Master's thesis
Calibration of FX options and pricing of barrier options

Anders Persson

June 4, 2013

Abstract

This paper examines the calibration of foreign exchange options during one year using the Black-Scholes, Heston and Bates model. First the translation from market specific quotes to volatility/strike pairs is considered. A calibration procedure using weighted least squares is then considered, using both unit and capped relative price weights. A comparison is made between the different models and the two weighting schemes; this comparison is made by studying relative error plots.

The market implied volatility surface is constructed using interpolation with a second order polynomial in delta, and flat forward interpolation in maturity dimension. The implied volatility surfaces are then constructed for the stochastic volatility models and the difference in implied volatilities is studied.

The pricing of up-and-out and up-and-in barrier options is performed. The pricing of European style contracts is carried out using a replicating strategy. The pricing of path-dependent contracts is carried out using Monte Carlo techniques where a Brownian bridge is constructed between discrete time points.

Acknowledgements

I would like to thank my supervisor Magnus Wiktorsson for his guidance, support and interesting discussions. I would also like to thank Stefan Hägnesten for introducing me to the FX market, contributing with market data and giving me valuable advice.

Contents

1 Introduction

The use of options, i.e. the right to buy or sell in the future has a long history. It has been argued that already the Greek philosopher Thales used options to speculate on his prediction that a year would yield a bumper crop of olives. Little did he know that over two millenia later (in 1973) a paper called 'The Pricing of Options and Corporate Liabilities' would be published by Fischer Black and Myron Scholes. In this paper the authors derived a partial differential equation which governs the price of a derivative. There are several assumptions made in the model, the most important one that the underlying asset follows a geometric Brownian motion with constant volatility. The assumption about constant volatility has - espicially in the aftermath of the market crash in 1987 - proved to be incorrect. In a Black-Scholes world one would expect the *implied volatility* to be constant for different strikes. Instead a so-called 'smile', or 'skew' can often be observed in markets.

Several models have been proposed to accomodate this. One quite natural extension is to assume volatility to change stochastically. In 1993 Heston proposed a model where the volatility changes according to a square-root diffusion process. An extension of this model was proposed by Bates in 1996 where a compound jump process is added, accomodating the fact that some market show not only continuous change, but sometimes quick and discontinuous changes.

It is clear that a model choice should be made with great care, and chosen according to which underlying asset is modelled. In this thesis options on currencies will be considered. There is a subtle difference between stock options and currency options in that it is not immediately obvious what the underlying asset is. When modelling a stock it is clear that the stock itself is the underlying asset. When it comes to currencies however, we model the exchange rate - and this is clearly not the underlying asset. Instead an investor sees the foreign bond (the foreign bank account) as the risky asset, and the exchange rate can be seen as a stochastic translation between the two currencies.

2 Problem formulation

The foreign exchange option market differs quite a bit compared to the equity or money market. For instance, the FX smile is not given directly, but is instead implicitly given as a set of restrictions implied by market instruments (Wystrup (2010)). In this thesis market instrument data on the Euro Dollar market will be analyzed. The first questions the author poses is the following:

- How does one translate FX instrument data into implied volatilities for options?

It is known that the assumptions made in a Black-Scholes world are not sufficient to explain market prices. It is therefore of interest to see how well we can explain market prices using a stochastic volatility model. Interesting questions to be answered are the following:

- How well do the Black-Scholes, Heston and Bates models explain market prices in and out of sample? More specifically, how well do these models explain prices when studying prices used in the calibration and studying prices for intermediate maturities not used in the calibration? Also, more technically: What are the main qualitative differences between using unit and relative price weights in a least-squares calibration procedure?
- How does one (for a given day) construct the market implied volatility surface in maturity/strike space? Furthermore, how can the implied volatility surface be constructed for our stochastic volatility models, and what can this tell us?

Although a calibration seems to capture the prices of some simple contracts well (even out of sample) one cannot be sure that the same model captures prices for more advanced contracts. Therefore the final question to be posed is:

- When pricing European and path-dependent barrier options, what are the differences between the different models and weighting schemes?

3 The foreign exchange market

3.1 Market overview

Foreign exchange (FX) is the buying of one currency and the selling of another. It is global and over-the-counter, meaning that there is no exchange the deal has to go through. The market is open 24 hours a day, is unregulated and very liquid. Investors will talk to several dealers for the best exchange rate on a specific currency pair. This can be compared to an exchange market, where investors will find out the best deal available to them.

In Kleist et al (2010) one can find out that in 2010 daily average turnover was (approximately) staggering 4 trillion dollars!

The American dollar (USD) plays a huge role in the foreign exchange market, and approximately 80% of all foreign exchange transactions have a dollar leg (Shamah (2003)). The most traded currency pairs are:

- EURUSD: the European euro against the American dollar,

- USDJPY: the American dollar aginst the Japanese yen,
- GBPUSD: the British pound against the American dollar, and
- USDCHF: the American dollar against the Swiss franc.

EURUSD is the most traded currency pair in the world. For this reason it seems like an important market to study, and this is the reason why this report is concerned with this currency pair.

3.2 FX contracts

A *spot transaction* is an agreement between two parties to exchange currencies. The delivery time is usually two business days (except for the Canadian dollar for instance which is one day).

A *forward transaction* is any transaction that settles on a date beyond spot. Two parties agreeing on a forward contract will exchange currencies in the future at a predetermined exchange rate.

An *option* gives the owner the right but not the obligation to buy or sell a specified quantity of a currency at a specified rate on a specified date (Shamah (2003)). The price of a contract is the price of one unit of foreign currency. An example from Wystrup (2010) is the following (some terms are explained in the next section):

3.3 FX terminology

The *exchange rate* (or spot), denoted by S_t represents the value in units of domestic (DOM) currency of one unit of foreign (FOR) currency. FORDOM is called a *currency pair*, and an example would be the following: if the spot rate for USDJPY is 98.22, this means that one dollar equals 98.22 Japanese yen (actual exchange rate April 28 2013).

The strike (denoted K) is the predetermined exchange rate for an option.

The premium currency is the currency in which the price is quoted and could be either domestic or foreign (in our case domestic).

The notional is the amount of currency which the holder of an option may exchange.

3.4 Example of a call option

Here is an example of a European call option:

An investor buys a EURUSD call with a spot of $S_t = 1.3900$, a strike of $K = 1.3500$. The premium of the option is 0.1024 USD. A notional of $1,000,000$ EUR is specified, meaning that the holder of the option will (if the exchange rate is above the strike) receive $1,000,000$ EUR and pay $1,350,000$ USD at maturity, and the current price of the option is $102,400$ USD.

4 Empirical facts

In this section we will study historical data on the exchange rate EURUSD, namely data between January 3 2005 and February 28 2013. In figure 1 the exchange rate and daily log returns during this period are shown.

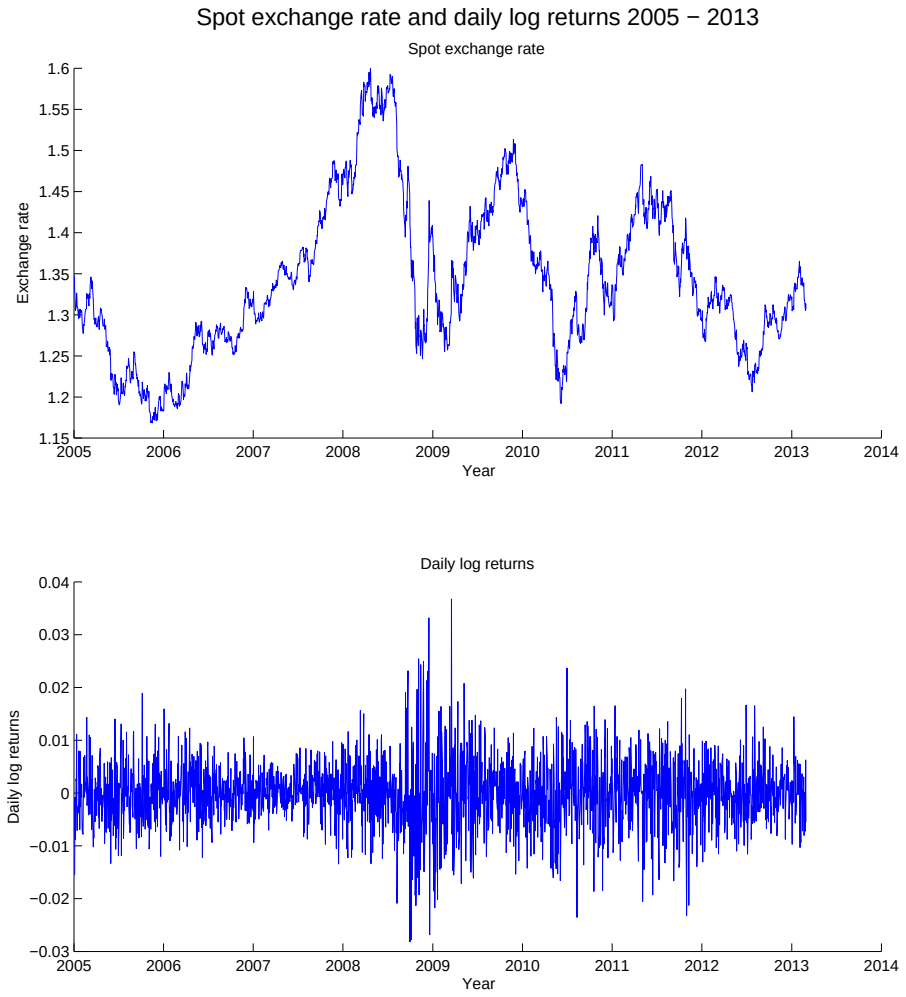

Figure 1: Exchange rate EURUSD and daily log returns, January 3 2005 - February 28 2013

Note the high daily log returns, in absolute terms, in the aftermath of the end of 2008, signifying a volatile market. In figure 2 a histogram is shown for the distribution of the daily log returns. The red curve is a normal probability density fitted to the data. A normal probability plot is also shown.

Note that the mass is more concentrated to the middle, and that the tails are fatter than indicated by the normal assumption. This can also be noted from the normal probability plot. This, together with the fact that volatility changes over time seems to run counter to crucial assumptions in the Black-Scholes model.

