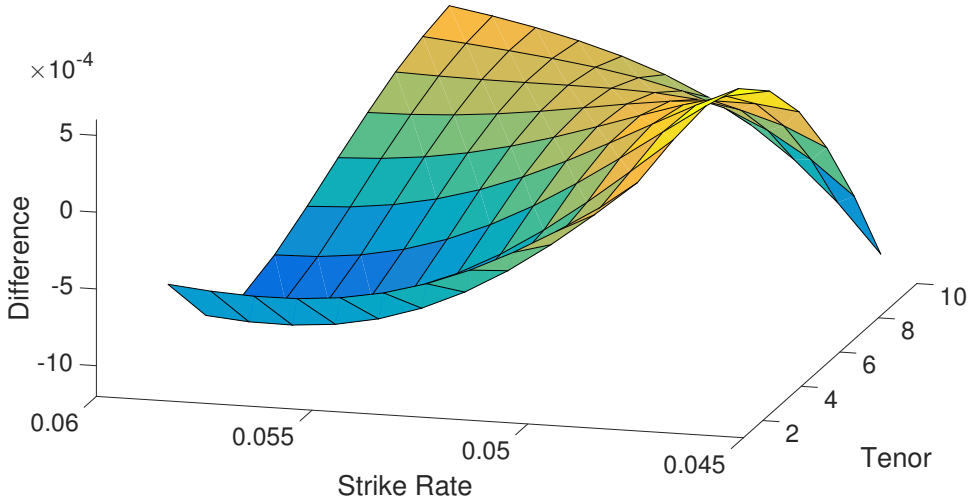

Fig. 3.2: Difference between the market and calibrated model cap surfaces

Parameter	Input Value	Full Calibration	Two Values Fixed
r_0^f	0.0291	0.0339	-
θ_f	0.0248	0.0507	-
κ_f	0.320	5.1600	0.3200
ξ_f	0.0317	0.2281	0.0317

Tab. 3.1: Results of calibration to the cap surface

3.2.2 Part 2: Calibration to the Option Surface

An input option surface is created using: mixed Monte Carlo pricing, the input parameters listed in Appendix A, multiple strikes and a variety of tenors. The strikes used range from 100 to 110 in steps of one and the tenors extend from one month to a year in steps of a month.

At this stage, the dynamics of r^d and r^f are assumed to be known after part one of the calibration. The remaining unknown parameters are the v parameters, i.e., v_0, κ, θ and ξ, as well as the elements of the correlation matrix $\rho_{sf}, \rho_{sv}, \rho_{sd}, \rho_{vd}$ and ρ_{vf}.

Correlation Matrix Assumed Known

Initially, one assumes knowledge of the correlation matrix. When pricing the call surface with 240 time steps per year and a quasi Monte Carlo sample of 10^4

Brownian-Bridged Sobol numbers for the calibration, one is able to recover the input volatility process parameters to an acceptable degree of error. The calibration involves minimising the SSD, again using `fmincon`, by changing v_0, κ, θ and ξ. The minimisation stops when the SSD reaches 1.18×10^{-8} . This minimisation (initial vector $x_0 = [0.01; 0.01; 0.1; 0.01]^T$) takes approximately six minutes. The Monte Carlo calibration takes longer than the cap surface calibration because each iteration of the pricing function requires the computation of a Monte Carlo estimate, whereas in Section 3.2.1 each repricing took place using an analytical formula. The resulting calibrated parameters are shown below in Table 3.2. The calibrated parameters are the same as the input market parameters to six decimal places.

The input price surface and calibrated model price surfaces are plotted in Figure 3.3. The model surface fits the market surface in such a way that one is unable to distinguish between the two, thus indicating an accurate model calibration. The difference between the input and calibrated model surfaces is plotted in Figure 3.4 (a). The difference arises because the minimisation stops when the change in SSD falls below a certain threshold and thus the parameters do not match to the last decimal. It is important to note the scale of the difference plot compared to the price surface. Noting this gives a further indication of the accuracy of the calibration.

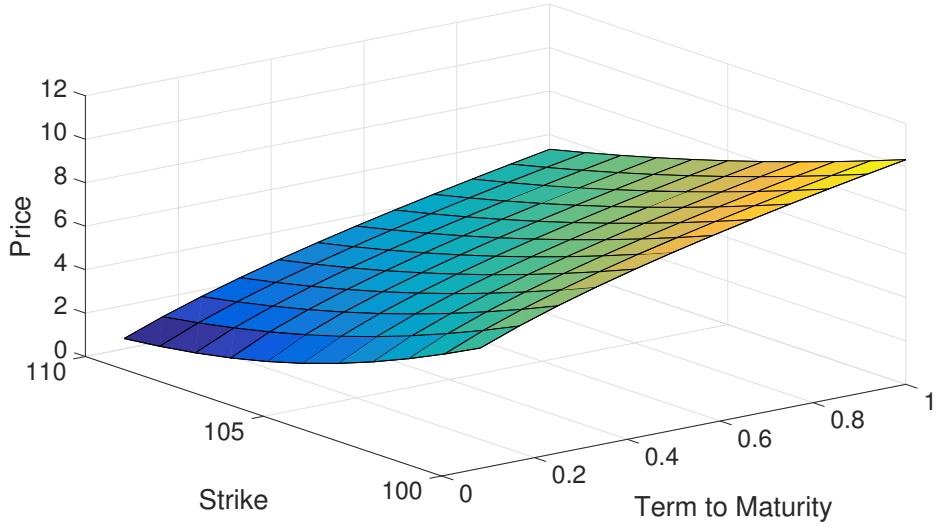

Fig. 3.3: The market and calibrated model option surfaces

Correlation Matrix Unknown

Now, the assumption of full knowledge of the correlation matrix is relaxed. The entire correlation matrix is unknown, except for ρ_{df}, which is assumed to have been estimated from historic data.

The Monte Carlo calibration searches for $v_0, \kappa, \theta, \xi, \rho_{sf}, \rho_{sv}, \rho_{sd}, \rho_{vd}$ and ρ_{vf}, nine parameters in total. This calibration is less stable than when one assumes knowledge of the correlation matrix. During this Monte Carlo calibration, one is calibrating for elements of a correlation matrix and a correlation matrix is required to be positive definite (by definition). A positive definite matrix has strictly positive eigenvalues, thus an extra constraint is placed on the minimisation so that the correlation matrix resulting from each iteration of the minimisation is forced to have strictly positive eigenvalues.

This minimisation is less successful than the previous calibration for four parameters, yet one is still able to recover the volatility parameters to an acceptable degree of accuracy. Despite using the same quasi-random numbers to generate both the input surface and the calibrated surface, one is not able to recover all of the elements of the correlation matrix. It is important to note that even though the parameters are not exactly the same, the resulting calibrated model is indistinguishable from the market surface. The ability to fit two surfaces with different parameters is evidence of over parameterisation of the model. Thus, the model may be able to calibrate to an acceptable degree of accuracy with different parameter sets.

The minimisation stops when the SSD reaches 3.06×10^{-9} . This minimisation (initial vector $x_0 = [0.01; 0.01; 0.1; 0.01; 0; 0; 0; 0; 0]^T$) takes 24 minutes. The resulting calibrated parameters are shown below in Table 3.2. An error plot is included in Figure 3.4(b). The error plot is a similar shape to F igure 3.4(a) but the difference at the longest maturity (1 year) is larger in magnitude. This is potentially due to the effect of incorrect correlation between the Brownian motions becoming more pronounced over time. The absolute difference is smaller (in the majority of the plot) than the previous calibration.

Parameter	Input Value	Correlations Known	Correlations Unknown
v_0	0.0275	0.0275	0.0275
θ	0.0232	0.0232	0.0232
κ	1.700	1.7021	1.6855
ξ	0.1500	0.1505	0.1492
ρ_{sf}	-0.1500	-	-0.0953
ρ_{sd}	-0.1500	-	-0.0125
ρ_{sv}	-0.100	-	-0.1007
ρ_{vf}	0.0500	-	0.0484
ρ_{vd}	0.1200	-	-0.0487

Tab. 3.2: Results of calibration to the input option surface

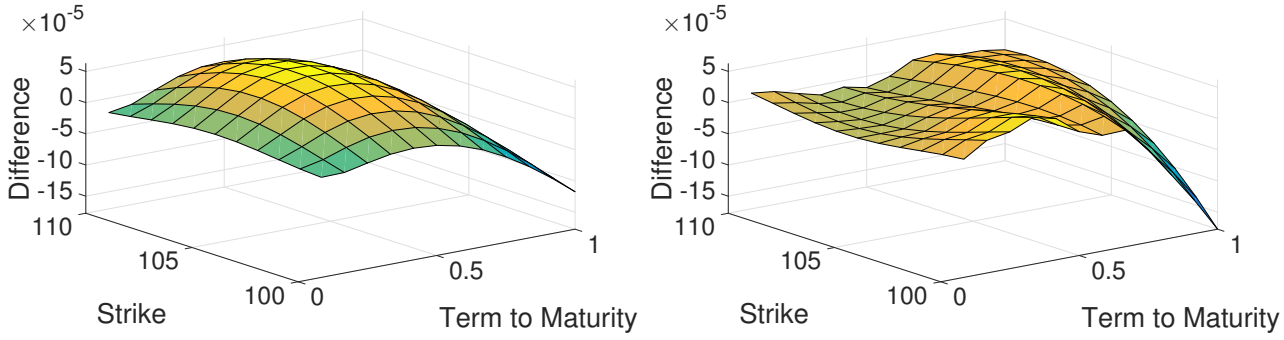

(a) Correlation matrix assumed known

(b) Correlation matrix unknown

Fig. 3.4: Difference between input and calibrated model option surface

3.3 Calibration to Market Data

This section describes the results of calibrating the model to market data. Section 3.3.1 describes the data used and Section 3.3.2 elaborates on the results of the calibration.

3.3.1 Market Data

This dissertation uses daily data spanning the period beginning on 01/01/2010 and ending on 31/12/2016. The data provider used was Bloomberg. The following data was used: USD cap volatility; EUR cap volatility; USD OIS curve; EUR OIS (EONIA) Curve; EURUSD spot exchange rate; and EURUSD option volatility data. A complete list of the Bloomberg tickers used can be found in Appendix B. The EUR cap volatility defines the Euro (EUR) Cap at-the-money (ATM) smile. The

EURUSD spot rate is the price of one EUR in terms of USD. The EURUSD option volatilities define the ATM smile for the option to buy one EUR for a fixed USD strike price.

The USD and EUR ATM cap volatilities are used to price the cap smile. The cap smile is then used to calibrate the domestic (EUR) and foreign (USD) short rate processes. The cap volatilities are used as input into Black's formula for caplets. The price of a cap is calculated as the sum of caplets struck at the prevailing fair swap rate for the relevant tenor structure $\mathcal{T}$.