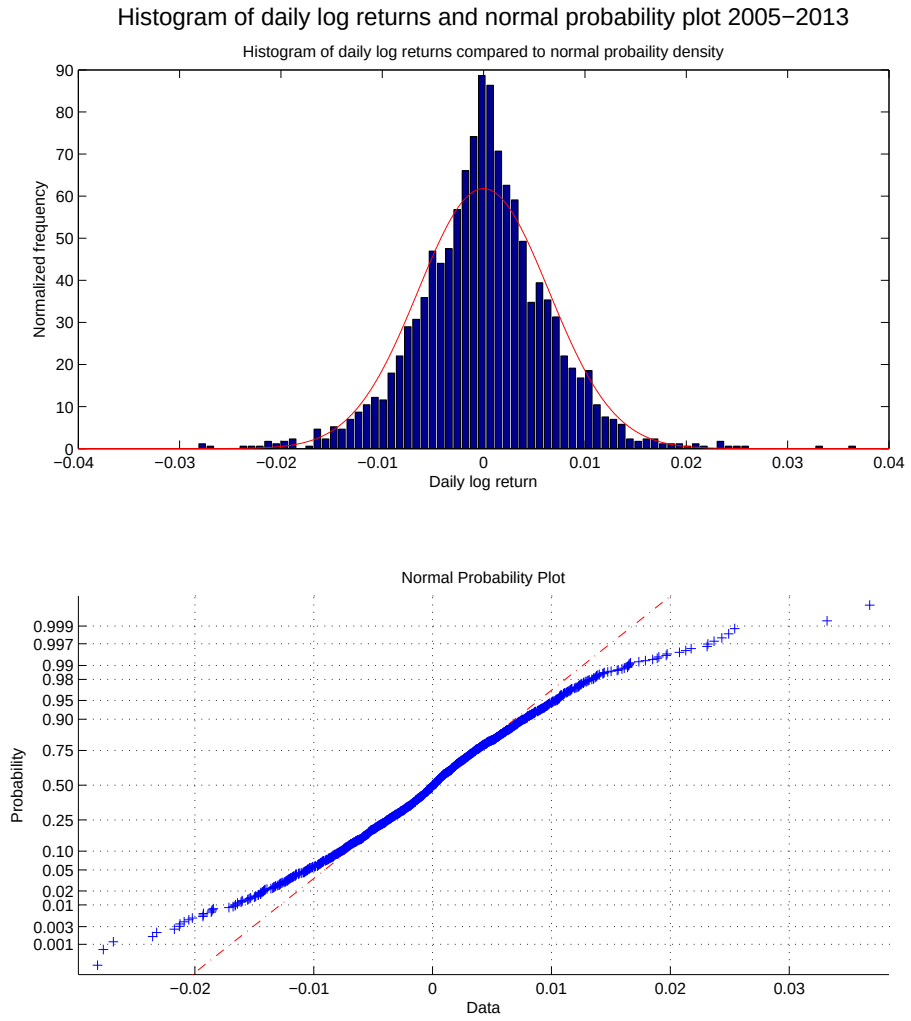

Figure 2: Upper plot: Histogram showing the distribution of the daily log returns together with a fitted normal probability density function. Lower plot: A normal probability plot.

5 FX derivative pricing in a Black-Scholes world

5.1 Exchange rate under risk neutral measure $\mathbb{Q}$

The price of European call and put options were derived by Black and Scholes 1973. In 1983 Garman and Kohlhagen applied this to currency options (Clark (2012)).

We will take the same approach as in Björk (2009) chapter 17, using similar notation. We denote by S_t the exchange rate, i.e. the amount of domestic currency you pay to obtain one unit of the foreign currency. We assume that the exchange rate follows a geometric Brownian motion and that we have one domestic and one foreign riskfree asset with constant interest rates:

$$dS_t = S_t \alpha dt + S_t \sigma d\bar{W}, \tag{5.1}$$

$$dB_d = r_d B_d dt, \tag{5.2}$$

$$dB_f = r_f B_f dt. \tag{5.3}$$

$\bar{W}_t$ here denotes a Wiener process under $\mathbb{P}$.

We are now interested to find out how we can price a contract of the form:

$$Z = \Phi(S_T), \tag{5.4}$$

where Φ is some function. We know from the risk neutral valution formula that

$$\Pi(t; Z) = e^{-r_d(T-t)} \mathbb{E}^{\mathbb{Q}}[\Phi(S_T)], \tag{5.5}$$

where $\mathbb{Q}$ denotes the risk netral measure. We now want to find out what this measure looks like. To this end we proceed in the following way: The possibility to invest in the foreign currency is equivalent to investing in a domestic asset with price process $\tilde{B_f}$:

$$\tilde{B_f} = B_f S_t. \tag{5.6}$$

Using Ito we obtain:

$$d\tilde{B_f} = B_f dS_t + S_t dB_f = \cdots = \tilde{B_f}(\alpha + r_f)dt + \tilde{B_f}\sigma d\bar{W}. \tag{5.7}$$

The model now consists of two assets: B_d and $\tilde{B_f}$. We know that under $\mathbb{Q}$ every domestic asset has r_d as the local rate of return. Hence the dynamics for $\tilde{B_f}$ under $\mathbb{Q}$ is given by:

$$d\tilde{B_f} = r_d \tilde{B_f} dt + \tilde{B_f}\sigma dW, \tag{5.8}$$

where W is a $\mathbb{Q}$-Wiener process. By using Ito on $S_t = \tilde{B_f}/B_f$ we obtain:

$$dS_t = \frac{1}{B_f} d\tilde{B_f} - \frac{\tilde{B_f}}{B_f^2} dB_f = \cdots = (r_d - r_f)S_t dt + S_t \sigma dW. \tag{5.9}$$

The solution can be written as (see Björk for details):

$$S_t = S_0 e^{(r_d - r_f - \frac{1}{2}\sigma^2)t + \sigma W_t}. \tag{5.10}$$

This is a geometric brownian motion with expected value $S_0 e^{(r_d - r_f)t}$.

In summary we now know how to price a contract of the form (5.4): Calculate (5.5), where $\mathbb{Q}$ follows (5.9).

An interesting note to make here is that the above formula is equivalent to pricing a dividend paying stock. If one interchanges r_f with q (the continuous-paid dividend) the formulas are equivalent.

5.2 FX outright forward rate $f(t,T)$

This contract has a value of zero at time t, and at time T there is an exchange between two parties at a pre-specified outright forward rate $f(t,T)$. More precisely, at time T a foreign notional of N will be exchanged against the amount $Nf(t,T)$ of domestic currency. The outright forward can be calculated according to the so called interest-rates parity

$$f(t,T) = S_t e^{(r_d - r_f)\tau}. \tag{5.11}$$

The interest-rates parity can be obtained using an arbitrage argument (demonstrated in table 1): At time zero we enter into a forward contract. Meanwhile we lend $Ne^{-r_f\tau}$ of the foreign currency (FOR) and deposit the money in a domestic bank account. At time T we have no stochastic terms, and we must have $f(t,T) = S_t e^{(r_d - r_f)\tau}$ to avoid arbitrage.

	t	T
Buy forward	0	$N(S_T - f(t,T))$
Lend FOR	$+NS_t e^{-r_f\tau}$	$-NS_T$
Deposit DOM	$-NS_t e^{-r_f\tau}$	$NS_t e^{(r_d - r_f)\tau}$
Sum	0	0

Table 1: Cash flow demonstrating that the fixed rate must be $f(t,T) = S_t e^{(r_d - r_f)\tau}$. Note that all cash flows are denoted in domestic currency.

5.3 FX forward value

An outright forward contract is, as described above, zero-valued to start with. When the exchange rate changes the value of the forward naturally changes. For a pre-specified exchange rate K the value of this contract can be calculated using the theory we derived above. The price can be calculated as:

$$\begin{aligned} v_f(t,T) = e^{-r_d\tau}\mathbb{E}^Q[S_T - K] = e^{-r_d\tau}\mathbb{E}^Q[S_T] - e^{-r_d\tau}K = \\ S_t e^{-r_f\tau} - Ke^{-r_d\tau}. \end{aligned} \tag{5.12}$$

This result will be helpful when we want to price a put option using put-call parity.

5.4 FX vanilla options

The value of a European call option will now be derived. After this the price of a put option will be derived using put-call parity.

According to the risk-neutral valuation formula we want to calculate

$$\begin{aligned} \Pi^C(t,S_t) =& e^{-r_d\tau}\mathbb{E}^Q[(S_T - K)^+] = \\ =& e^{-r_d\tau}\left(\mathbb{E}^Q[S_T \mathbb{1}_{\{S_T > K\}}] - K\mathbb{E}^Q[\mathbb{1}_{\{S_T > K\}}]\right). \end{aligned} \tag{5.13}$$

We note that in both expectations we have the event $S_T > K$:

$$
\begin{aligned}
S_T > K \Leftrightarrow & S_t e^{(r_d - r_f - \frac{1}{2}\sigma^2)\tau + \sigma\sqrt{\tau}G} > K \Leftrightarrow \\
& \Leftrightarrow G > -\left(\frac{\ln \frac{S_t}{K} + (r_d - r_f - \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}} \right) = -\delta,
\end{aligned} \tag{5.14}
$$

where $G \in N(0,1)$.

The second expectation can be calculated:

$$
\mathbb{E}^Q[\mathbb{1}_{\{S_T > K\}}] = Q(S_T > K) = \int_{x > -\delta} e^{-x^2/2} \frac{1}{\sqrt{2\pi}} dx = N(\delta), \tag{5.15}
$$

where N denotes the normal cumulative distribution function.

The first expectation in (5.13) can be calculated as:

$$
\begin{aligned}
\mathbb{E}^Q[S_T \mathbb{1}_{\{S_T > K\}}] = & \int_{x > -\delta} S_t e^{(r_d - r_f - \frac{1}{2}\sigma^2)\tau + \sigma\sqrt{\tau}x} e^{-x^2/2} \frac{1}{\sqrt{2\pi}} dx = \\
= & S_t e^{(r_d - r_f)\tau} \int_{x > -\delta} e^{-\frac{1}{2}(x - \sigma\sqrt{\tau})^2} \frac{1}{\sqrt{2\pi}} dx = \\
= & S_t e^{(r_d - r_f)\tau} \int_{y > -(\delta + \sigma\sqrt{\tau})} e^{-y^2/2} \frac{1}{\sqrt{2\pi}} dy = \\
= & S_t e^{(r_d - r_f)\tau} N(\delta + \sigma\sqrt{\tau}).
\end{aligned} \tag{5.16}
$$

Using the same notation as in Wystrup (2010) one finally obtains:

$$
\Pi^C(t, S_t) = e^{-r_f \tau} S_t N(d_+) - e^{-r_d \tau} K N(d_-), \tag{5.17}
$$

where

$$
d_\pm = \frac{\ln\left(\frac{f(t,T)}{K}\right) \pm \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}. \tag{5.18}
$$

The price of a put option could be obtained using a similar approach, but an easier approach is to use put-call parity - the fact that a call minus a put equals a forward with a given strike K. The put equals a call minus the forward, and we obtain:

$$
\begin{aligned}
\Pi^P(t, S_t) = & \Pi^C(t, S_t) - v_f = \\
= & e^{-r_f \tau} S_t N(d_+) - e^{-r_d \tau} K N(d_-) - (S_t e^{-r_f \tau} - K e^{-r_d \tau}) = \\
= & - S_t e^{-r_f \tau}(1 - N(d_+)) + e^{-r_d \tau} K (1 - N(d_-)) = \\
= & - S_t e^{-r_f \tau} N(-d_+) + e^{-r_d \tau} K N(-d_-).
\end{aligned} \tag{5.19}
$$

5.5 Implied volatility

The *vega* of an option measures the sensitivity in the price when the volatility changes (note that the volatility is a parameter and not a variable). From Wystrup (2006) page 19 the vega for a call option is calculated as: $\frac{\partial \Pi}{\partial \sigma} = S_t e^{-r_f \tau} \sqrt{\tau} N'(d_+)$. We see that this expression is always positive, meaning that

the price of a European call option is an increasing function of σ. For every price we can find out which volatility *implies* this price. From put-call parity it follows that the implied volatility is the same for a put option with the same strike and time to maturity. The implied volatility can be seen as the market's view on how volatile the market is going to be in the future - not be confused with the historic volatility, which measures the volatility in the past.

In a Black-Scholes world, since the volatility is assumed to be constant, the volatility should not be dependent on the strike of the option. However, on most markets one observes a *smile* or a *skew*, i.e. the implied volatility is dependent on the strike level. This shows that the assumption of constant volatility is incompatible with market data. A remedy for this is to assume that volatility changes stochastically instead.

6 Stochastic volatility models

An extension to the assumption of constant volatility is to allow time dependence of volatility of the form $\sigma = \sigma(t)$. When taking the term structure into account one still can't account for the fact that different strikes give different implied volatilities. Dupire (1993) proposed a local volatility model, where volatility is both time and state dependent. He showed that it was possible to find $\sigma = \sigma(S_t, t)$ that accounts for the dynamics of the whole volatility surface (Clark (2012)).

A perhaps more realistic assumption is that volatility is random in its behaviour. The two stochastic volatility models considered in this thesis are the Heston and Bates models.

6.1 The Heston model

In 1993 Heston proposed a model where the volatility itself follows a random process, a so called square-root process. Under the risk-neutral measure $\mathbb{Q}$ the model takes the following form:

$$\begin{aligned} dS_t &= (r_d - r_f)S_t dt + \sqrt{V_t} S_t dW_t^1, \\ dV_t &= \kappa(\theta - V_t)dt + \sigma\sqrt{V_t} dW_t^2, \\ dW_t^1 dW_t^2 &= \rho dt, \end{aligned} \tag{6.1}$$

where $\{S_t\}_{t\geq 0}$ and $\{V_t\}_{t\geq 0}$ are the exchange rate and volatility processes respectively. $\{W_t^1\}_{t\geq 0}$ and $\{W_t^2\}_{t\geq 0}$ are two Brownian motions with correlation parameter ρ. It is instructive to analyze the parameters and their impact on the behaviour of the process, and also important so that one can notice when parameter values seem to be out of line.

ρ measures the correlation between the two processes. If the correlation is negative this means that volatility tends to increase when returns are negative. This implies that the probability mass on the left side is more scattered than on the right side of the distribution. ρ also has an impact on the volatility skew. If ρ is negative, the volatility smile will have a skew with higher volatility for lower strikes than for higher strikes (see Moodley (2005)).

σ measures the volatility of volatility. If σ is zero, the volatility is deterministic, and increasing this parameter will increase the kurtosis on both sides of the distribution.

θ is the long-run mean of the variance process. Note that when $V_t > \theta$ the drift becomes negative, and when $V_t < \theta$ the drift becomes positive.

κ measures the rate of reversion, i.e. the speed at which the process goes back to θ again. One could make an economic argument here also that κ in a way measures volatility clustering; if the parameter is low volatility tends to stay the same for some time.

Except from these four parameters one also needs to estimate the initial variance, here denoted V_0. This is not a parameter, but a process. In a market where the exchange rate follows the Heston model perfectly, we would expect the parameters to be constant, and V_0 to change.

In the Black-Scholes model the price of a European call and put option can be derived in closed form. The reason for this is that the distribution of the log-return can be derived is closed form. This is not the case for many stochastic volatility models. Another approach to the pricing-problem is to use inverse Fourier methods. The rationale behind it is to derive a pricing formula where one inverts the characteristic function of the log-spot $\ln S_t$. The (at least to the author) astounding fact is that the characteristic function can indeed be derived in closed form for some models - the Heston model is one of them. Using fourier methods the closed form solution of a European call option is:

$$\Pi^C(t, S_t) = S_t P_1 e^{-r_f \tau} - K e^{-r_d \tau} P_2, \tag{6.2}$$

where

$$\begin{aligned}
P_1 &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re\left(\frac{e^{-i\phi \ln K} f(\phi - i)}{i\phi f(-i)} \right) d\phi, \\
P_2 &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re\left(\frac{e^{-i\phi \ln K} f(\phi)}{i\phi} \right) d\phi, \\
f_{\text{Heston}}(\phi) &= e^{A+B+C}, \\
A &= i\phi \ln S_t + i\phi(r_d - r_f)\tau, \\
B &= \frac{\theta\kappa}{\sigma^2} \left((\kappa - \rho\sigma i\phi - d)\tau - 2\ln\left(\frac{1 - ge^{-d\tau}}{1 - g} \right) \right), \\
C &= \frac{\frac{V_0}{\sigma^2}(\kappa - \rho\sigma i\phi - d)(1 - e^{-d\tau})}{1 - ge^{-d\tau}}, \\
d &= \sqrt{(\rho\sigma i\phi - \kappa)^2 + \sigma^2(i\phi + \phi^2)}, \\
g &= \frac{\kappa - \rho\sigma i\phi - d}{\kappa - \rho\sigma i\phi + d},
\end{aligned} \tag{6.3}$$

where $f_{\text{Heston}}(\phi)$ denotes the characteristic function (see Gilli and Schumann (2010)).

6.2 The Bates model

The Bates model is an extension of the Heston model, adding jumps following a compound Poisson process with jump intensity λ. The dynamics under $\mathbb{Q}$ is:

$$
\begin{aligned}
dS_t &= (r_d - r_f - \lambda\mu_J)S_t dt + \sqrt{V_t}S_t dW_t^1 + J_t S_t dN_t, \\
dV_t &= \kappa(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^2, \\
dW_t^1 dW_t^2 &= \rho dt.
\end{aligned} \tag{6.4}
$$

A compound Poisson process is a poisson process where the jump sizes follow a particular distribution, in this case:

$$
\log(1 + J_t) \in N\left(\log(1 + \mu_J) - \frac{\sigma_J^2}{2}, \sigma_J^2\right). \tag{6.5}
$$

The drift has to be modified for arbitrage reasons; for further details the reader is referred to the litterature on stochastic calculus for Poisson processes. When simulating (6.4) on the log scale the dynamics is the same as for Heston, except for the drift and jumps with a magnitude decided by a normal random variable.

The pricing can in this case be carried out according to (6.2). The modification compared to the Heston case is that the characteristic function should be multiplied by a term accounting for the jump (Gilli and Schumann (2010)).

$$
f_{\text{Bates}} = f_{\text{Heston}} e^{-\lambda\mu_J i\phi\tau + \lambda\tau\left((1+\mu_J)^{i\phi} e^{\frac{1}{2}\sigma_J^2 i\phi(i\phi-1)} - 1\right)}. \tag{6.6}
$$

The reason why the characteristic function is formed as the characteristic function for Heston multiplied by a jump-term is the following: In the Bates model the log-return is formed as a sum of two independent random variables; one accounting for the stochastic volatility part and one accounting for the jump-part. The characteristic function for the sum of two independent random variables is the multiplication of the two characteristic functions (see Gut (2009)).

7 Simulation

7.1 Monte Carlo

When pricing path-dependent options it is of interest to be able to simulate from a model using Monte Carlo methods. In this section the framework for the Monte Carlo sampler will be presented, as well as one variance reduction technique.

7.1.1 The basic Monte Carlo sampler

Let g be a real-valued function, and X a random variable with density function f defined under the probability measure $\mathbb{Q}$. A technical assumption is further that $g(X)$ has finite expectation μ and variance σ^2. It is of interest to determine the quantity

$$
\mathbb{E}^Q(g(X)) = \int_{x\in\mathbb{R}} f(x)g(x)dx. \tag{7.1}
$$

Given N independent realizations of X - $x_1, x_2, \cdots, x_N$, the basic Monte Carlo sampler is formed as

$$S_N = \frac{1}{N}\sum_{i=1}^{N} g(x_i). \tag{7.2}$$

According to the law of large numbers, (7.2) $\xrightarrow{p}$ (7.1), $N \to \infty$, where p denotes convergence in probability. The law of large numbers tells us that in the long run we will guess correctly. An even more interesting result for practical purposes is the central limit theorem (see Gut (2009)):

$$S_N \xrightarrow{d} N(\mu, \frac{\sigma^2}{N}) \text{ as N } \to \infty. \tag{7.3}$$

This tells us approximately how the error is distributed for large N. We can now contruct a confidence interval for S_N with confidence level α:

$$\mathbb{P}\left(\mu - \lambda_{\alpha/2}\frac{\sigma}{\sqrt{N}} < S_N < \mu + \lambda_{\alpha/2}\frac{\sigma}{\sqrt{N}}\right) = 1 - \alpha. \tag{7.4}$$

The above expression however presupposes that we know the variance of $g(X)$ - which we normally do not. One can therefore instead say that due to the central limit theorem we know that our sample mean asymptotically is approximately normal with mean μ, and variance $\tilde{\sigma}^2$. The variance is calculated as the unbiased sample variance:

$$\tilde{\sigma}^2 = \frac{1}{N-1}\sum_{i=1}^{N}(g(x_i) - S_N). \tag{7.5}$$

The confindence interval with confidence level α is then: $I_\alpha = [S_N - \lambda_{\alpha/2}\tilde{\sigma}, S_N + \lambda_{\alpha/2}\tilde{\sigma}]$

7.1.2 Control variates

In this section a variance reduction technique called control variates will be introduced.

It is of interest to decide the expectation of a random variable X. We can draw samples from X, and also from another stochastic variable Y which is correlated with X. One can then form the following statistic:

$$\tilde{X} = X - b(Y - \mathbb{E}(Y)). \tag{7.6}$$

The expectation is the same as that of X, and the variance is:

$$\mathbb{V}(\tilde{X}) = \mathbb{V}(X) + b^2\mathbb{V}(Y) - 2b\sqrt{\mathbb{V}(X)\mathbb{V}(Y)}\rho_{XY} \triangleq f(b). \tag{7.7}$$

Note that this is a convex function since $f^{(2)}(b) > 0$. Hence the optimum can be found by setting the derivative equal to zero. Solving this equation one finds that $b^* = \sqrt{\mathbb{V}(X)/\mathbb{V}(Y)}\rho_{XY}$. Inserting this into (7.7) one finds that the variance is:

$$\mathbb{V}(\tilde{X}) = f(b^*) = \mathbb{V}(X)(1 - \rho_{XY}^2). \tag{7.8}$$

An important remark to make here is that in order to get a variance reduction the correlation between the random variables should be significantly different

from zero. It should also be relatively easy to simulate from Y - otherwise it might not be computationally efficient.

Here comes an explanation for how control variates can be used: An up-and-out call option (explained more in detail in chapter 12) gives the same payoff as a European call option given that a specific barrier has not been breached during the lifetime of the option. An up-and-in call option will give the same payoff as a European call option given that a specific barrier *has* been breached during this time period.

When pricing an up-and-in call option a natural control variate to use is the payoff from a European call option, which intuitively should be correlated with the barrier. In this case it is computationally cheap, since the price can be calculated once, and the payoff is easily calculated. The computational gain will increase the closer the barrier is to the strike, since the payoffs will in this case more and more look like the payoffs from a vanilla option.

In the case of an up-and-out barrier, the use of the same control variate is a little more dubious. The reason for this is that when we simulate the exchange rate and we reach the barrier, we immediately want to terminate that simulation since we know the final payoff will be zero. In this case it is hence (for barriers relatively close to the strike) computationally expensive to simulate the control variate.

7.2 Brownian bridge

When pricing for example an up-and-out call option a problem arises from the fact that one only can obtain a finite number of points in the sample - namely that the barrier might knock out the option between two discrete time points. A remedy for this is to use a Brownian bridge construction. The aim in this section is to derive the distribution of the maximum of a Brownian motion, given its endpoints. In this way we are (by simulating on a log-scale) able to calculate the probability that a geometric Brownian motion crosses the barrier in between two time points. In the Heston and Bates cases this will not be exact, since volatility changes. However, if the grid points are relatively close we can 'freeze' the volatility during this short period of time.