Let us assume the same definitions as in Section 3.1.2 and denote the cap volatility by σ_{cap}. The time t Black price per unit nominal of a caplet can be written as

$$\text{BlkCaplet}(t, t_{i-1}, t_i, K, \sigma_{\text{cap}}) = \left(f(t, t_{i-1}, t_i)\Phi(d_1) - K\Phi(d_2)\right)(t_i - t_{i-1})P(t, t_{i-1}),$$

where

$$\begin{aligned} d_1 &= \frac{\ln[f(t, t_{i-1}, t_i)/K] + \frac{1}{2}\sigma_{\text{cap}}^2(t_{i-1} - t)}{\sigma_{\text{cap}}\sqrt{t_{i-1} - t}}, \\ d_2 &= d_1 - \sigma_{\text{cap}}\sqrt{t_{i-1} - t}. \end{aligned}$$

The OIS curves provide nominal annual, continuously compounded rates (NACC) for tenors ranging from one day to ten years. The OIS curves are used to create a ZCB curve. If the current time is $t = t_0$ and the OIS rate corresponding to a tenor of t_i is denoted $r(t, t_i)$, then the discount factor or ZCB $P(t, t_i)$ corresponding to the same tenor is:

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

If there is no quoted rate corresponding to the tenor of interest, then the ZCB of interest is estimated by linearly interpolating using the ZCB curve. If the OIS curve has observable OIS rates $r(t, t_i), r(t, t_j)$ for tenors t_i, t_j but one requires $P(t, s)$ such that $s \in (t_i, t_j)$ and $i < j$, then

$$P(t, s) = \frac{(s - t_i)P(t, t_j) + (t_j - s)P(t, t_i)}{t_j - t_i}.$$

This produces neither the smoothest nor an arbitrage-free ZCB curve, but the focus of this dissertation is not building an arbitrage-free ZCB curve. The resultant curve is sufficiently smooth for the purposes of the dissertation. The USD and EUR OIS rates are calculated using an ACT/360 day count convention. Note that the fixed leg of the interest rate swaps are calculated using a 30/360 day count convention. This is market convention.

The EURUSD option volatility data is used to price European-style FX call options, hereafter *call options*, according to Black's option formula. For a call option written on the EURUSD spot exchange rate S_t at time t, with strike K; maturity T; and volatility σ_{call}, Black's formula per unit nominal of the option is

$$\text{BlkCall}(t, T, K, \sigma_{\text{call}}) = S_t \exp\{-q(t,T)(T-t)\}\Phi(d_1) - K\exp\{-r(t,T)(T-t)\}\Phi(d_2),$$

where

$$d_1 = \frac{\ln[S_t \exp\{(r(t,T) - q(t,T))\,(T-t)\}/K] + \frac{1}{2}\sigma_{\text{call}}^2(T-t)}{\sigma\sqrt{(T-t)}},$$
$$d_2 = d_1 - \sigma_{\text{call}}\sqrt{T-t}.$$

Above, $r(t, t_i)$ is the continuously compounded (NACC) EUR (EONIA) OIS rate observed at time t for tenor $(t_i - t)$ and $q(t, t_i)$ is the NACC USD OIS rate observed at time t for tenor $(t_i - t)$.

3.3.2 Results of Calibration to Market Data

The short rate calibrations are conducted assuming that the prevailing OIS curves present proxies for the initial short rate and mean reversion level, so the calibration involves searching for κ_d (or κ_f) and ξ_d (or ξ_f). The same calibrations are repeated, but without implying the initial short rates and mean reversion levels. Thus the calibration involves searching for $r_0^d, \theta_d, \kappa_d$ and ξ_d (or $r_0^f, \theta_f, \kappa_f$ and ξ_f). The calibrated short rate parameters are then used to calibrate the model to the call option smile. The table in Appendix 3.3 displays the parameters resulting from the market calibration on 14 March 2014. The values marked with an asterisk (*) are implied from the relevant OIS curve.

The case where the OIS curves are used as proxies for initial short rate and mean reversion levels is called the *two parameter calibration* because two short rate parameters are being calibrated. When all four parameters of the short rate processes are being calibrated it is called the *four parameter calibration*.

Figure 3.5 shows the results of a calibration to the USD cap smile. The calibrated model cap smile for a two parameter calibration is plotted using circles ($\circ$) and the cap smile resulting from a 4 parameter calibration is plotted using asterisks ($*$). Figure 3.6 shows the resulting smiles when calibrating to the EUR cap smile. The markers are the same as in Figure 3.5. Although there are cap volatilities where the strike is at a premium to the fair swap rate, only ATM quoted cap volatilities are

used in the calibration of the USD and EUR short rate processes. Since only ATM caps are used in the calibration, there is a single strike rate per tenor and hence a cap smile results, not a surface.

It is clear that the CIR model for short rates is unable to match the market cap smile perfectly in both the USD and EUR case. In particular, the short end of the calibrated smile in the EUR and USD cases provides a poor match to the market. This is a flaw that can result in the model not being arbitrage-free. An arbitrage-free extension of the CIR model, the CIR++, is presented in Brigo and Mercurio (2001). The CIR++ model is constructed in such a way to match the term structure of interest rates exactly.

Figure 3.5 and Figure 3.6 display the ability of the four parameter calibration to match the cap smile more accurately than the two parameter calibration. This is not surprising, since the four parameter calibration allows more flexibility in the smile because of the extra two parameters. This suggests that the assumption of the OIS curves as proxies for the initial short rate and mean reversion levels is not conducive to accurate calibrations. The four parameter calibration results in the calibrated parameters losing economic interpretation. If the calibrated parameters for the USD (r_d process) are interpreted in the CIR parameter sense, then the USD short rate process has an initial value of 0.0015% and a mean reversion level of 1106.7550%. These values bear little resemblance to the values implied from the OIS curve of 0.08% and 2.476% for the initial short rate value and the mean reversion level implied from the USD OIS curve.

Figure 3.7 shows the results of the calibration to the call option smile following the short rate process calibrations mentioned above. Note that there is a market smile of prices and not a surface because the call option volatilities are quoted at-the-money (ATM) for varying tenors. In this particular two-part calibration, the calibration using the results from the two parameter calibration is able to match the observed call option smile just as well as the four parameter calibration case. Both calibrations match the entire observed market smile. The close fit in Figure 3.7 shows promise for hedging purposes. Not all calibrations to the call option smile are as successful as the one shown in Figure 3.7. In any event, the inability to match every point on the smile will permit arbitrage in the model.

Since the model is able to calibrate to the USD cap smile, EUR cap smile and consequently the call option smile, one is able to hedge options observed in the market.

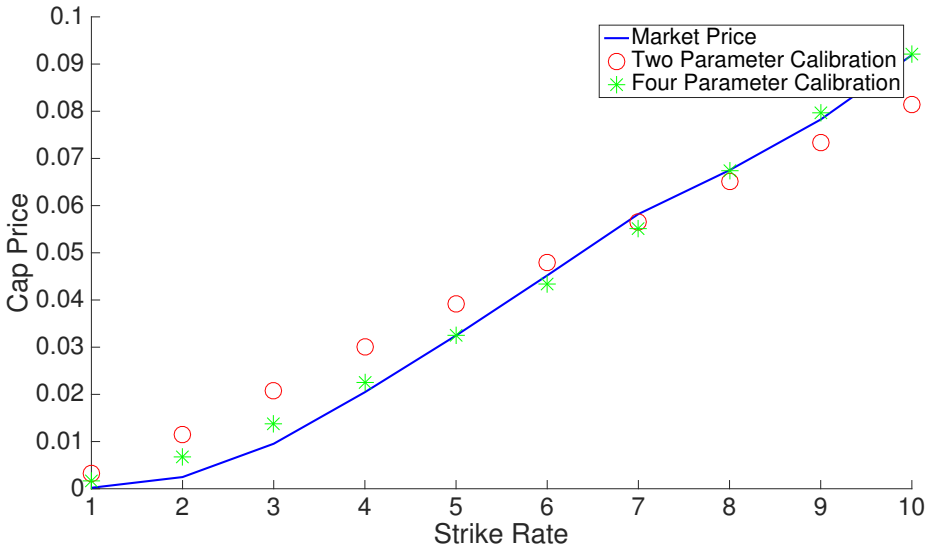

Fig. 3.5: USD market cap smile vs. calibrated model smile

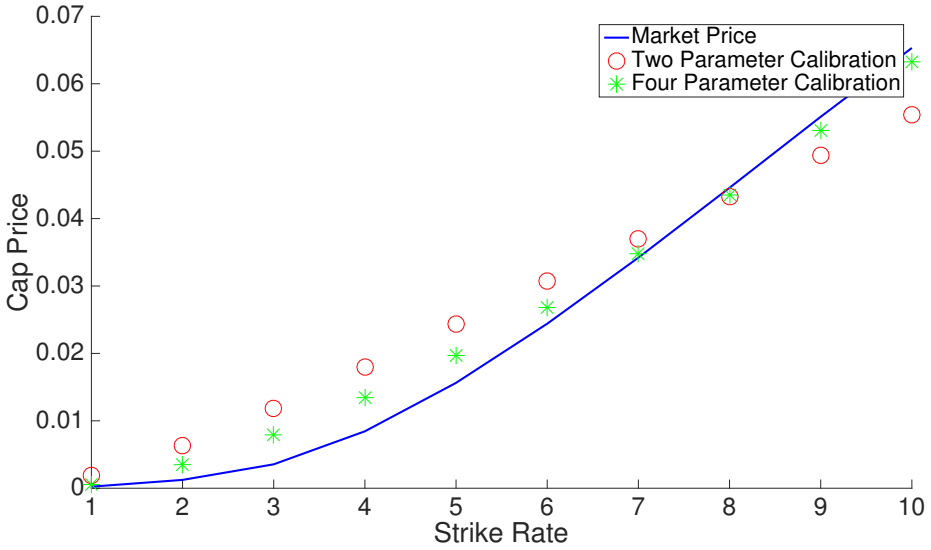

Fig. 3.6: EUR market cap smile vs. calibrated model smile

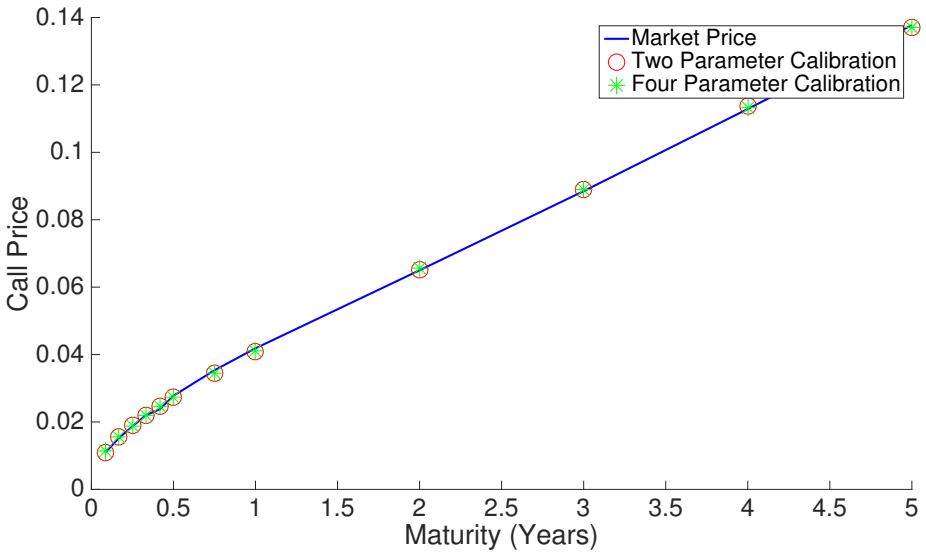

Fig. 3.7: Market call option smile vs. calibrated model smile

Parameter Name	Two Parameter Calibration	Four Parameter Calibration
r_0^d	0.0008*	0.00001
θ_d	0.02478*	11.0676
κ_d	0.3655	0.0004
ξ_d	0.4506	0.2676
r_0^f	0.0016*	0.00001
θ_f	0.0230*	18.3202
κ_f	0.1403	0.0002
ξ_f	0.2293	0.1065
v_0	0.0047	0.0057
θ	0.0048	0.0061
κ	4.3379	1.4849
ξ	0.08836	0.3916
ρ_{sf}	0.2681	-0.0196
ρ_{sd}	-0.0305	-0.0181
ρ_{sv}	-0.1780	-0.0329
ρ_{vf}	-0.00732	-0.0628
ρ_{vd}	-0.0419	-0.0406

Tab. 3.3: Parameter values resulting from market calibration

Chapter 4

Hedging

This chapter begins by describing the delta-hedging process. Thereafter, the results of delta-hedging European-style FX call options, hereafter *call options*, against model-generated data are presented. Subsequently, European-style EURUSD call options are hedged in the market. Section 4.1 describes the process and results of delta-hedging the call options against model-generated data. The purpose of the hedging exercise is to demonstrate the performance of delta-hedging in the model world, using mixed Monte Carlo and the calibration process. The same parameters are used in Section 4.1 as in Cozma and Reisinger (2015). The before-mentioned parameters are listed in Appendix A. Section 4.2 presents the results of delta-hedging EURUSD call options daily in the market.

4.1 Hedging Against Model-Generated Data

The input call option surface is generated using the same quasi-random numbers as in Chapter 3. The Monte Carlo calibration also uses the quasi-random numbers, as in Section 3.2.2. Pricing the input call option surface and Monte Carlo calibration both use a sample of 10^4 low discrepancy numbers.

In the Monte Carlo calibration, the parameter values from the calibration on the previous day are used as initial vectors for subsequent calibrations. Nine call options are hedged. All nine permutations of one of in-the-money (ITM), at-the-money (ATM) and out-the-money options with maturities of one month (1M), three months (3M) and six months (6M) are hedged (i.e., ATM-1M, ITM-3M etc.). The call options are delta-hedged daily in a self-financing portfolio. It is assumed that there are 20 working days per month, and hence 240 days per year. A time series of the calibrated parameters is recorded so that parameter stability can be evaluated over time.

The delta of a call option is the sensitivity of the value of the option to a change in the spot FX rate. The deltas of call options are calculated using a central difference method with a perturbation of $h = 10^{-5}$ to the spot FX rate, which is input into the mixed Monte Carlo pricing function, along with the other calibrated model parameters. If one represents the mixed Monte Carlo pricing function as $f(S_0, x_{cal})$ where S_0 is the prevailing spot FX rate and x_{cal} represents a vector of calibrated model parameters, then the delta (Δ) of the call option is calculated as

$$\Delta = \frac{f(S_0 + h, x_{cal}) - f(S_0 - h, x_{cal})}{2h}.$$

The hedge portfolio is constructed to be self-financing. A portfolio is self-financing if there is no withdrawal or deposit of cash from the portfolio, thus any change in value of the portfolio is driven by changes in the spot FX rate or the return earned on a ZCB. The process of delta-hedging the call option is detailed below and displayed in Algorithm 1. In this exercise, one assumes the ability to trade in fractions of a share. In addition, transaction costs are ignored and it is assumed possible to borrow and lend at the same rate, which is equivalent to selling or buying a ZCB at the same rate.

Consider a call option that has value denoted by X, on share S. The sensitivity of X to a change in S is the delta of the option. The value of delta of the option is denoted by Δ. The option under consideration has a strike of K and a maturity of T. Let there be n discrete time steps from the initiation of the option (t_0) until the maturity of the option $T = t_n$, i.e., time between now and the maturity of the option is discretised into $\{t_0, t_1, ..., t_{n-1}, t_n = T\}$. The values of the share, option and delta of the option at t_i are S_i, X_i and Δ_i, respectively. At each t_i there is a market for zero coupon bonds available for maturities t_j for $j \in \{i+1, i+2, ..., n-1, n\}$. A zero coupon bond pays one upon maturity. At t_i the value of a zero coupon bond with maturity t_j is represented by $ZCB_{i,j}$.

At t_0 the seller of the options sells one option on S for X_0. The seller uses the proceeds from the sale to purchase Δ_0-many shares, which costs the seller $\Delta_0 \cdot S_0$. The difference of $\Delta_0 \cdot S_0 - X_0$ is made up by selling $\frac{\Delta_0 \cdot S_0 - X_0}{ZCB_{0,1}}$-many zero coupon bonds for maturity t_1.

At t_i for $i \in \{1, 2, ..., n-1\}$ the seller wishes to hold Δ_i-many shares. Thus, the seller liquidates (sells) his or her holding in Δ_{i-1}-many shares for $\Delta_{i-1} \cdot S_i$ in cash. This cash is used to purchase Δ_i-many shares for a cost of $\Delta_i \cdot S_i$ as well as settle the outstanding zero coupon bonds with maturity t_i. Any shortfall (excess) is netted,

made zero, by selling (buying) zero coupon bonds with maturity t_{i+1}.

At t_n the seller of the option owes the purchaser $(S_n - K)^+$ i.e. the greater of $S_n - K$ and 0. This obligation is funded by liquidating the holding in Δ_{n-1}-many shares for $\Delta_{n-1} \cdot S_n$ cash as well as the position in zero coupon bonds with maturity t_n. Any shortfall (surplus) is recorded as a loss (profit).

Algorithm 1 contains pseudocode that displays the hedging procedure as well as how the PnL is recorded. For this illustration, *holding* is the number of shares of S currently held; $call_i$ is the price of the call option at time t_i; $delta_i$ is the delta of the call at time t_i; *bank* represents the value of the holding in either cash or ZCBs (this can be negative if one shorts ZCBs); *port* is the value of the hedge portfolio; $ZCB_{i,j}$ is the price of a ZCB at time t_i that matures at time t_j; S_i is the share price at time and *PnL* is the resulting profit or loss from the hedge. The indices i represent the time steps t_i

Algorithm 1: Algorithm for delta-hedging a call option

```
for i ← 0 to n do
    if i = 0 then
        bank = call_i;
        holding = delta_i;
        bank = bank - delta_i × S_i;
        port = holding×S_i + bank;
        bank = bank/ZCB_{i,i+1};
    else if 1 ≤ i ≤ n − 1 then
        bank = bank + holding × S_i - delta_i×S_i;
        holding = delta_i;
        port = holding×S_i + bank;
        bank = bank/ZCB_{i,i+1};
    else
        port = holding×S_i + bank;
        PnL = port - max(S_i - K,0);
    end
end
```

4.1.1 Hedging with Full Knowledge of Parameters

In this subsection, the nine options mentioned in Section 4.1 are delta-hedged according to the method outlined in Algorithm 1. The resultant profit and loss (PnL) is recorded for 10^4 simulations of the experiment in order to build a PnL frequency distribution. The frequency distribution displays the hedge slippage under the ideal scenario of exact calibration. The realisation of 10^4 PnLs takes approximately 1.5 days of computing time.

The options are delta-hedged with full knowledge of all parameters. This entails knowledge of the volatility parameters, domestic short rate parameters, foreign short rate parameters and the correlation matrix. This experiment provides benchmark PnL distributions, shown in Figure 4.1, that display the hedge-slippage one expects from using the four-factor model, mixed Monte Carlo and delta-hedging. The PnL distributions in Figure 4.1 present the best possible performance of delta-hedging portfolio under this model before the confounding effects of calibration.

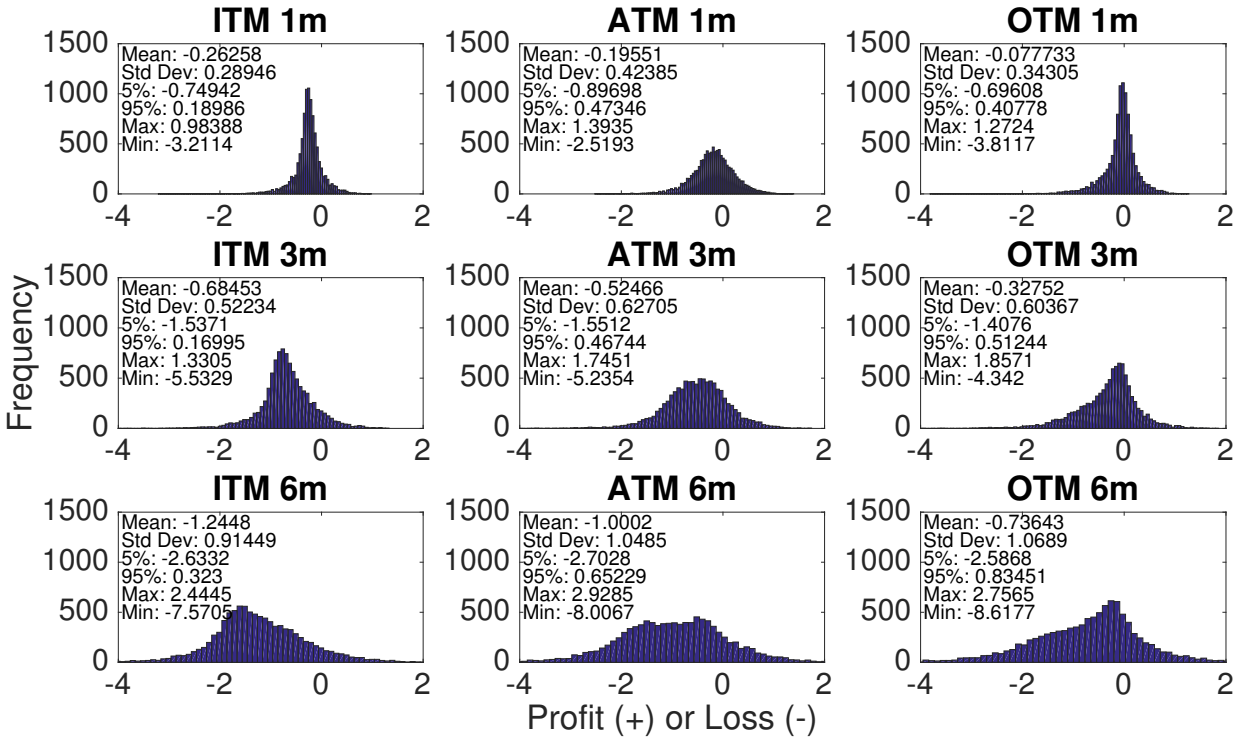

Fig. 4.1: PnL from daily delta-hedging with known parameters

A single path of the delta-hedging strategy is shown in Figure 4.2 for each of the options considered. In this realisation, all of the hedging strategies exhibit a loss

on maturity of the option. The next step is to include calibration into the hedging exercise. Calibration of the short rate processes and volatility process is introduced to the exercise in Section 4.1.2.