Lemma 7.2.1. *Let X and Y be two multivariate normal variables. Then $Y \mid X$ is normal, and the conditional expectation and variance are given by*

$$\begin{aligned} \mathbb{E}[Y \mid X] &= \mathbb{E}[Y] + \frac{\mathbb{C}[Y,X]}{\mathbb{V}[X]}(X - \mathbb{E}[X]) \\ \mathbb{V}[Y \mid X] &= \mathbb{V}[Y] - \frac{\mathbb{C}[Y,X]^2}{\mathbb{V}[X]} \end{aligned} \tag{7.9}$$

Proof See Gut (2009). □

Lemma 7.2.2. *Let W_t be a standard Brownian motion. Furthermore, let $X_t = X_0 + \mu t + W_t$, $0 < t \leq u$. Then it holds that*

$$X_t \mid X_u \in N\left(X_0 + \frac{t}{u}(X_u - X_0), \frac{t(u-t)}{u}\right). \tag{7.10}$$

Proof The distribution of $X_t \mid X_u$ is normal, since X_t and X_u are multivariate normal. The expectation and variance follows from Lemma 7.2.1. □

The interesting conclusion here is that the conditional distribution is independent of the drift. This means that it suffices to calculate the distribution of the maximum of a Brownian motion *without* drift.

Theorem 7.2.1. *Let W_t be a standard Brownian motion. Then it holds that*

$$\mathbb{P}\left(\sup_{0<t\leq T} W_t > b \mid W_T = x\right) = e^{\frac{-2b(b-x)}{T}}. \tag{7.11}$$

Proof Define $\tau_b = \inf\{t : W_t = b\}$. Clearly, the event that the supremum of W_t crosses b is equivalent to the event that $\tau_b < T$. We get:

$$\mathbb{P}\left(\tau_b \leq T \mid W_T = x\right) = \frac{\mathbb{P}(\tau_b \leq T, W_T \in dx)}{\mathbb{P}(W_T \in dx)}, \tag{7.12}$$

where dx denotes an infinitesimal surrounding of x (the notation is clearly improper, but quite convenient). We note that for $b < x$ we have that $\{\tau_b \leq T\} \subset \{W_T \in dx\}$, and in this case the probability equals one. We now consider the case $b > x$ and calculate the numerator and denominator:

$$\begin{aligned}
\text{Numerator} &= \mathbb{P}(\tau_b \leq T)\mathbb{P}(W_T \in dx \mid \tau_b \leq T) = \\
&= \mathbb{P}(\tau_b \leq T)\mathbb{P}(W_T \in 2b - dx \mid \tau_b \leq T) = \\
&= \mathbb{P}(W_T \in 2b - dx, \tau_b \leq T) = \\
&= \mathbb{P}(W_T \in 2b - dx) = \\
&= \frac{1}{\sqrt{T}}\phi\left(\frac{2b-x}{\sqrt{T}}\right)dx. \\
\text{Denominator} &= \frac{1}{\sqrt{T}}\phi\left(\frac{x}{\sqrt{T}}\right)dx,
\end{aligned} \tag{7.13}$$

where ϕ is the density function for a standard normal random variable. The major trick in this calculation is the second equality, where it is used that conditional that the barrier is breached before time T, the Brownian motion is equally likely to go up or down. Dividing the numerator with the denominator in (7.13) gives (7.11). □

Let us now assume that we have a process X_t, $0 < t \leq T$ of the following form:

$$X_t = X_0 + \mu t + \sigma W_t = X_0 + \sigma\left(\frac{\mu t}{\sigma} + W_t\right) \triangleq X_0 + \sigma B_t. \tag{7.14}$$

We note that B_t starts at zero, has a drift of μ/σ and ends at $(X_T - X_0)/\sigma$. Using the fact that the maximum of a Brownian bridge is independent of the drift, and the result from the previous theorem we now get:

$$
\begin{aligned}
&\mathbb{P}\left(\sup_{0<t\leq T} X_t \geq x \mid X_0, X_T\right) = \\
&= \mathbb{P}\left(\sup_{0<t\leq T} X_0 + \sigma B_t \geq x \mid X_0, X_T\right) = \\
&= \mathbb{P}\left(\sup_{0<t\leq T} B_t \geq \frac{x - X_0}{\sigma} \mid B_T = \frac{X_T - X_0}{\sigma}\right) = \\
&= \exp\left(\frac{-2\tilde{b}(\tilde{b} - \tilde{x})}{t}\right), \\
&\tilde{x} = \frac{x - X_0}{\sigma}, \tilde{b} = \frac{X_T - X_0}{\sigma}.
\end{aligned}
\tag{7.15}
$$

7.3 Simulating Black-Scholes

As earlier derived, under the domestic risk-neutral measure $\mathbb{Q}$, S_t follows a geometric brownian motion:

$$
dS_t = (r_d - r_f)S_t d_t + \sigma S_t dW_t. \tag{7.16}
$$

If $S(t) = S_t$, the spot price at time $T > t$ can be obtained in closed-form:

$$
S(T) = S_t e^{((r_d - r_f - \frac{1}{2}\sigma^2)(T-t) + \sigma(W(T) - W(t)))}. \tag{7.17}
$$

We now define the process $X_t = \log S_t$. The solution for X_T is:

$$
\begin{aligned}
X_T &= X_t + (r_d - r_f - \frac{1}{2}\sigma^2)(T-t) + \sigma(W(T) - W(t)) \stackrel{d}{=} \\
&\stackrel{d}{=} X_t + (r_d - r_f - \frac{1}{2}\sigma^2)\tau + \sigma\sqrt{\tau}G,
\end{aligned}
\tag{7.18}
$$

where $G \in N(0,1)$, and $\stackrel{d}{=}$ denotes equal in distribution. We do not really need to simulate Black-Scholes since closed-form solutions for barrier options are available. However, it is a good way to check that the Brownian bridge construction works correctly.

7.4 Simulating Heston

Several simulation techniques for the Heston (and other stochastic volatility models) have been proposed in the litterature. The easiest of these include the Euler and Milstein scheme. In this thesis a simulation scheme where the variance is simulated exactly (and hence cannot become negative) will be used. An error will occur when approximating an integral.

To start with, let's assume that an exact simulation procedure is given for the variance; more precisely, let's assume we know the distribution of $V(t_{k+1} \mid t_k)$. In (6.1) there is a correlation between the Brownian motions. To simplify one can put: $\rho d\tilde{W}_1 + \sqrt{1-\rho^2}d\tilde{W}_2 := dW_1$ and $d\tilde{W}_1 := dW_2$, where $\tilde{W}_1$ and $\tilde{W}_2$ are two independent Brownian motions. The new pair of random variables are equidistributed with the previous ones (normal with same expected value,

variance and covariance). By rewriting the variance in integral form, and solving for the spot rate one obtains:

$$S(t_{k+1}) = S(t_k)\exp\left((r_d - r_f)\tau - \frac{1}{2}\int_{t_k}^{t_{k+1}} V(s)ds + \int_{t_k}^{t_{k+1}} \sqrt{V(s)}\rho d\tilde{W}_1(s) + \int_{t_k}^{t_{k+1}} \sqrt{V(s)}\sqrt{1-\rho^2}d\tilde{W}_2(s)\right), \tag{7.19}$$

$$V(t_{k+1}) = V(t_k) + \int_{t_k}^{t_{k+1}} \kappa(\theta - V(s))ds + \int_{t_k}^{t_{k+1}} \sigma\sqrt{V(s)}d\tilde{W}_1(s). \tag{7.20}$$

From (7.20) one gets:

$$\int_{t_k}^{t_{k+1}} \sqrt{V_s}d\tilde{W}_1(s) = \frac{1}{\sigma}\left(V(t_{k+1}) - V(t_k) - \int_{t_k}^{t_{k+1}} \kappa(\theta - V(s))ds\right). \tag{7.21}$$

Plugging this into (7.19) and rearranging the terms one gets:

$$S(t_{k+1}) = S(t_k)\exp\left((r_d - r_f - \frac{\rho\kappa\theta}{\sigma})\tau + \frac{\rho}{\sigma}(V(t_{k+1}) - V(t_k))\right) \times \times \exp\left(\left(\frac{\rho\kappa}{\sigma} - \frac{1}{2}\right)\int_{t_k}^{t_{k+1}} V(s)ds + \sqrt{1-\rho^2}\sqrt{\int_{t_k}^{t_{k+1}} V(s)ds}G\right), \tag{7.22}$$

where $G \in N(0,1)$. The last term in (7.22) is obtained from the fact that $\int_{t_k}^{t_{k+1}} \sqrt{V(s)}\sqrt{1-\rho^2}d\tilde{W}_2$ is normally distributed with mean zero and variance $\int_{t_k}^{t_{k+1}} V(s)\sqrt{1-\rho^2}ds$. In practice it is easier to simulate the exchange rate on a log-scale. On the final day in the simulation, the exchange rate is simply calculated by taking the exponential of this.

Now assume that a sequence of variances $V_{t_1}, \cdots, V_{t_k}, V_{t_{k+1}}, \cdots, V_N$ are known. Furthermore the time-grid is equidistant and equals Δt. The integral $\int_{t_k}^{t_{k+1}} V(s)ds$ is very hard to simulate exactly. The exact simulation involves calculating the characteristic function, containing the modified Bessel function of the first kind, and then inverting this to obtain the distribution function G. To draw from this distribution one can draw a uniform number U, and then with a Newton-Raphson root search find the solution to the equation $G-U=0$ (See Haastrecht and Pelsser (2008)). In summary, this is a complicated and (presumably) very time consuming way to generate these numbers.

Instead of this a simplified approach is taken, by approximating the integral by $\frac{V(t_{k+1})+V(t_k)}{2}\Delta t$.

The simulation of the variance process boils down to finding the conditional distribution for $V(t_{k+1}) \mid V(t_k)$. This random variable is distributed as a constant times a non-central chi-square distributed random variable. The interested reader is referred to Andersen et. al (2010) where the details are given.

7.5 Simulating Bates

Recall that in the Bates model there is an extra drift term compared to the Heston model. Also, written on the log scale there is a normal distributed random variable (a jump) added, following a Poisson process with intensity λ. In order to simulate this process we can simply simulate the process without the jumps separately from the jumps - since these are independent. The first part is simulated as in the Heston case, but with the extra drift term. The jumps can be simulated by first simulating the number of jumps, which we know is distributed as $Poi(\lambda\tau)$, where $\tau = T - t$ is the time to maturity. Conditional on the number of jumps, the jump times are uniformly distributed on the interval $[t, T]$ (see Gut (2009)). Schematically this can be summed up as:

$$\begin{aligned} &1.\ \text{Draw } N \in Poi(\lambda\tau), \\ &2.\ \text{Draw } N \text{ uniform numbers } U_1, \cdots, U_N \text{ on } [0,1], \\ &3.\ \text{The time points are formed as: } U_{(1)}\tau, \cdots, U_{(N)}\tau, \end{aligned} \tag{7.23}$$

where $U_{(i)}$ denotes the order statistic, i.e. $U_{(i)}$ is the i:th largest number when sorting the numbers in ascending order.

8 FX conventions and quotations

8.1 Delta

On the FX market there are several ways to define the delta. The delta can be measured as spot or forward delta, and also as premium adjusted or not premium adjusted. These different deltas are merely mentioned here so that the reader has at least heard of them. In this thesis only the spot delta needs to be defined: The spot delta measures the sensitivity of the vanilla option when the spot changes:

$$\Delta_S(K, \sigma, \phi) \triangleq \frac{\partial \Pi}{\partial S} = \Pi_S = \phi e^{-r_f \tau} N(\phi d_+), \tag{8.1}$$

where $\phi = +1$ for a call option and $\phi = -1$ for a put option. The last equality in (8.1) holds in a Black-Scholes world.