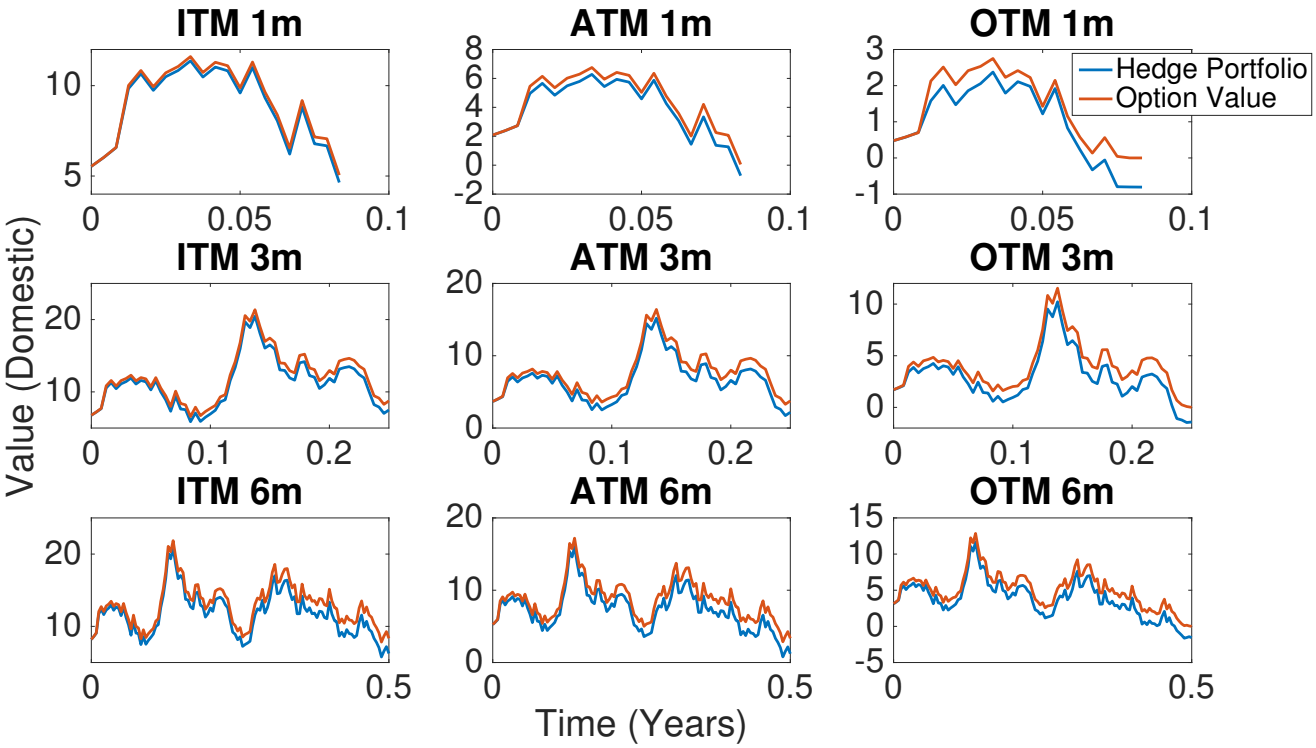

Fig. 4.2: Time series of hedge portfolio vs. option value

4.1.2 Hedging with Knowledge of Correlation Matrix

In this subsection, the nine options mentioned in Section 4.1 are delta-hedged according to the method outlined in the same section. The resultant profit and loss is recorded for a single realisation of the experiment. An individual realisation of PnL takes approximately 13 hours to complete and thus it is not practical to create a PnL distribution as in Section 4.1.1.

The options are hedged with knowledge of the correlation matrix and the remaining parameters are calibrated as in Chapter 3. The correlation matrix is assumed known because in Table 3.2 it is clear that, through the Monte Carlo calibration, one is able to calibrate the volatility parameters accurately, but not the elements of the correlation matrix. The inaccuracy resulting from calibrating for the elements of the correlation matrix confounds the overall hedging results.

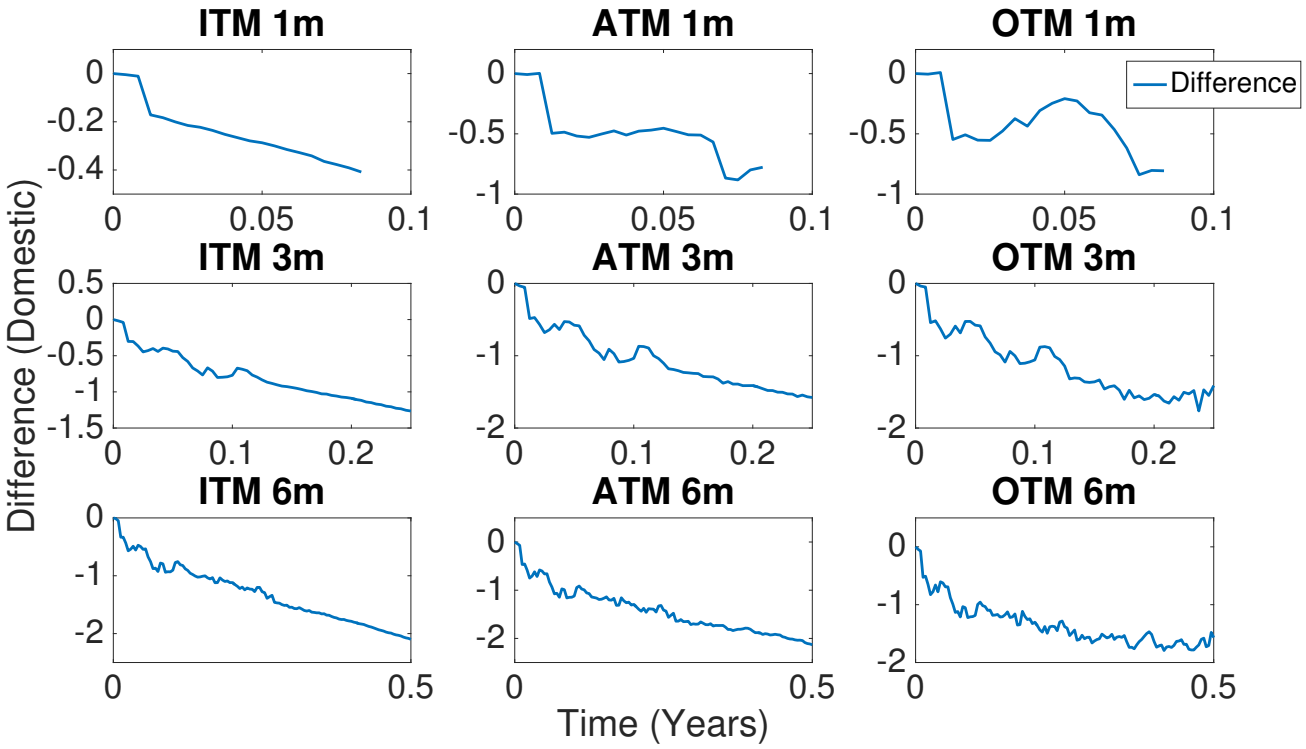

Fig. 4.3: Time series of difference between hedge portfolio and option value

A time series of the calibrated parameter values is recorded and plotted below in Figure 4.1.2 (a) and (b) to display the stability of the calibrated parameters over the life of the 6 month option. Figure 4.1.2 (a) shows the calibrated values compared with the market values of the domestic short rate parameters over the six month life of the longest-dated option being hedged. The results for the foreign short rate are identical. The input parameters are plotted using solid lines and the calibrated parameter values are plotted using dots that are the same colour as the corresponding input parameter. It is clear that the calibration is able to recover the model parameters to an acceptable degree of error on each day of the hedging exercise. From Figure 4.1.2 (b) one observes the relatively stable calibration of the volatility process. The calibrated parameter values occasionally drift from the market parameter values, but return to the market value almost immediately.

Since the calibration process appears stable for a known correlation matrix, the calibration is extended to include the elements of the correlation matrix in Section 4.1.3.

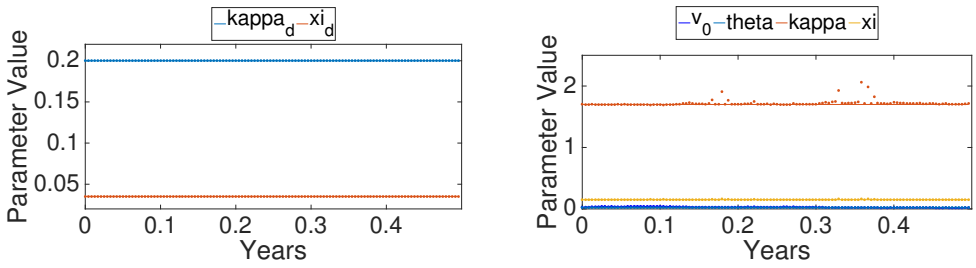

(a) Calibrated domestic short rate parameter values

(b) Calibrated volatility parameter values

Fig. 4.4: Monte Carlo calibrated parameter values when the correlation matrix is known

4.1.3 Hedging with No Knowledge of Parameters

In this subsection, the nine options mentioned in Section 4.1 are delta-hedged according to the method outlined in the same section. The resultant profit or loss is recorded for a single simulation of the experiment. It is not practical to create a PnL distribution since an individual realisation of PnL takes approximately one day to complete.

The options are hedged with no previous knowledge of any parameters. All parameter values are calibrated as described in Chapter 3.

The calibrated parameters values are recorded and plotted over time to display the calibration stability. Two cases of parameter stability are investigated and compared: when the parameters resulting from previous calibration are used as initial guesses for the calibration; and when a standard initial guess of $\mathbf{x}_0 = [0.1, 0.1, 1, 0.1, 0, 0, 0, 0, 0]^{\mathrm{T}}$ for subsequent calibrations. Both calibrations take equally long to complete a full realisation and the resulting time series are almost identical.

The calibration is consistently able to recover both parameters for the short rate processes as before. The calibrated volatility parameters are stable and accurate, in comparison to the market parameters, until approximately 4.5 months after inception. At this point the calibrated values for the speed of mean reversion κ drift from the input value and does not return. This diversion from the market parameter values is visible in Figure 4.1.3. Figure 4.1.3 (b) shows the calibrated values of the correlation matrix compared with the input values. Two out of five elements of the correlation matrix are calibrated accurately, with some stability, over time.

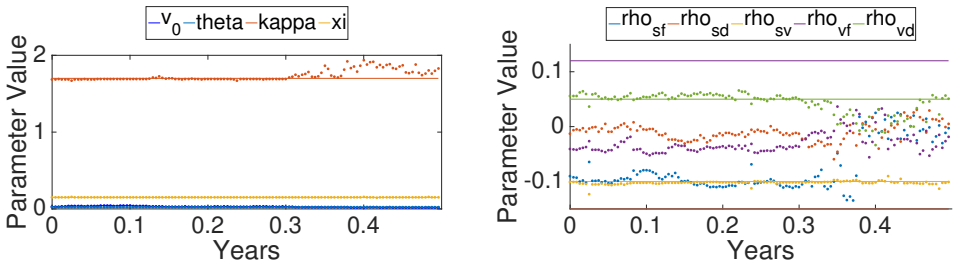

(a) Time series of volatility parameters

(b) Time series of calibrated correlation coefficients

Fig. 4.5: Time series of Monte Carlo calibrated parameters

Again, the stability is disrupted 4.5 months after inception, as with the κ parameter in Figure 4.1.3 (a). After 4.5 months, there appears to be no accuracy or stability in the calibrated values except for the correlation ρ_{sv} between the Brownian motion driving the spot FX rate and the Brownian motion that drives the volatility of the spot FX rate. This suggests that stability in the Monte Carlo calibration is distorted over time.

While the calibration is not able to recover all correlation coefficients accurately, the calibration is able to match the call option surface to an acceptable degree of error. This is sufficient to hedge. The full calibration exercise is extended to market data in Section 4.2.

4.2 Hedging Against Market Data

The first subsection, Section 4.2.1, describes the results of calibrating the model to market data for the 279 business days starting on 3 January 2011 and ending on 26 January 2012. Thereafter, the results of hedging 1M, 3M and 6M ATM EURUSD European-style FX call options, hereafter *options*, over the period 3 January 2012 to 3 July 2012 are explored in Section 4.2.2.

4.2.1 Results of Calibrations

The four processes and correlation matrix are calibrated daily to the prevailing USD and EUR cap smile and consequently the observable call option smile. The short rate calibrations involve calibrating for all four parameters. The initial rate and mean reversion level are not implied from the ZCB curve. The time series of SSDs arising from the calibration to the cap smiles and the call option smiles are shown

in Figure 4.6. In the case of the call option market, since there is a large spread of SSDs over the year, with one extreme outlier, box plots of the SSDs are shown in Figure 4.7. Each box plot is made up of SSDs obeying a shrinking upper limit. The number of points in the sample is denoted by n.

From the box plots it is evident that 218 out of the 279 calibrations to the call option smile reached SSDs of less than 0.001, with approximately 75% of those being less than 0.0001. Hence, the calibration can result in a calibrated model smile that matches the observed market option smile to an acceptable degree of error.

The accuracy of the Monte Carlo call smile calibrations vary. The most accurate calibration obtains an SSD of 4.249×10^{-7} whereas the least accurate calibration attains an SSD of 15.510, which is an outlier since the next largest SSD is approximately 5. The majority of the call smile calibrations, shown in Figure 4.6 (a), reach SSDs less that 10^{-4}.The results of the most accurate calibration are shown in Figure 4.12.

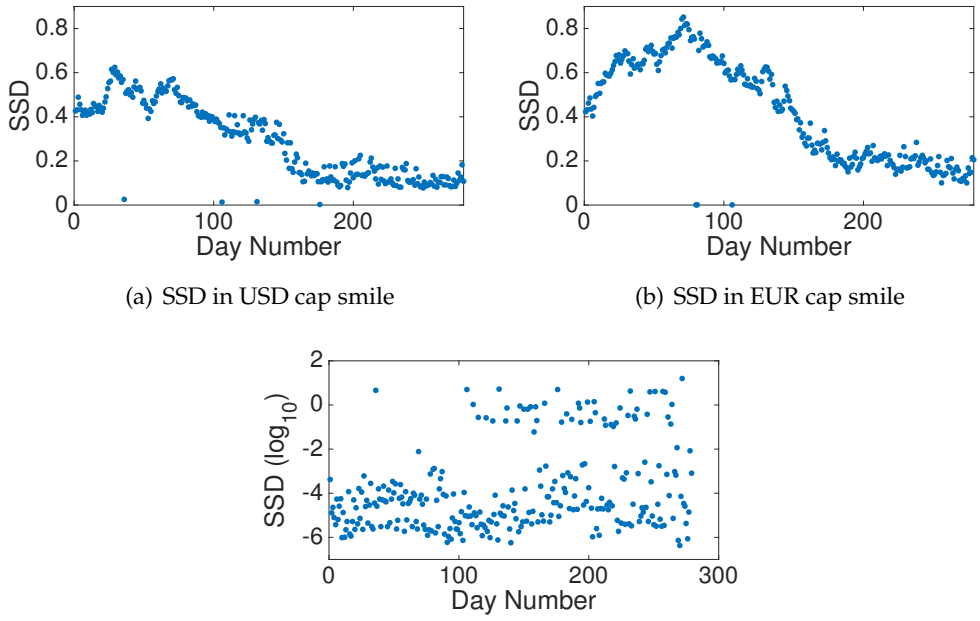

(a) SSD in USD cap smile

(b) SSD in EUR cap smile

(c) SSD in option smile

Fig. 4.6: Time series of SSD for each instrument calibrated

The calibrated parameter values are recorded and the stability is investigated. Each calibration begins from an arbitrary, but constant, initial guess. The calibrated pa-

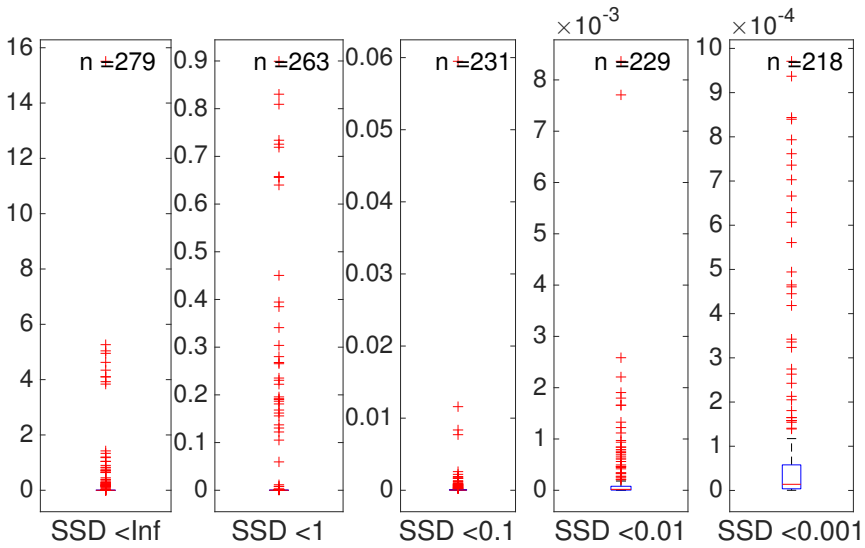

Fig. 4.7: Box plot of SSD for option smile calibration

rameter values from the previous day are not used as an initial guess for the calibration. This approach is used so that a poor calibration does not result in bias for subsequent calibrations. The time series of calibrated parameter values are displayed in Figures 4.8, 4.9, 4.10 and 4.11.

The upper bound for all the non-correlation coefficient parameters is set at 100; the lower bound is set at 0. The upper bound is arbitrarily chosen in order to allow a reasonably unconstrained calibration in line with initial calibrations. It is evident that in some cases during the 279 day period that the calibrated parameters reach this upper bound. If parameters such as the initial short rate (or volatility) or mean reversion levels reach this upper bound, the economic interpretation of the parameters is lost since a parameter value of 100 equates to 10 000%.

When inspecting Figure 4.8 it is crucial to note the $\log_{10}$ scale. There appears to be a trend in each of the four parameters being calibrated. This shows promise in terms of parameter stability, except that there are many data points that - frequently - do not follow the trends.

In Figure 4.9 there is potential for a trend in the initial short rate and the speed of mean reversion. The trends look smooth - in $\log_{10}$ scale - until the halfway mark where both sequences experience jumps.

The time series for the initial volatility level parameter appears the most stable out the of the volatility process parameters. This holds until halfway through the period where the parameter values - along with those for the other parameters - become erratic.

Figure 4.11 shows the time series of calibrated correlation coefficients. The mean of each calibrated correlation coefficient is plotted in red. It is visible from Figure 4.11 that the calibrated values for all of the correlation coefficients fluctuate drastically during the period of observation. It is interesting to note that the coefficients very rarely reach the upper and lower bounds of 1 and -1, respectively. Correlations between Brownian motions are notoriously difficult to estimate. In practice these correlations would most likely be estimated from historic data in a similar way to how ρ_{df} is estimated in this dissertation. It was decided to calibrate the correlation coefficients instead. One can be sceptical of the quality of the calibrated correlation coefficients because there is no stability in any of the calibrated values, i.e., the sign of a calibrated correlation coefficient can change from day to day. This is not detrimental to the results because the aim is not coefficient discovery, but rather to fit the option smile. One is able to do so in the vast majority of the cases under investigation according to Figure 4.6 and the box plots in Figure 4.7. Since one is able to match the observed market smile, one is able to hedge.

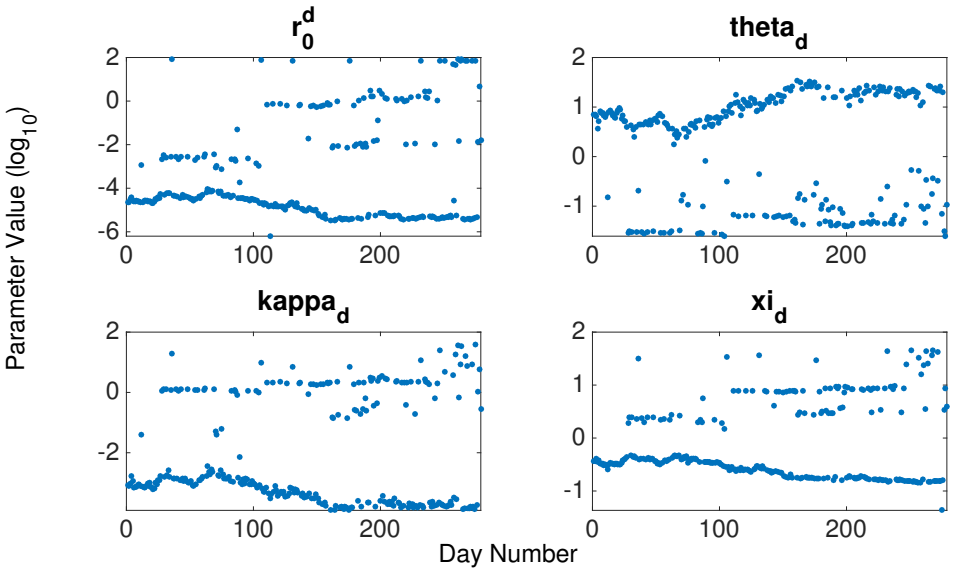

Fig. 4.8: Time series of calibrated USD short rate parameter values

As mentioned previously, the calibration to the call smile that resulted in the small-