For a given delta we want to find the corresponding strike. Solving (8.1) for K gives for a call option:

$$K = f(t,T) e^{-N^{-1}(e^{r_f \tau}\Delta_S)\sigma^2(\Delta_S)\sqrt{\tau} + \frac{1}{2}\sigma^2\tau}. \tag{8.2}$$

Using put-call parity and differentiating with respect to S_t one obtains the put-call delta parity:

$$\Delta_S(K, \sigma, +1) - \Delta(K, \sigma, -1) = e^{-r_f \tau}. \tag{8.3}$$

This result says that the difference between a delta for a call and put option with the same strikes is equal to a discount factor. This is a useful result, since it makes it possible for us to translate a delta for a put option into the delta for a call option.

An important note to make here is that the delta is sometimes casually quoted in absolute terms. A '$25DP$' is read: a put option with a delta of 0.25, when in fact it is implicitly meant -0.25.

8.2 At-the-money (ATM) definitions

At-the-money aims at choosing a strike such that the value of an option is 'in between' in-the-money (ITM) and out-of-the money (OTM) measured suitably. Defining ATM is not as straightforward as one might first think and there are four different ways to define ATM (Wystrup (2010)). A straight forward approach is to set $K = S_t$, i.e. to choose the strike as the current exchange rate. A similar approach is to choose the strike as the forward level, i.e. the expected exchange rate under the risk neutral measure. Yet another definition is to choose the strike such that a call with this strike equals the put value with the same strike. The last definition (which is used for the data in this thesis) is to choose the strike such that the delta of a call option equals the delta of a put option (again in absolute terms; the delta for a put option is always negative). This definition is constructed such that an investor can buy a delta-neutral straddle, i.e. a long call and a long put such that the delta of the contract equals zero. This can be seen as a way to 'buy volatility', since the payoff depends on how volatile the exchange rate is. The different ATM definitions are summarized below:

$$\begin{aligned} \text{ATM-spot} \quad & K = S_t, \\ \text{ATM-fwd} \quad & K = f, \\ \text{ATM-value-neutral} \quad & \text{K such that call value = put value,} \\ \text{ATM-}\Delta\text{-neutral} \quad & \text{K such that call delta = put delta.} \end{aligned} \tag{8.4}$$

8.3 Quotation

On the equity market for instance option prices are quoted for different strikes and maturities. For a given stock, prices are given for a whole range of different strike levels. In this way it is possible for the trader to retrieve the implied volatily smile by simply translating prices into volatilities. The prices of FX options are not quoted as directly as one might think. Instead prices are quoted in terms of deltas and implied volatilities. The volatilities are not quoted directly for vanilla options, but instead from specific contracts in order to capture both the convexity and skew of the volatility smile. This is indeed very confusing, but it is market practice to do so. For a given maturity volatilities for FX options can be quoted in four ways: ATM, risk reversals, butterflies and strangles. The quotes used in this thesis are ATM, risk reversals and butterlies. These contracts are constructed as a combination of vanilla options with a specific delta (0.25 for example).

8.3.1 ATM options

The most liquid options traded are the ATM options. The strike for these options is set to the ATM level - note however that one must know which ATM convention is used. Since in our case we have the ATM-Δ-neutral definition, the

corresponding delta is $e^{-r_f\tau}/2$. The corresponding strike can then be retrieved from (8.2).

8.3.2 Risk reversals

A risk reversal is a contract consisting of a long position of an out-of-the money (OTM) call option and a short position of an OTM put option where the options have the same deltas, typically 0.10 or 0.25 (note again the sign convention for put options). The risk reversal is used as a measure of the skew in the volatility smile. The risk reversal is quoted as the difference between the implied volatility of a call and the implied volatility of a put with the same delta in absolute terms. For a generic delta Δ we have:

$$\sigma_{\Delta RR} = \sigma_{\Delta DC} - \sigma_{\Delta DP}. \tag{8.5}$$

If, for example the $\sigma_{25RR} < 0$, this means that OTM call options are cheaper than OTM put options.

8.3.3 Butterflies

A butterfly is a symmetric product giving maximum payoff if the spot rate is close to the at-the-money (ATM) level at expiry. The product can be replicated using an OTM call with a given delta, long an ITM call with the strike given as the same as a put with the same given delta and short two ATM call options. The butterfly measures the convexity of the volatility smile. The butterfly is quoted as:

$$\sigma_{\Delta BF} = \frac{\sigma_{\Delta DC} + \sigma_{\Delta DP}}{2} - \sigma_{ATM}. \tag{8.6}$$

A graphic representation of how the risk reversal measures the skew and the butterfly measures the degree of convexity can be seen in figure 3 - which was taken from Wystrup (2006).

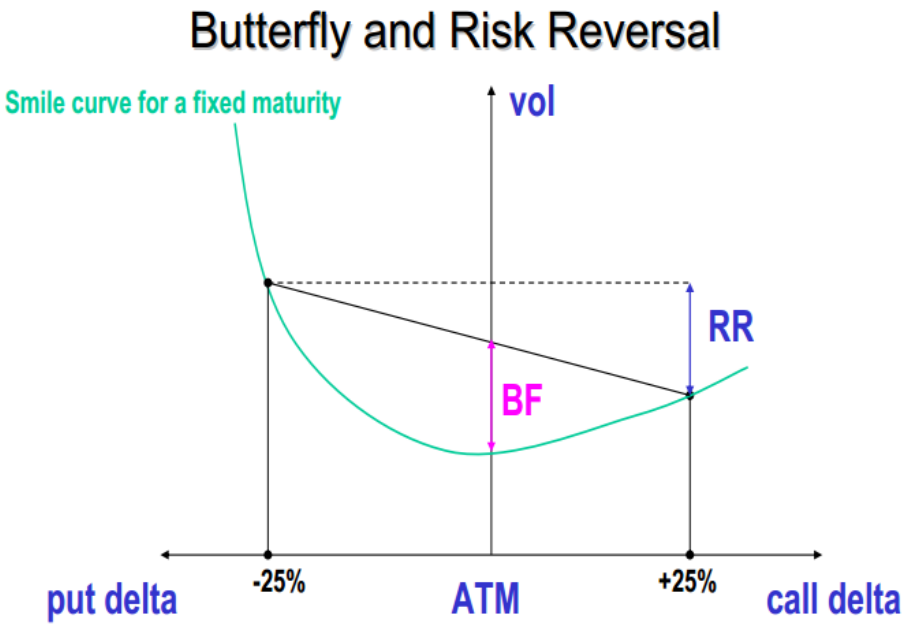

Figure 3: Graphic represention of the butterfly and risk reversal. Note that the more negative the risk reversal becomes, the more prominent skew. Also, the bigger butterfly, the more curvature in the smile (or mathematically: the higher second derivative).

9 Data set

The option data was obtained from Nordea in Copenhagen, and consists of quotes on ATM, risk reversals and butterflies for the EURUSD exchange pair. These are quoted with deltas 0.10 and 0.25. Hence we have quotes for σ_{ATM}, σ_{25RR}, σ_{10RR}, σ_{25BF} and σ_{10BF}. We now want to find out which implied volatilities this corresponds to for call and put options. Combining (8.5) and (8.6) we can for a generic Δ solve for $\sigma_{\Delta DC}$ and $\sigma_{\Delta DP}$:

$$\begin{aligned} \sigma_{\Delta DC} &= \sigma_{\Delta BF} + \frac{1}{2}\sigma_{\Delta RR} + \sigma_{ATM}, \\ \sigma_{\Delta DP} &= \sigma_{\Delta BF} - \frac{1}{2}\sigma_{\Delta RR} + \sigma_{ATM}. \end{aligned} \tag{9.1}$$

Using put-call-delta parity one can then obtain five delta/volatility pairs for vanilla call options. Using (8.2) the delta and volatility can be plugged in to obtain the strikes. The prices of the five options are then calculated by the Black-Scholes formula for call options.

The data consists of 263 trading days ranging from February 28 2012 to February 28 2013. For each day quotes are given for 1M (M denotes months), 2M, 3M, 6M, 9M and 12M maturities.

Normally one would have both a bid and an ask price - for simplicity only midvolatilities (i.e. the average of the bid and the ask implied volatilities) have

been used in this study. It seems to be market practice to use midvolatilities, but it quite frankly seems strange, since there is no linear relation between volatility and price.

Interest rates have been obtained from the British Bankers' Association via the federal reserve bank of St. Louis
(http://research.stlouisfed.org/fred2/categories/33003/downloaddata),
and consists of USD Libor and Euribor interest rates with different tenors. These rates are quoted as annualized rates (quoted as r_a). This means that the interest obtained for a period of n days is $r_a \frac{n}{360}$. In our case the $n/360$ will simply become $1/12, 2/12$ et cetera. This can then be translated into effective annual interest rate, and then into continuously compounded interest rates r_c. The conversion is:

$$r_c = \ln\left(1 + r_a \frac{n}{360}\right) \frac{360}{n}. \tag{9.2}$$

In order to check the interpolation capacity within the different models we split our data into one estimation and one validation set. The estimation set consists of the data with maturities 1M, 3M, 6M and 12M; the validation set consists of the data with maturities 2M and 9M.

10 Calibration

10.1 The calibration procedure

A calibration procedure aims at, given a model and a parameter set, finding parameter values such that the prices implied by the model do not differ too much from the prices given by the market. Daily data is used in the calibration and a weighted least square scheme is chosen. Mathematically put, given a model and a parameter set $\hat{\theta}$, we choose $\hat{\theta}$ as

$$\hat{\theta} = \arg\min \sum_{i=1}^{N} w_i (\Pi_i^{market}(\tau_i, K_i) - \Pi_i^{model}(\tau_i, K_i))^2, \tag{10.1}$$

where w_i is a weight and N denotes the number of options. In our case we have five options and four maturities giving $N = 20$. The choice of w_i naturally changes the performance of the calibration. In this study two different weights will be used:

$$\begin{aligned} \text{Unit weights: } & w_i = 1, \\ \text{Relative price weights: } & w_i = (\min(1/\Pi_i^{market}, H))^2, \end{aligned} \tag{10.2}$$

where H is used to cap the weights so that OTM options do not get too much influence on the calibration. $H = 70$ will be used in this report; more on this choice when the calibration of the stochastic volatility models are discussed.

These choices of weight schemes will naturally have a different impact on the calibration. Using unit weights one aims at trying to minimize the overall discrepancy in absolute squared price difference. Using the relative price weights will instead minimize the relative squared price difference. Inuitively, if we have unit weights more effort is going to be put into correcting the options with high

prices. Hence, with this weighting scheme we can expect that the relative price error will be bigger for OTM options than for ITM options. Using a weighting scheme minimizing the squared relative prices takes the OTM options relatively more into consideration.

The calibration is done for all 263 days, and for each day a set of parameters is obtained for each model. In order to validate the calibration an error plot is shown for 2M and 9M using both unit and relative price weights. The error plots display an empirical 95% confidence interval for the relative price errors for different moneyness levels. The relative price error is measured as the difference between market price and model price divided by market price. Hence a negative value means the model overestimates the price and a positive value that the model underestimates the price. These are shown with blue bars. A red line connects the average error for each moneyness level.

The implementation in MATLAB is done using the function LSQNONLIN, which finds the parameters that minimizes the squared sum of a vector. In our case we have a vector consisting of 20 indices where each index is the difference between the market price and the model's price.