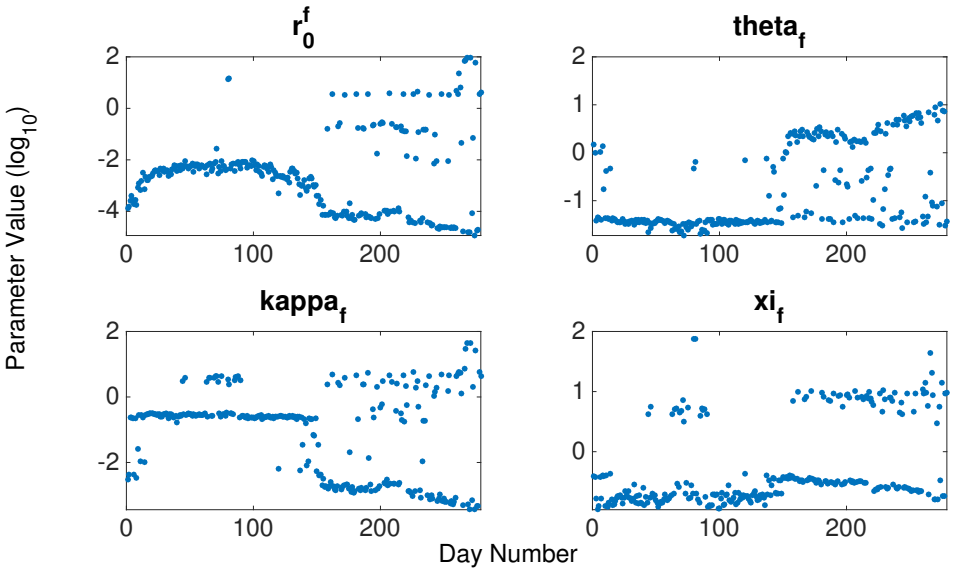

Fig. 4.9: Time series of calibrated EUR short rate parameter values

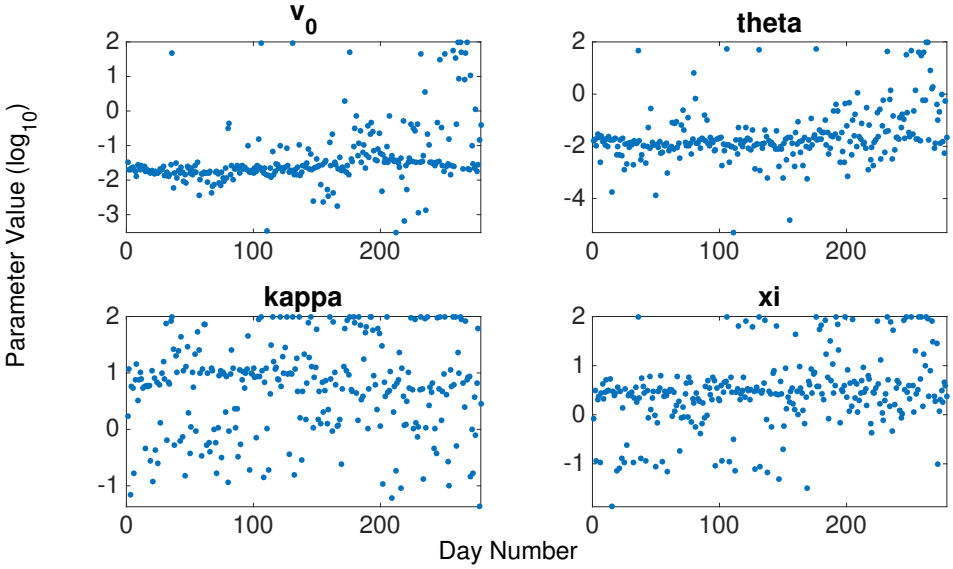

Fig. 4.10: Time series of calibrated volatility parameter values

est SSD is displayed in 4.12. Even though the calibrated model call prices match the observed market prices the closest out of the 279 observations, the short rate calibrations do not result in calibrated model cap prices that lie on the market smile.

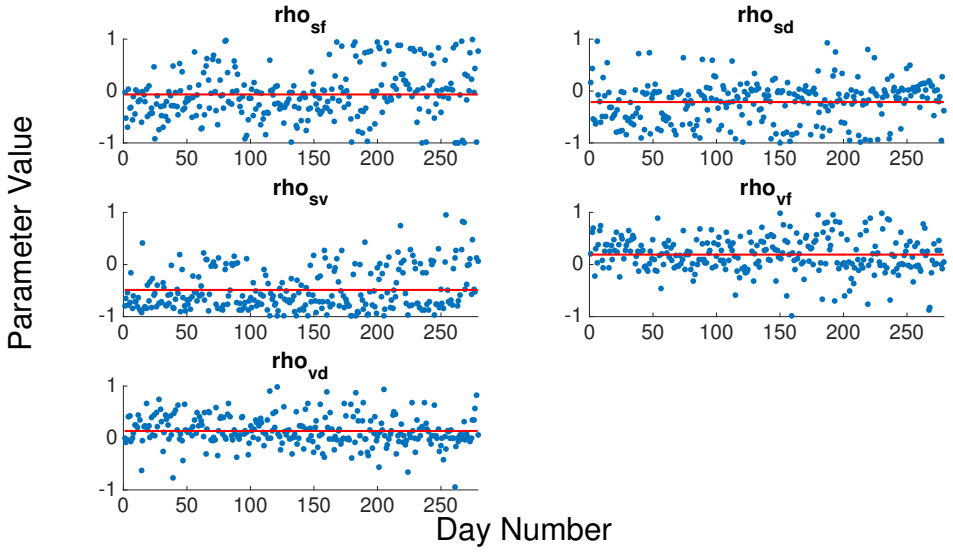

Fig. 4.11: Time series of correlation matrix values

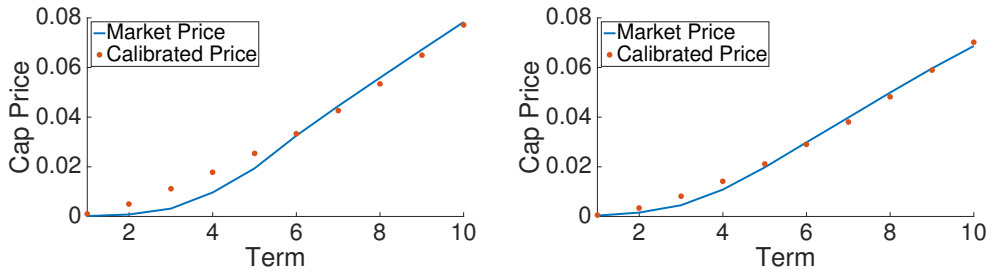

(a) USD cap calibration that results in best option calibration

(b) EUR cap calibration that results in best option calibration

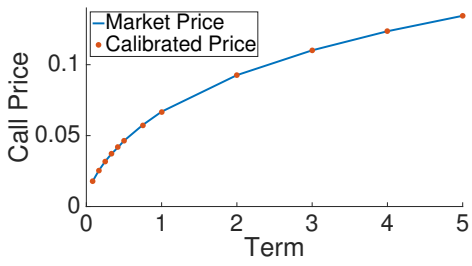

(c) Results of best option smile calibration

Fig. 4.12: Time series of SSD for each instrument calibration

4.2.2 Results of Hedging

This subsection describes the results of delta-hedging 1M, 3M and 6M ATM options. The options were sold on 3 January 2011 and expired on 3 February 2011, 4 April 2011 and 4 July 2011. It is assumed that there is no bid/ask spread and that there are no transaction costs.

The initial contracts are sold at a price of USD 2.13, USD 4.11 and USD 5.31 for the right but not the obligation to buy EUR100 for USD 133.61 on the expiry of each option. The hedged portfolios, of EUR in cash and USD-denominated ZCBs, are rebalanced daily to ensure a holding of Δ-many EUR. The Δs of the call options are calculated using a central finite difference scheme with a perturbation of 10^{-5} and mixed Monte Carlo pricing. The calibrated model parameters are used as inputs to the mixed Monte Carlo pricing function. The risk of a poorly constructed hedge arises from an inaccurate calibration to the call option smile on a given day.

In all three cases of the options sold, and consequently hedged, on 3 January 2011 a profit is realised at the expriry of the options. The profits realised are USD 0.38, USD 0.79 and USD 0.26 per contracts of size EUR 100. The values of the hedge portfolios in comparison to the value of the options, which are valued using the calibrated parameters, are shown in Figure 4.13. Upon inspection of the performance of the 6M hedge portfolio, four outlying option values are clear. There are two spikes that are significantly larger (referred to hereafter as *upwards outliers*) and two consecutive option values that are smaller (referred to hereafter as *downwards outliers*) than the neighbouring preceding and following option values. Inclusion of these outliers makes it difficult to evaluate the performance of the hedge. The upwards outliers occur on the 21 February and 30 May, respectively, and the downwards outliers occur on the 22 and 25 of April.

The SSDs of the call smile calibration on the 21 February and 21 May (the days of the upwards outliers) are 4.6244 and 5.0404 respectively. These two SSDs are the two largest of the period 3 January 2011 – 4 July 2011. The outlying option values can be explained as a result of poor calibration to the call option smile on those dates.

The SSDs of the call smile calibration on the 22 and 25 April are 0.0012 and 0.0013, respectively. These SSDs are over 1000 times smaller than those of the upwards outliers, hence explaining the smaller deviation, in terms of magnitude. The downwards outliers do form part of the largest 8% of SSDs for the period. Thus, the

outlying option values can again be explained by a poor fit of the calibrated option smile to the market call option smile.

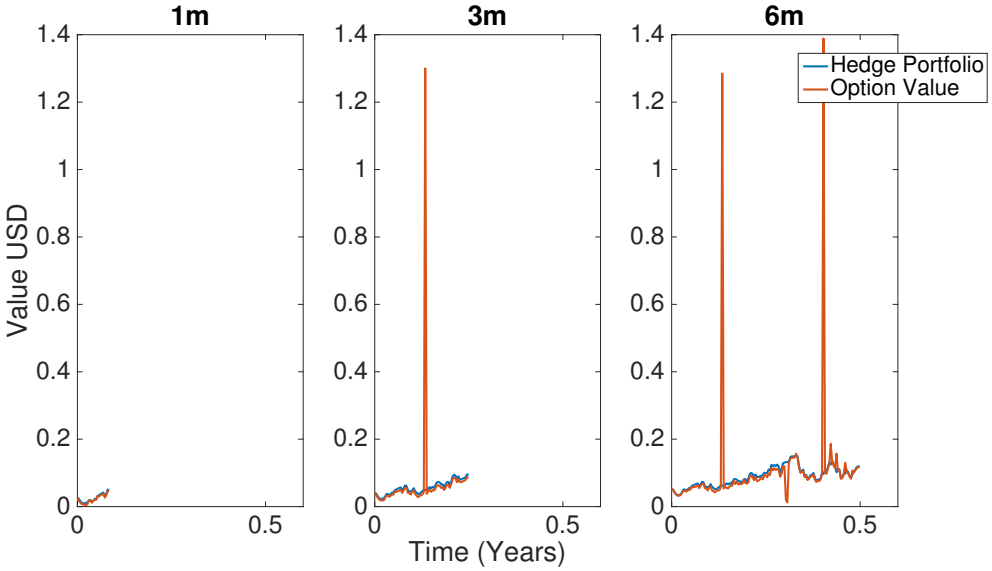

Fig. 4.13: Hedge portfolios of options sold on 3 January 2011

If one removes the four before mentioned outliers from the plot, the resulting hedge performance becomes as displayed in Figure 4.14. One can see that the hedge portfolio moves closely in line with the value of the option. Additionally, the differences between the hedge portfolios and the option values are shown in Figure 4.15.

All three of the options expire ITM, yet the seller of the option makes a profit on the hedge portfolio. The terminal spot FX prices are USD 136.34, USD 142.21 and USD 145.39 per EUR 100. The strike for all three options is USD 133.61 per EUR 100. The evolution of the EURUSD spot FX rate is displayed in Figure 4.16.

The delta of each option is recorded and plotted in Figure 4.17. The deltas arising on upwards and downwards outliers are plotted using asterisks ($*$).