10.2 Calibrating Black-Scholes

There is only one parameter to fit in the Black-Scholes model: the volatility σ. Since the data is quoted in implied volatility a soundness check can be performed after the calibration - a volatility deviating too much from the ATM volatility for instance is not reasonable. A search interval of $[0, 2]$ was used with an initial guess of 0.1. This is done for the first day, in the following days the initial guess is chosen as the previous day's optimal value. The result is shown in figure 4. It is notable that the estimates of the volatility when using unit weights is higher than the estimates when using relative price weights. An explanation for this could be that when using relative weights, more effort is put into OTM call options. For the data we have at hand the OTM call options have lower implied volatility than the ITM call options, hence contributing more when using relative weights.

When studying figure 5 and 6 one notices that the confidence interval becomes narrower for the longer maturity, especially for the OTM options. An explanation for this is naturally that for an OTM option with short time to maturity there is much uncertainty in the price compared to an option with longer time to maturity. More interestingly, comparing the difference between the unit and relative weights, one can notice that the confidence interval for unit weights is wider for the OTM options and more narrow for the ITM options compared to the relative weights.

In figure 7 the market call prices for 9 months are shown, together with the Black-Scholes prices are shown using both unit and relative weights. Note that the prices for the OTM call options are better approximated when using relative weights. The inverse relation can be seen for ITM options, especially for the ATM options.

10.3 Calibrating Heston

The calibration procedure gets much trickier for stochastic volatility models since more parameters need to be estimated, and the computational effort is

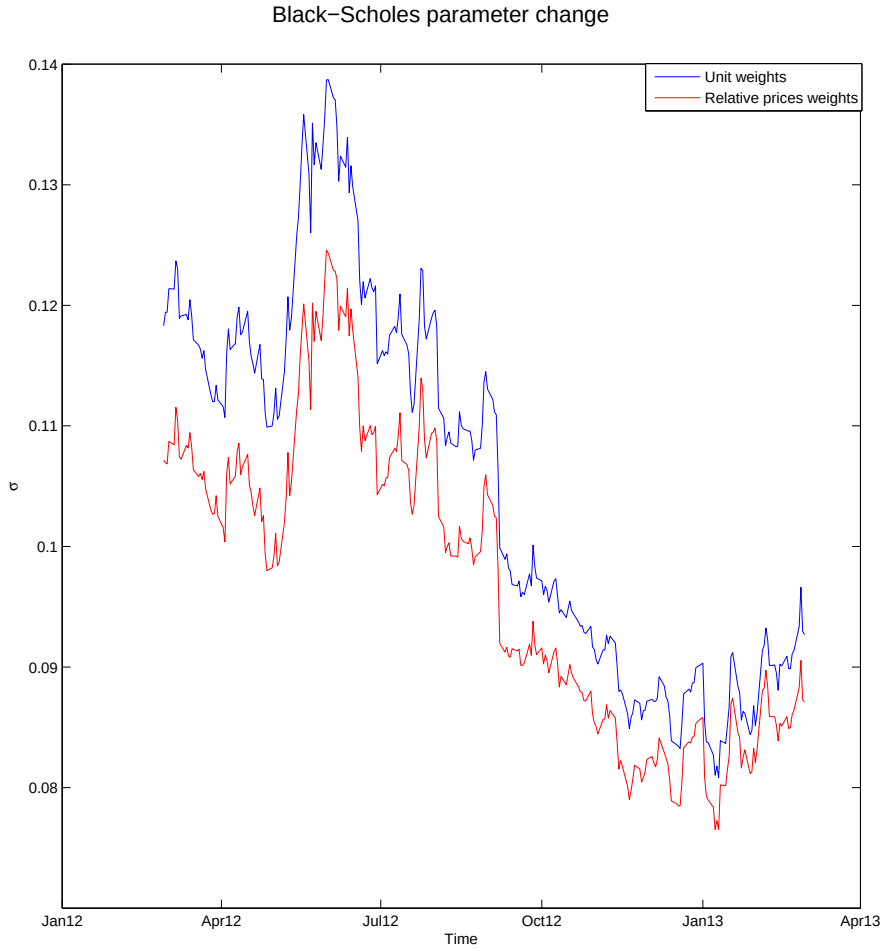

Figure 4: The volatility development for the Black-Scholes model using the two different weighting schemes. Note that the volatility is estimated lower for relative price weights than for unit weights.

much higher. The prices are calculated according to (6.2), using numerical integration in MATLAB. Note that the upper limit in the integral is infinity, meaning we need to truncate the upper limit. According to Moodley (2005), an upper limit of 100 is more than enough for practical purposes. In order to be on the safe side an upper limit of 500 was used. The function QUADL was used for the numerical integration. The function uses Lobatto quadrature, which is similar to Gaussian quadrature. The method approximates the integral by a weighted sum of function evaluations in specific points, chosen suitably. The tolerance level was by default set to 10^{-6}, and this tolerance level is set to 10^{-10} instead.

The so called Feller condition states that if $V_0 > 0$ and $2\kappa\theta > \sigma^2$, then the variance will never become negative (see e.g. Clark (2012)). In order to incorporate this condition in our calibration we can start by forming a new variable $F = 2\kappa\theta - \sigma^2$. We can then force F to be positive in our optimization

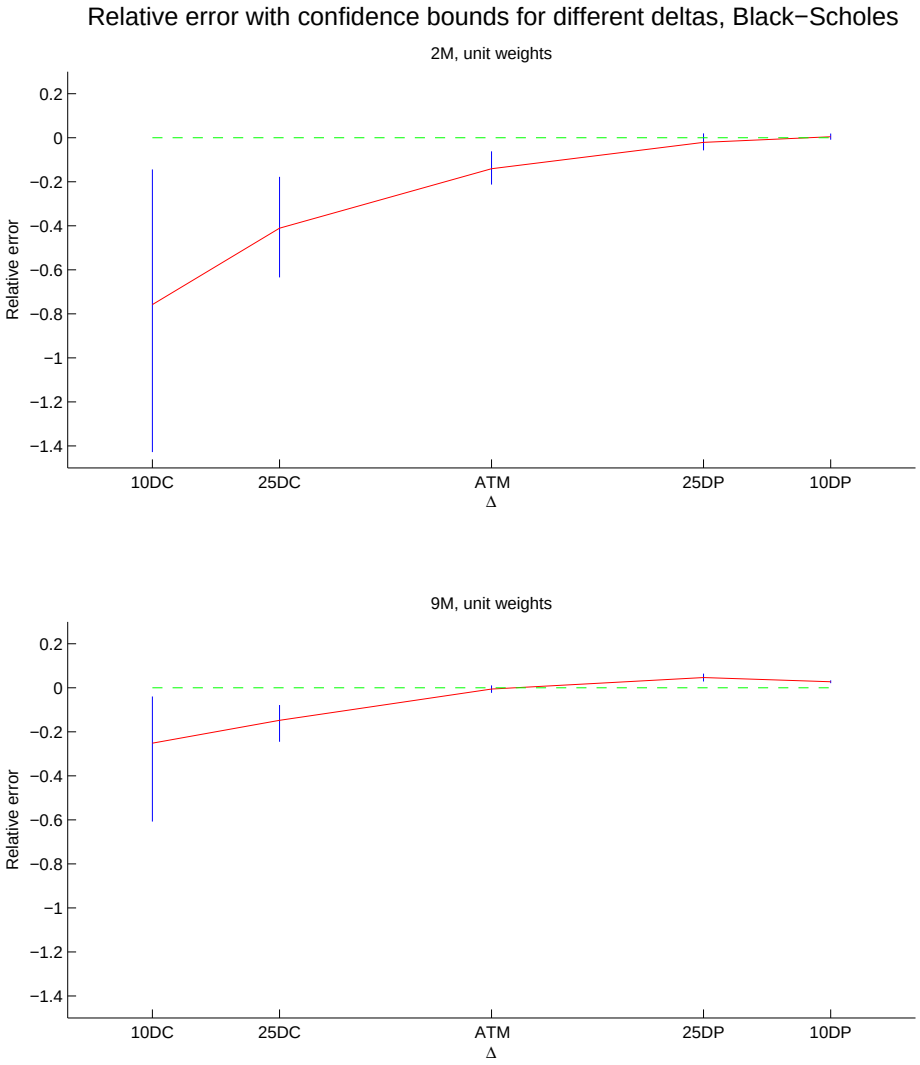

Figure 5: Unit weights, BS: 95% confidence interval for the relative price errors for different moneyness levels. Note that the error is scattered for low moneyness levels and more clustered for higher moneyness levels. The error is overall negative meaning that the model tends to overestimate the price.

routine. Hence a positive bound for F was set, a transformation according to $\kappa = (F + \sigma^2)/(2\theta)$ can then be sent to the pricer, and after the parameters are obtained, κ is reconstructed from F.

A search for the optimal parameters was done for the first day, and for the next days the previous day's optimum values are used as initial guess for that day. The optimizer needs to know boundaries on the parameter values, spcified by the user. These boundaries can be obtained by a combination of economic and mathematical arguments. For instance, the correlation must mathematically be between -1 and 1, and a mean reversion rate of 150% seems econom-

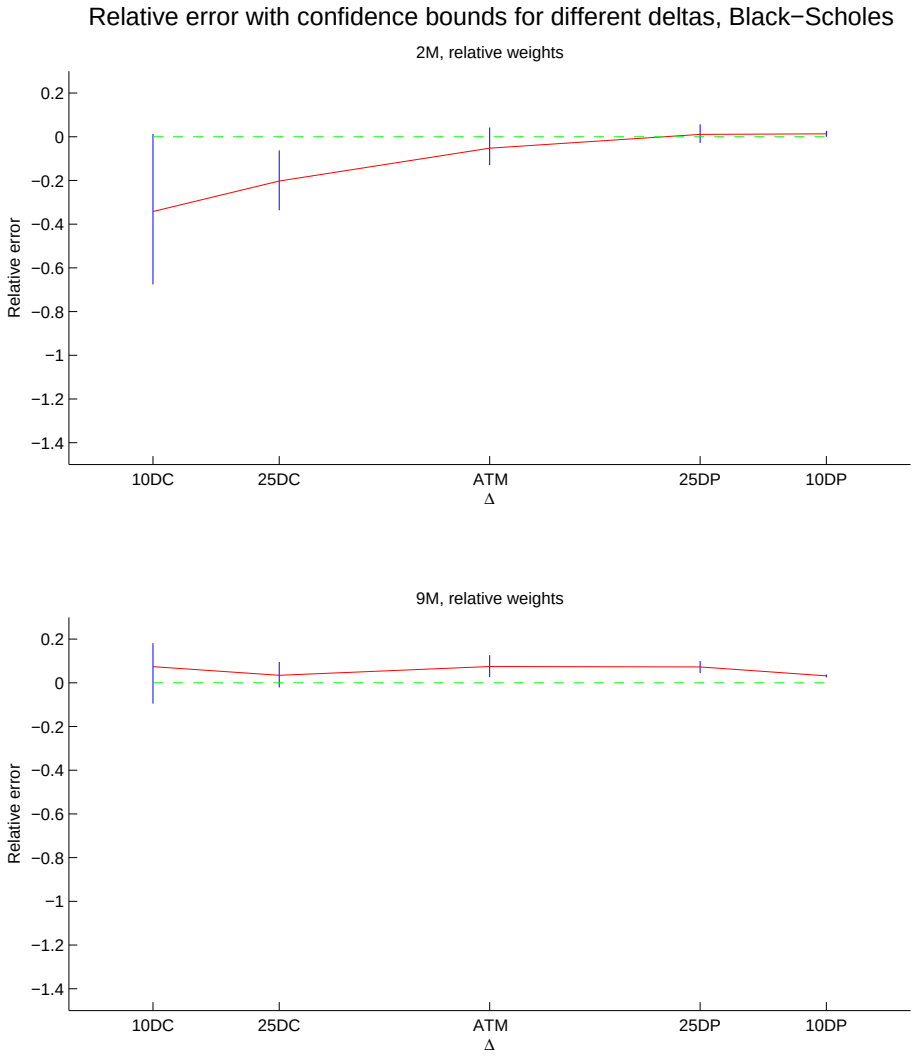

Figure 6: Relative price weights BS: 95% confidence interval for the relative price errors for different moneyness levels. Note that the error is scattered for low moneyness levels and more clustered for higher moneyness levels. The error is overall negative for 2M and positive for 9M. Hence in this case the price is overestimated for 2M and underestimated for 9M.