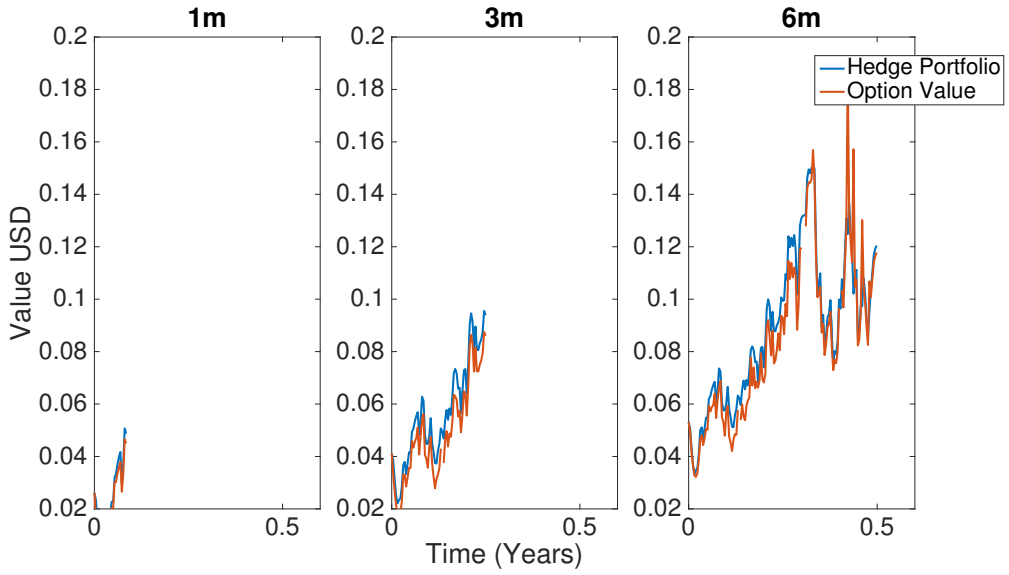

Fig. 4.14: Hedge portfolios of options sold on 3 January 2011 without outliers

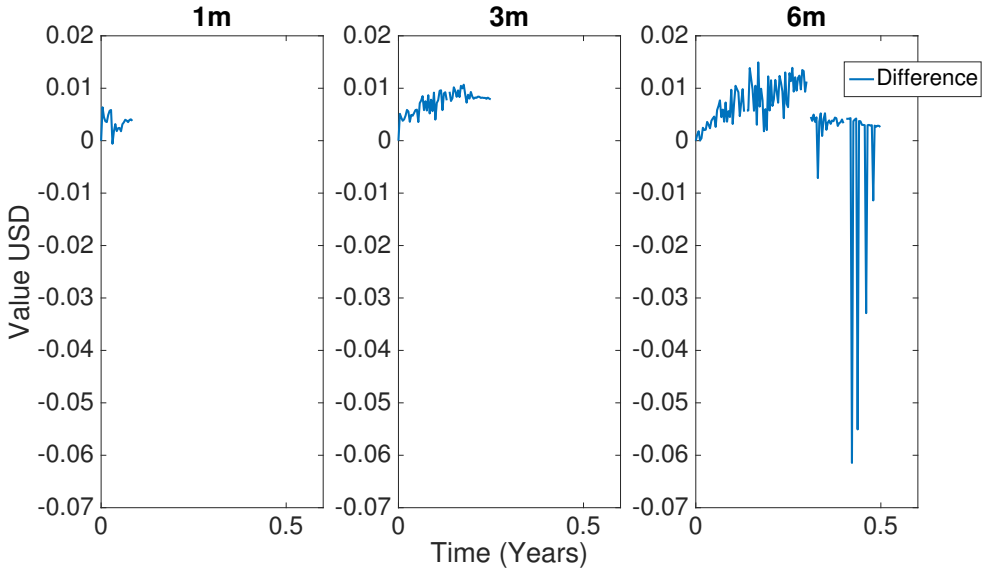

Fig. 4.15: Difference between hedge portfolios and value of options sold on 3 January 2011

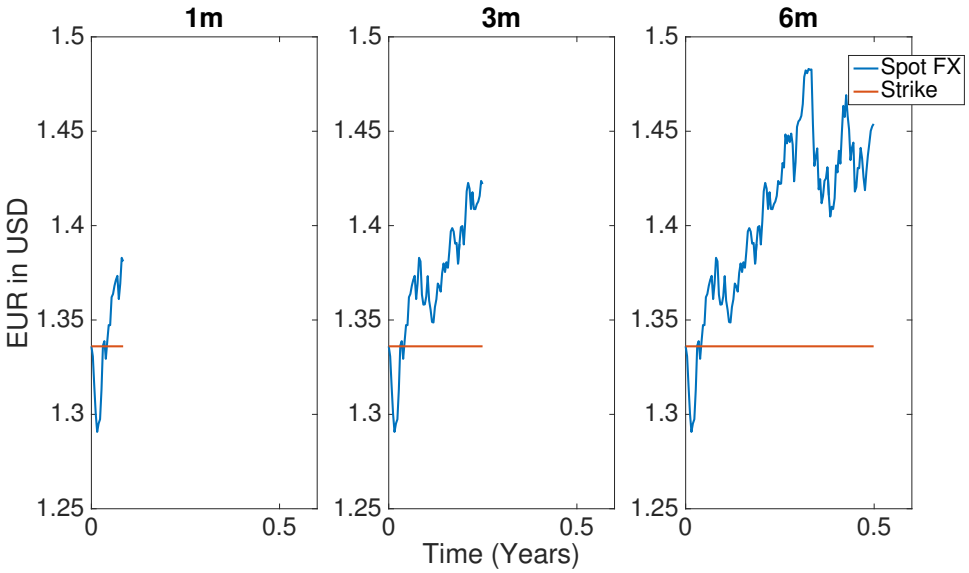

Fig. 4.16: Evolution of the EURUSD spot FX rate from 3 January 2011

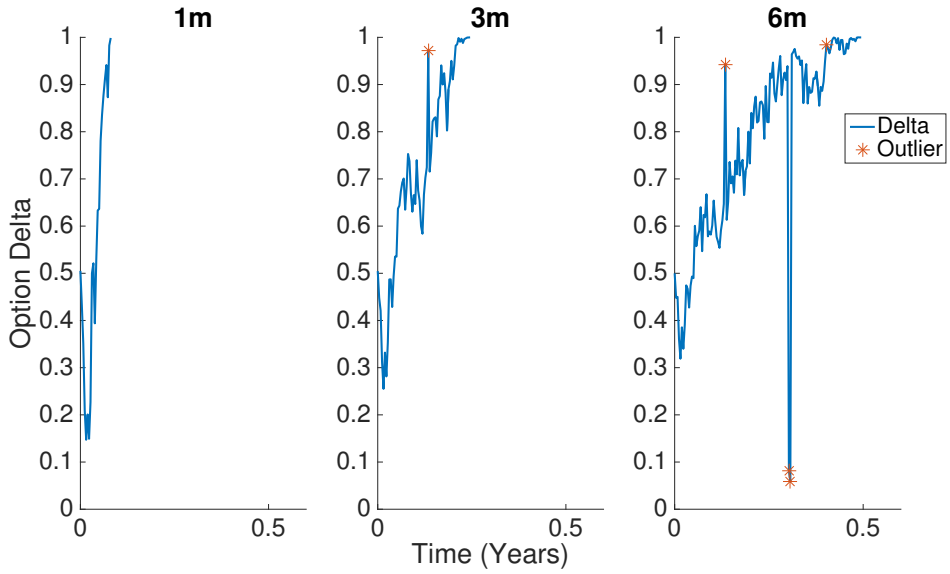

Fig. 4.17: Evolution of the delta of options sold on 3 January 2011

Chapter 5

Conclusion

One is able to calibrate the four-factor Heston-CIR-CIR model, described in Cozma and Reisinger (2015) in a two-part calibration process. The first part is a calibration to the domestic and foreign cap surfaces, under the important assumption of negligible country risk. The second part is Monte Carlo calibration to the FX call option surface using mixed Monte Carlo pricing. One can consistently recover input parameter values for the short rate processes as well as the volatility process, when calibrating to model-generated data. The above-mentioned two-step calibration is unable to recover all of the elements of the correlation matrix. The correlation matrix should be estimated using an alternative method.

The length of time to perform a full calibration is a hindrance to producing a PnL distribution for daily calibration. The Monte Carlo calibration is the main contributor to the duration of the calibration. This is not a problem in practice since one requires a full calibration to be conducted once a day.

It is possible to calibrate the model to EUR cap, USD cap, EURUSD option and EURUSD FX data. The quality of the cap smile calibrations is consistent with no obvious outliers over the investigated period. This is not the case for the option smile calibration. There are a few poor calibrations, but in 75% of the market calibrations the SSDs that result from calibrating to the FX option smile are below 10^{-5}. Thus, the model is able to match the option market smile to an acceptable degree of accuracy in the majority of the days investigated.

Since one is able to fit the model to the market option surface, one is also able to delta-hedge and value market-observed options using mixed Monte Carlo pricing. The hedge portfolio follows the value of the option contracts closely, except when there in the event of a poor calibration on a given day.

Future work can look at alternative - more accurate - methods of estimating the parameter values of the correlation matrix. Reducing the number of variables being sought may stabilise the Monte Carlo calibration, which proved accurate when calibrating for volatility process parameters only. Furthermore, one could investigate faster methods of calibrating such a model so that PnL distributions can be created when calibrating, and subsequently hedging, the model against model-generated data. Finally, there is potential to incorporate the CIR++ model into the four-factor model so as to prevent arbitrage.

Bibliography

Ahlip, R. and Rutkowski, M. (2013). Pricing of foreign exchange options under the Heston stochastic volatility model and CIR interest rates, *Quantitative Finance* **13**(6): 955–966.

Billingsley, P. (1995). Probability and Measure. Wiley Series in Probability and Mathematical Statistics.

Brigo, D. and Mercurio, F. (2001). *Interest Rate Models - Theory and Practice*, Springer Finance, Springer.

Carr, P., Gabaix, X. and Wu, L. (2009). Linearity-Generating Processes, Unspanned Stochastic Volatility, and Interest-Rate Option Pricing, *Working Paper* .

Cox, J. C., Ingersoll Jr, J. E. and Ross, S. A. (1985). A theory of the term structure of interest rates, *Econometrica: Journal of the Econometric Society* pp. 385–407.

Cozma, A. and Reisinger, C. (2015). A mixed Monte Carlo/PDE variance reduction method under the Heston-CIR model, *arXiv preprint arXiv:1509.01479* .

Glasserman, P. (2003). *Monte Carlo methods in financial engineering*, Vol. 53, Springer Science & Business Media.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of financial studies* **6**(2): 327–343.

Jamshidian, F. (1989). An Exact Bond Option Formula, *The Journal of Finance* **44**(1): 205–209.
URL: *http://dx.doi.org/10.1111/j.1540-6261.1989.tb02413.x*

Loeper, G. and Pironneau, O. (2009). A mixed PDE/Monte-Carlo method for stochastic volatility models, *Comptes Rendus Mathematique* **347**(9): 559–563.

McWalter, T. (2016). *Numerical Methods in Finance*, African Institute of Financial Markets and Risk Management (AIFMRM), University of Cape Town.

Ouwehand, P. (2015). *Interest Rate Modelling*, African Institute of Financial Markets and Risk Management (AIFMRM),University of Cape Town.

Van Haastrecht, A., Lord, R., Pelsser, A. and Schrager, D. (2009). Pricing long-maturity equity and FX derivatives with stochastic interest rates and stochastic volatility, *Insurance: Mathematics and Economics* **45**(3).

Appendix A

Cozma Reisinger Parameters

Parameter Name	Parameter Value
S_0	105
r_0^f	0.0291
κ_f	0.32
θ_f	0.0248
ξ_f	0.0317
r_0^d	0.0524
κ_d	0.2
θ_d	0.0475
ξ_d	0.0352
v_0	0.0275
θ	0.0232
κ	1.7
ξ	0.15
ρ_{sf}	-0.15
ρ_{sv}	-0.1
ρ_{sd}	-0.15
ρ_{df}	0.25
ρ_{vd}	0.12
ρ_{vf}	0.05

Tab. A.1: Parameters used in Cozma and Reisinger (2015)

Appendix B

Market Data

The market data used in this dissertation spans from 01/01/2010 to 31/12/2016. The data is extracted from Bloomberg with following reference codes. The description of the reference codes is also supplied.

B.1 Discount Curves

Bloomberg Code	Description
YCSW0042 Index	USD OIS Curve
YCSW0133 Index	EUR OIS (EONIA) Curve

Tab. B.1: OIS Discount Curve tickers

B.2 USD Cap Surface

Bloomberg Code	Description
USCV1 Curncy	USD ATM Cap Vol 1Y
USCV2 Curncy	USD ATM Cap Vol 2Y
USCV3 Curncy	USD ATM Cap Vol 3Y
USCV4 Curncy	USD ATM Cap Vol 4Y
USCV5 Curncy	USD ATM Cap Vol 5Y
USCV6 Curncy	USD ATM Cap Vol 6Y
USCV7 Curncy	USD ATM Cap Vol 7Y
USCV8 Curncy	USD ATM Cap Vol 8Y
USCV9 Curncy	USD ATM Cap Vol 9Y
USCV10 Curncy	USD ATM Cap Vol 10Y
USCV101 Curncy	USD ATM Cap Vol 1Y 1% premium
USCV102 Curncy	USD ATM Cap Vol 2Y 1% premium
USCV103 Curncy	USD ATM Cap Vol 3Y 1% premium
USCV104 Curncy	USD ATM Cap Vol 4Y 1% premium
USCV105 Curncy	USD ATM Cap Vol 5Y 1% premium
USCV106 Curncy	USD ATM Cap Vol 6Y 1% premium
USCV107 Curncy	USD ATM Cap Vol 7Y 1% premium
USCV108 Curncy	USD ATM Cap Vol 8Y 1% premium
USCV109 Curncy	USD ATM Cap Vol 9Y 1% premium
USCV1010 Curncy	USD ATM Cap Vol 10Y 1% premium
USCV201 Curncy	USD ATM Cap Vol 1Y 2% premium
USCV202 Curncy	USD ATM Cap Vol 2Y 2% premium
USCV203 Curncy	USD ATM Cap Vol 3Y 2% premium
USCV204 Curncy	USD ATM Cap Vol 4Y 2% premium
USCV205 Curncy	USD ATM Cap Vol 5Y 2% premium
USCV206 Curncy	USD ATM Cap Vol 6Y 2% premium
USCV207 Curncy	USD ATM Cap Vol 7Y 2% premium
USCV208 Curncy	USD ATM Cap Vol 8Y 2% premium
USCV209 Curncy	USD ATM Cap Vol 9Y 2% premium
USCV2010 Curncy	USD ATM Cap Vol 10Y 2% premium

Tab. B.2: Black Volatilities used to create USD Cap Surface

Bloomberg Code	Description
USCV301 Curncy	USD ATM Cap Vol 1Y 3% premium
USCV302 Curncy	USD ATM Cap Vol 2Y 3% premium
USCV303 Curncy	USD ATM Cap Vol 3Y 3% premium
USCV304 Curncy	USD ATM Cap Vol 4Y 3% premium
USCV305 Curncy	USD ATM Cap Vol 5Y 3% premium
USCV306 Curncy	USD ATM Cap Vol 6Y 3% premium
USCV307 Curncy	USD ATM Cap Vol 7Y 3% premium
USCV308 Curncy	USD ATM Cap Vol 8Y 3% premium
USCV309 Curncy	USD ATM Cap Vol 9Y 3% premium
USCV3010 Curncy	USD ATM Cap Vol 10Y 3% premium
USCV401 Curncy	USD ATM Cap Vol 1Y 4% premium
USCV402 Curncy	USD ATM Cap Vol 2Y 4% premium
USCV403 Curncy	USD ATM Cap Vol 3Y 4% premium
USCV404 Curncy	USD ATM Cap Vol 4Y 4% premium
USCV405 Curncy	USD ATM Cap Vol 5Y 4% premium
USCV406 Curncy	USD ATM Cap Vol 6Y 4% premium
USCV407 Curncy	USD ATM Cap Vol 7Y 4% premium
USCV408 Curncy	USD ATM Cap Vol 8Y 4% premium
USCV409 Curncy	USD ATM Cap Vol 9Y 4% premium
USCV4010 Curncy	USD ATM Cap Vol 10Y 4% premium
USCV501 Curncy	USD ATM Cap Vol 1Y 5% premium
USCV502 Curncy	USD ATM Cap Vol 2Y 5% premium
USCV503 Curncy	USD ATM Cap Vol 3Y 5% premium
USCV504 Curncy	USD ATM Cap Vol 4Y 5% premium
USCV505 Curncy	USD ATM Cap Vol 5Y 5% premium
USCV506 Curncy	USD ATM Cap Vol 6Y 5% premium
USCV507 Curncy	USD ATM Cap Vol 7Y 5% premium
USCV508 Curncy	USD ATM Cap Vol 8Y 5% premium
USCV509 Curncy	USD ATM Cap Vol 9Y 5% premium
USCV5010 Curncy	USD ATM Cap Vol 10Y 5% premium

Tab. B.3: Black Volatilities used to create USD Cap Surface cont.

B.3 EUR Cap Surface

Bloomberg Code	Description
EUCV1 Curncy	EUR ATM Cap Vol 1Y
EUCV2 Curncy	EUR ATM Cap Vol 2Y
EUCV3 Curncy	EUR ATM Cap Vol 3Y
EUCV4 Curncy	EUR ATM Cap Vol 4Y
EUCV5 Curncy	EUR ATM Cap Vol 5Y
EUCV6 Curncy	EUR ATM Cap Vol 6Y
EUCV7 Curncy	EUR ATM Cap Vol 7Y
EUCV8 Curncy	EUR ATM Cap Vol 8Y
EUCV9 Curncy	EUR ATM Cap Vol 9Y
EUCV10 Curncy	EUR ATM Cap Vol 10Y
EUCV101 Curncy	EUR ATM Cap Vol 1Y 1% premium
EUCV102 Curncy	EUR ATM Cap Vol 2Y 1% premium
EUCV103 Curncy	EUR ATM Cap Vol 3Y 1% premium
EUCV104 Curncy	EUR ATM Cap Vol 4Y 1% premium
EUCV105 Curncy	EUR ATM Cap Vol 5Y 1% premium
EUCV106 Curncy	EUR ATM Cap Vol 6Y 1% premium
EUCV107 Curncy	EUR ATM Cap Vol 7Y 1% premium
EUCV108 Curncy	EUR ATM Cap Vol 8Y 1% premium
EUCV109 Curncy	EUR ATM Cap Vol 9Y 1% premium
EUCV1010 Curncy	EUR ATM Cap Vol 10Y 1% premium
EUCV201 Curncy	EUR ATM Cap Vol 1Y 2% premium
EUCV202 Curncy	EUR ATM Cap Vol 2Y 2% premium
EUCV203 Curncy	EUR ATM Cap Vol 3Y 2% premium
EUCV204 Curncy	EUR ATM Cap Vol 4Y 2% premium
EUCV205 Curncy	EUR ATM Cap Vol 5Y 2% premium
EUCV206 Curncy	EUR ATM Cap Vol 6Y 2% premium
EUCV207 Curncy	EUR ATM Cap Vol 7Y 2% premium
EUCV208 Curncy	EUR ATM Cap Vol 8Y 2% premium
EUCV209 Curncy	EUR ATM Cap Vol 9Y 2% premium
EUCV2010 Curncy	EUR ATM Cap Vol 10Y 2% premium

Tab. B.4: Black Volatilities used to create EUR Cap Surface

Bloomberg Code	Description
EUCV301 Curncy	EUR ATM Cap Vol 1Y 3% premium
EUCV302 Curncy	EUR ATM Cap Vol 2Y 3% premium
EUCV303 Curncy	EUR ATM Cap Vol 3Y 3% premium
EUCV304 Curncy	EUR ATM Cap Vol 4Y 3% premium
EUCV305 Curncy	EUR ATM Cap Vol 5Y 3% premium
EUCV306 Curncy	EUR ATM Cap Vol 6Y 3% premium
EUCV307 Curncy	EUR ATM Cap Vol 7Y 3% premium
EUCV308 Curncy	EUR ATM Cap Vol 8Y 3% premium
EUCV309 Curncy	EUR ATM Cap Vol 9Y 3% premium
EUCV3010 Curncy	EUR ATM Cap Vol 10Y 3% premium
EUCV401 Curncy	EUR ATM Cap Vol 1Y 4% premium
EUCV402 Curncy	EUR ATM Cap Vol 2Y 4% premium
EUCV403 Curncy	EUR ATM Cap Vol 3Y 4% premium
EUCV404 Curncy	EUR ATM Cap Vol 4Y 4% premium
EUCV405 Curncy	EUR ATM Cap Vol 5Y 4% premium
EUCV406 Curncy	EUR ATM Cap Vol 6Y 4% premium
EUCV407 Curncy	EUR ATM Cap Vol 7Y 4% premium
EUCV408 Curncy	EUR ATM Cap Vol 8Y 4% premium
EUCV409 Curncy	EUR ATM Cap Vol 9Y 4% premium
EUCV4010 Curncy	EUR ATM Cap Vol 10Y 4% premium
EUCV501 Curncy	EUR ATM Cap Vol 1Y 5% premium
EUCV502 Curncy	EUR ATM Cap Vol 2Y 5% premium
EUCV503 Curncy	EUR ATM Cap Vol 3Y 5% premium
EUCV504 Curncy	EUR ATM Cap Vol 4Y 5% premium
EUCV505 Curncy	EUR ATM Cap Vol 5Y 5% premium
EUCV506 Curncy	EUR ATM Cap Vol 6Y 5% premium
EUCV507 Curncy	EUR ATM Cap Vol 7Y 5% premium
EUCV508 Curncy	EUR ATM Cap Vol 8Y 5% premium
EUCV509 Curncy	EUR ATM Cap Vol 9Y 5% premium
EUCV5010 Curncy	EUR ATM Cap Vol 10Y 5% premium

Tab. B.5: Black Volatilities used to create EUR Cap Surface cont.

B.4 EUR/USD Option Surface

Bloomberg Code	Description
EURUSDV1M Curncy	EUR/USD 1 Month ATM Option Volatility
EURUSDV2M Curncy	EUR/USD 2 Month ATM Option Volatility
EURUSDV3M Curncy	EUR/USD 3 Month ATM Option Volatility
EURUSDV4M Curncy	EUR/USD 4 Month ATM Option Volatility
EURUSDV5M Curncy	EUR/USD 5 Month ATM Option Volatility
EURUSDV6M Curncy	EUR/USD 6 Month ATM Option Volatility
EURUSDV7M Curncy	EUR/USD 7 Month ATM Option Volatility
EURUSDV8M Curncy	EUR/USD 8 Month ATM Option Volatility
EURUSDV9M Curncy	EUR/USD 9 Month ATM Option Volatility
EURUSDV10M Curncy	EUR/USD 10 Month ATM Option Volatility
EURUSDV11M Curncy	EUR/USD 11 Month ATM Option Volatility
EURUSDV1Y Curncy	EUR/USD 1 Year ATM Option Volatility
EURUSDV2Y Curncy	EUR/USD 2 Year ATM Option Volatility
EURUSDV3Y Curncy	EUR/USD 3 Year ATM Option Volatility
EURUSDV4Y Curncy	EUR/USD 4 Year ATM Option Volatility
EURUSDV5Y Curncy	EUR/USD 5 Year ATM Option Volatility
EURUSDV6Y Curncy	EUR/USD 6 Year ATM Option Volatility
EURUSDV7Y Curncy	EUR/USD 7 Year ATM Option Volatility
EURUSDV8Y Curncy	EUR/USD 8 Year ATM Option Volatility
EURUSDV9Y Curncy	EUR/USD 9 Year ATM Option Volatility
EURUSDV10Y Curncy	EUR/USD 10 Year ATM Option Volatility

Tab. B.6: Black Volatilities used to create EUR/USD Option Surface