ically unreasonable. The boundaries on the parameters and V_0 were chosen as:

$$F : [0, 20], \theta : [0, 1], \sigma : [0, 5], \rho : [-1, 1], V_0 : [0, 1]. \tag{10.3}$$

In order to find parameters for the first day, the parameter space was partitioned into a hypercube. More specifically, each parameter was allowed to take three distinct values, partioned uniformly over the allowed interval for that parameter. Since each parameter can take three different values, there are $3^5 = 243$ different parameter combinations, and each combination was used as the start-

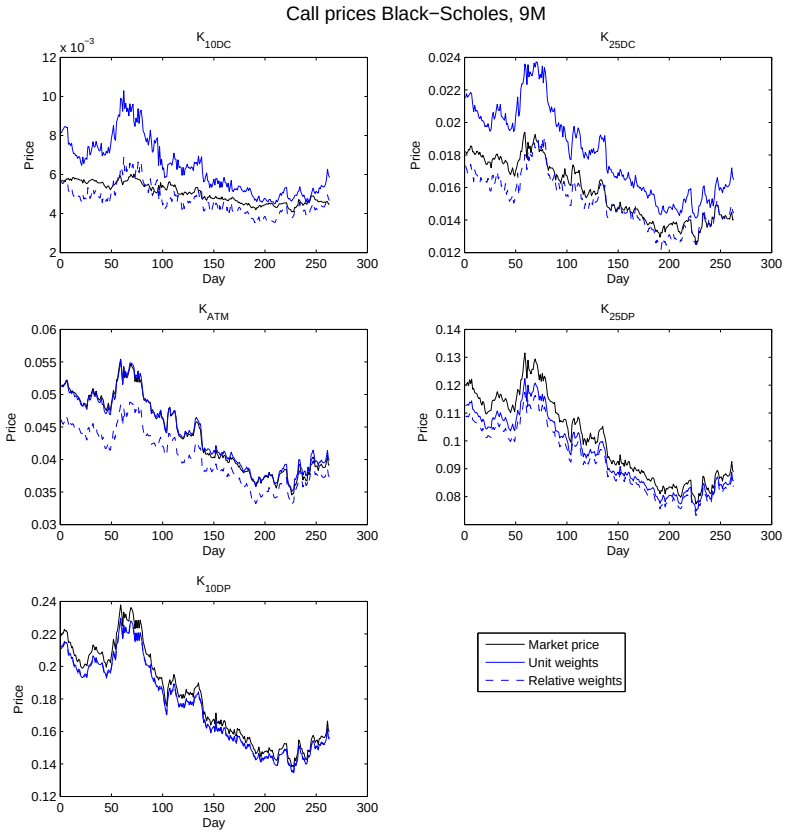

Figure 7: Market call prices for 9M maturity and Black-Scholes prices using both unit and relative weights.

ing guess to the optimizer. The result is shown in figure 8 where the parameter development is shown for all days using both unit and relative weights.

One can notice right away that both weighting schemes agree on V_0. For the parameters we notice the much more erratic behaviour for the relative price weights. The reason for this is probably that since the relative weighting scheme takes OTM options (where there is more uncertainty compared to ITM options) more into account, the parameters are more likely to vary from day to day.

Another interesting note to make is that the estimates for ρ are similar, but are offset compared to each other. The unit weights seem to predict a more prominent left tail skew in the return distribution. The unit weight scheme takes ITM options (low strike) more into account than the relative weight scheme, for this reason possibly estimating the skew as more prominent.

In figures 9 and 10 the error plots for the calibration are depicted. Note that for OTM options the confidence bound covers the zero more or less in the middle of the confidence bound for relative weights. For the unit weights an underestimation is consistently being made for the prices. More effort is put into 'fixing' the prices for ITM call options.

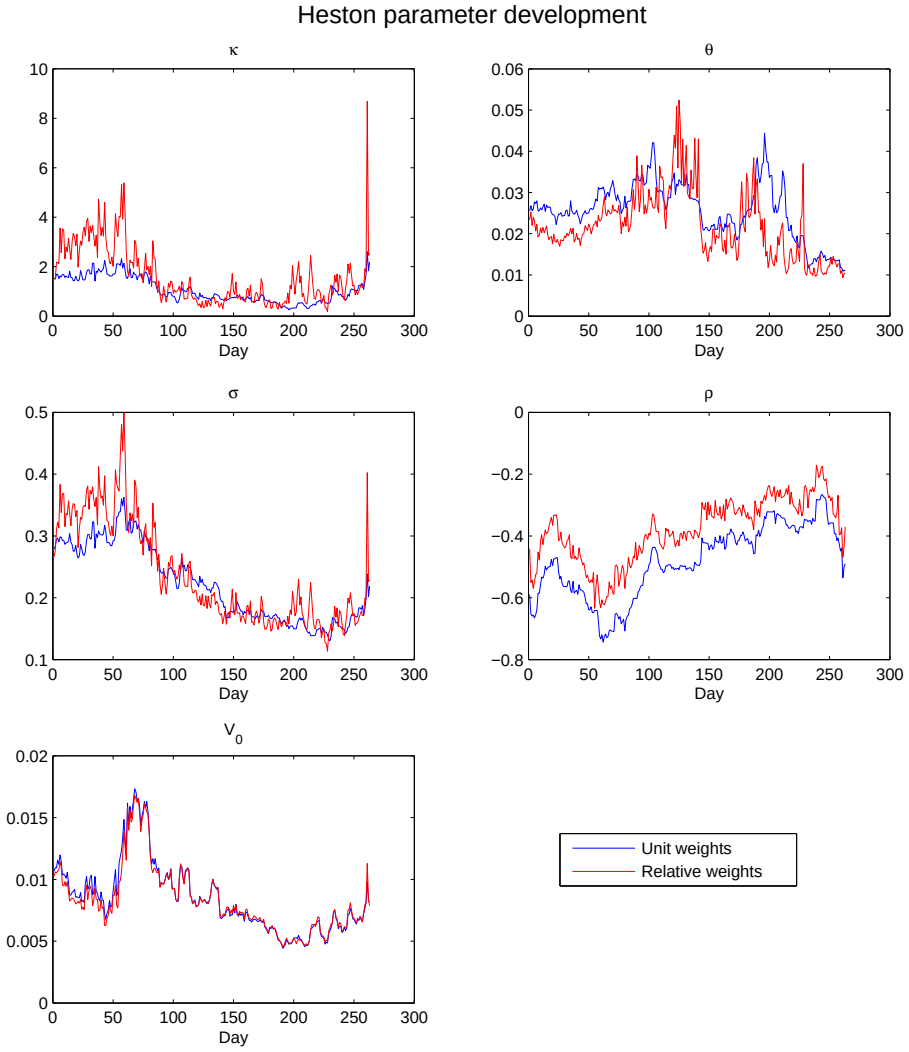

Figure 8: Parameter development for Heston with unit and relative weights.

10.4 Calibrating Bates

The calibration procedure is analogous as in the Heston case. The same upper limit in the integration and the same tolerance was used. Since the two models are similar it seems reasonable to assume that the parameter values they have in common are similar. Therefore the initial guess for the first day was the same for the mutual parameters and V_0. The boundaries on λ, μ_J and σ_J were chosen as:

$$\lambda : [0, 10], \mu_J : [-3, 3], \sigma_J : [0, 1]. \tag{10.4}$$

When a parameter set was found for the first day, that parameter set was used as the initial point for the next day.

In figure 11 the parameter development is shown for unit and relative weights.

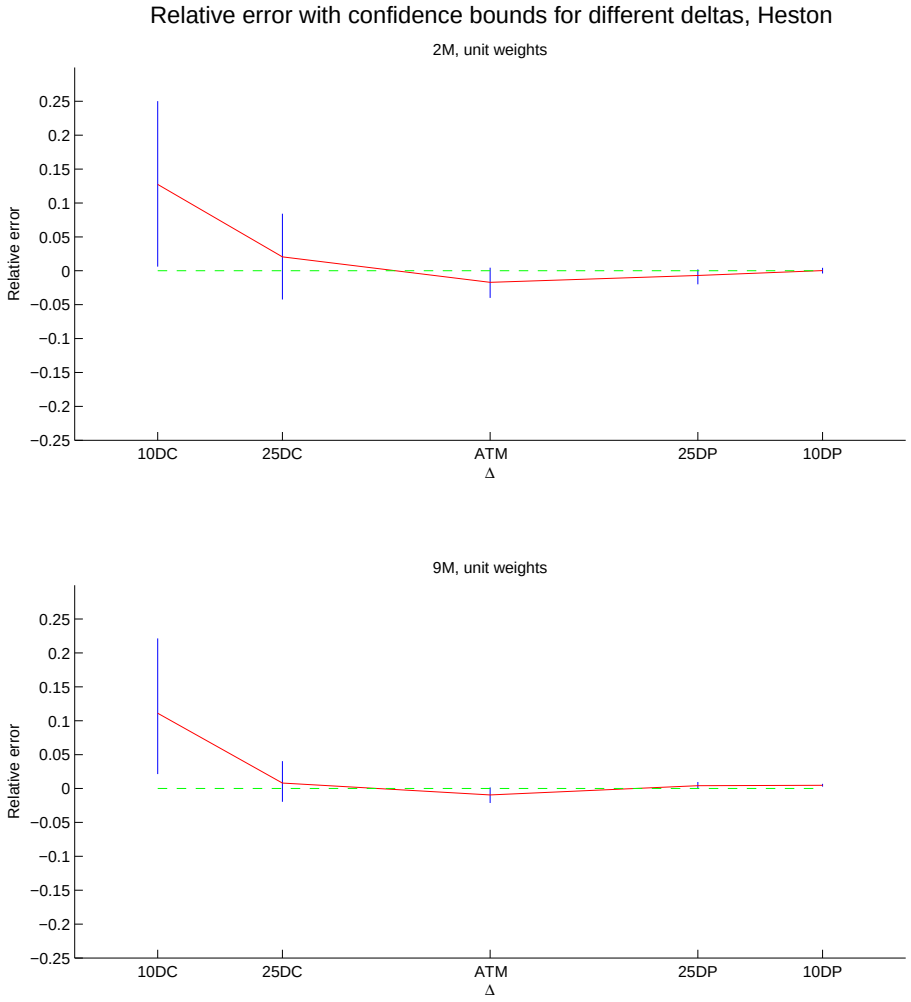

Figure 9: Unit weights, Heston: 95% confidence interval for the relative price errors for different moneyness levels.

The parameter development is rather smooth for the unit weights. Note however that the parameter development for the relative weights are extremely erratic. It seems economically unsound with a long run variance for some days suddenly jumping to 100%, and the correlation suddenly jumping to minus one. Apart from having economically unsound values, it also seems strange that the parameter values should change so abruptly from one day to another. After all, the model is specified to have constant parameter values.

This is where the choice of H in the weighting scheme seems crucial. If one chooses H as very small (say 0.1) the weighting scheme will simply become the same as for unit weights (since we will choose the same weights for all options). If one chooses H very big, the weights will become extremely big for cheap options, and these will most likely play too big a role in the calibration. The approach taken to find a suitable choice of H was to study how much the prices

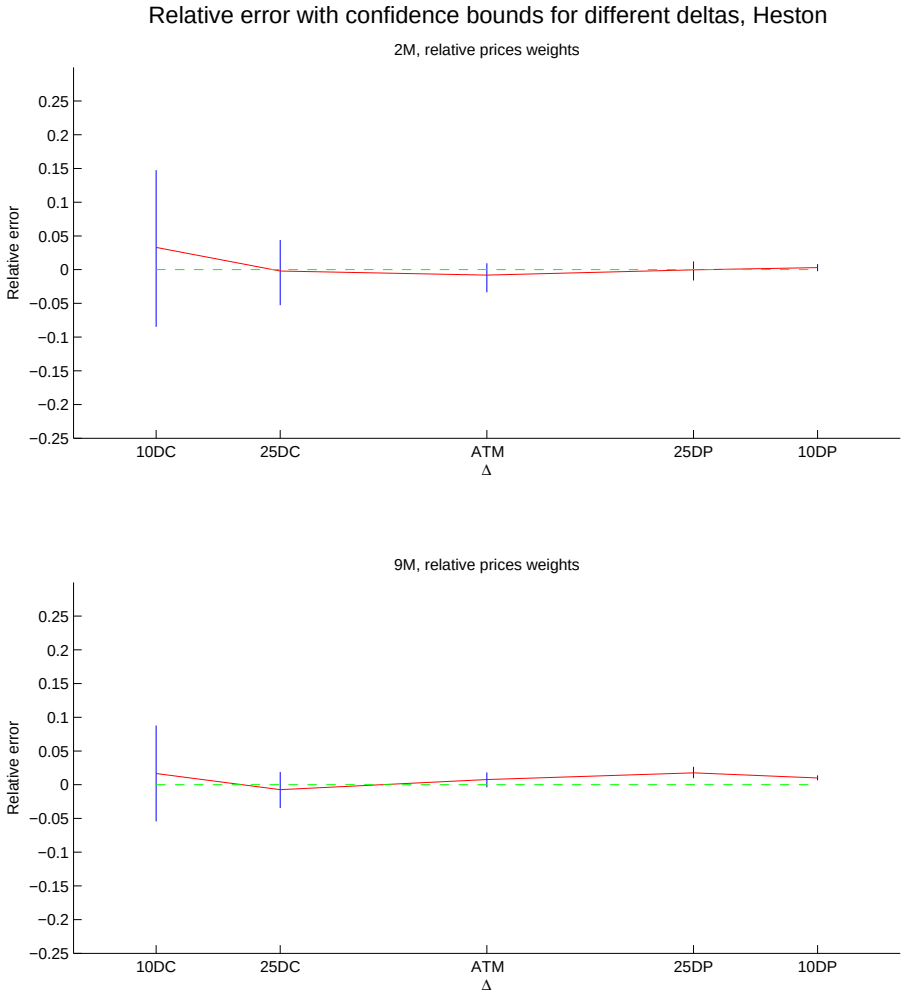

Figure 10: Relative weights, Heston: 95% confidence interval for the relative price errors for different moneyness levels.

varied for different moneyness levels and maturities. An average was calculated, and the inverse of these prices were sorted and studied. The result is shown in figure 12, where the dashed red line indicates the choice of $H = 70$. One can notice that the cheapest options give extremely high weights compared to the more expensive options. Since we still want to see the effect of relative weights it seems unreasonable to put the cap boundary too low, since the whole point of relative weights then will be lost. An ad hoc choice of $H = 70$ was made for these reasons.

It can be mentioned that several different choices of H was tried to study the different calibration performances. Even when the cap boundary was chosen as low as 25 the calibration for the relative weights derailed completely for some days. The calibration performance for both Black-Scholes and Heston was rather insensitive to the choice of H.

Since the calibration for the relative weights did not produce plausible pa-

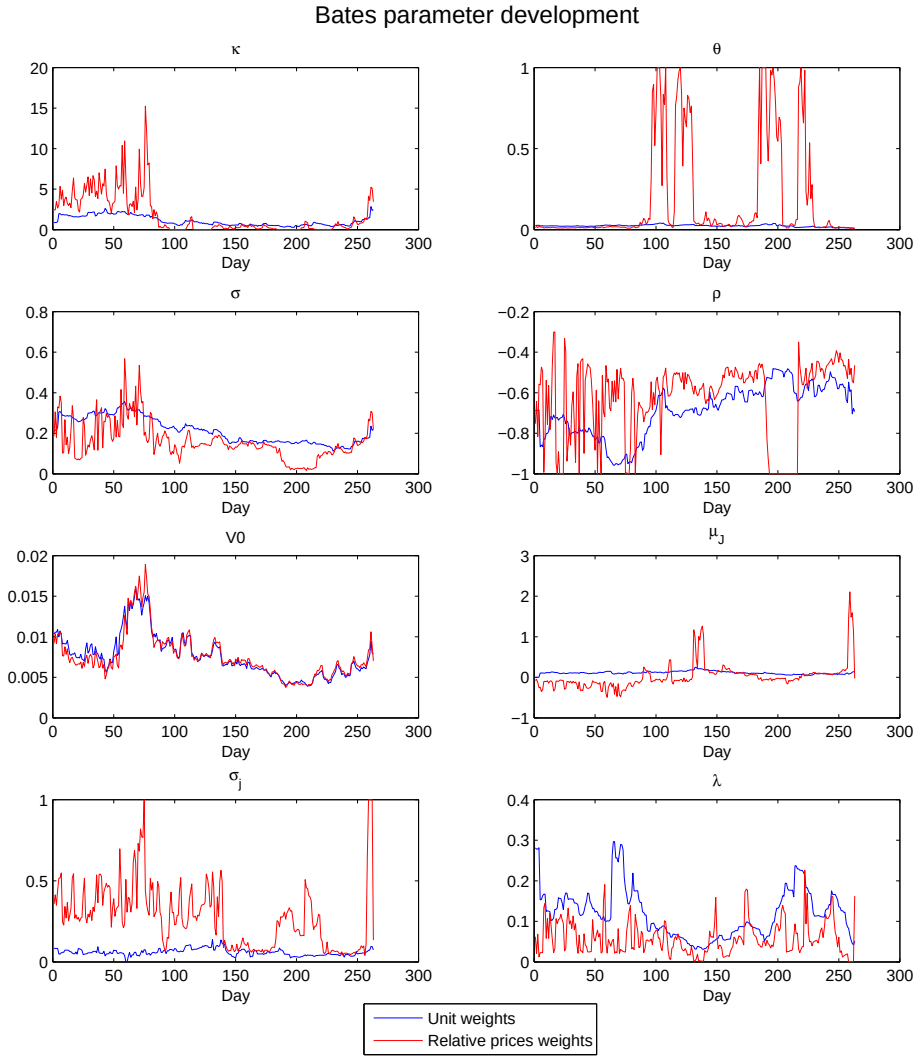

Figure 11: Parameter development for Bates with unit and relative weights.

rameter values, we conclude that it doesn't seem possible to calibrate the data to Bates using relative weights. From now on only unit weights will be considered for Bates.

In figure 13 the 95% confidence interval for the relative price errors are shown for different moneyness levels.

Note the similarity between Heston and Bates. In the Appendix the calibration result is shown quantitatively, where the mean of the errors are shown for the models and all maturities. One notices that Heston and Bates perform similarly.

In figure 14 the call prices for 9M are shown for Heston and Bates. Note the big difference for OTM options where Bates seems to capture the prices better while simultaneously not performing worse for ITM options.

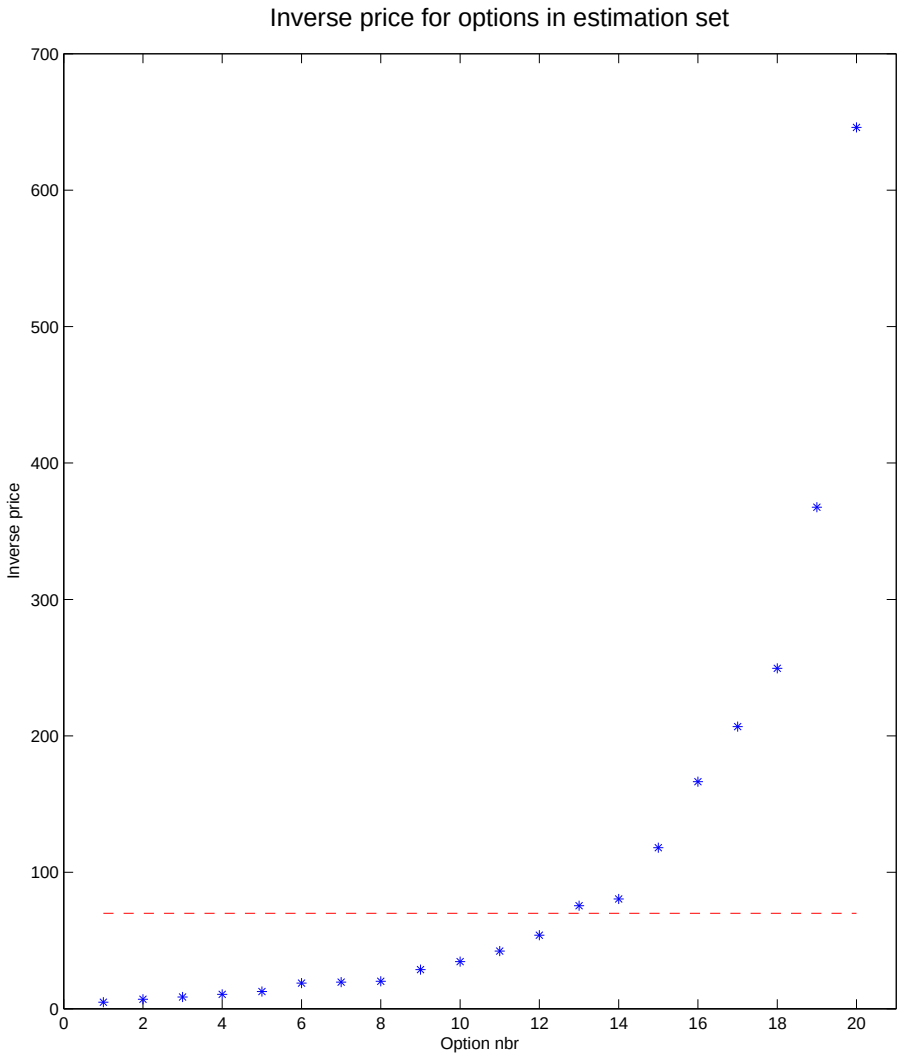

Figure 12: A mean for all prices was calculated, and the inverse was sorted and plotted. The red line indicates the choice of $H = 70$ in this calibration.

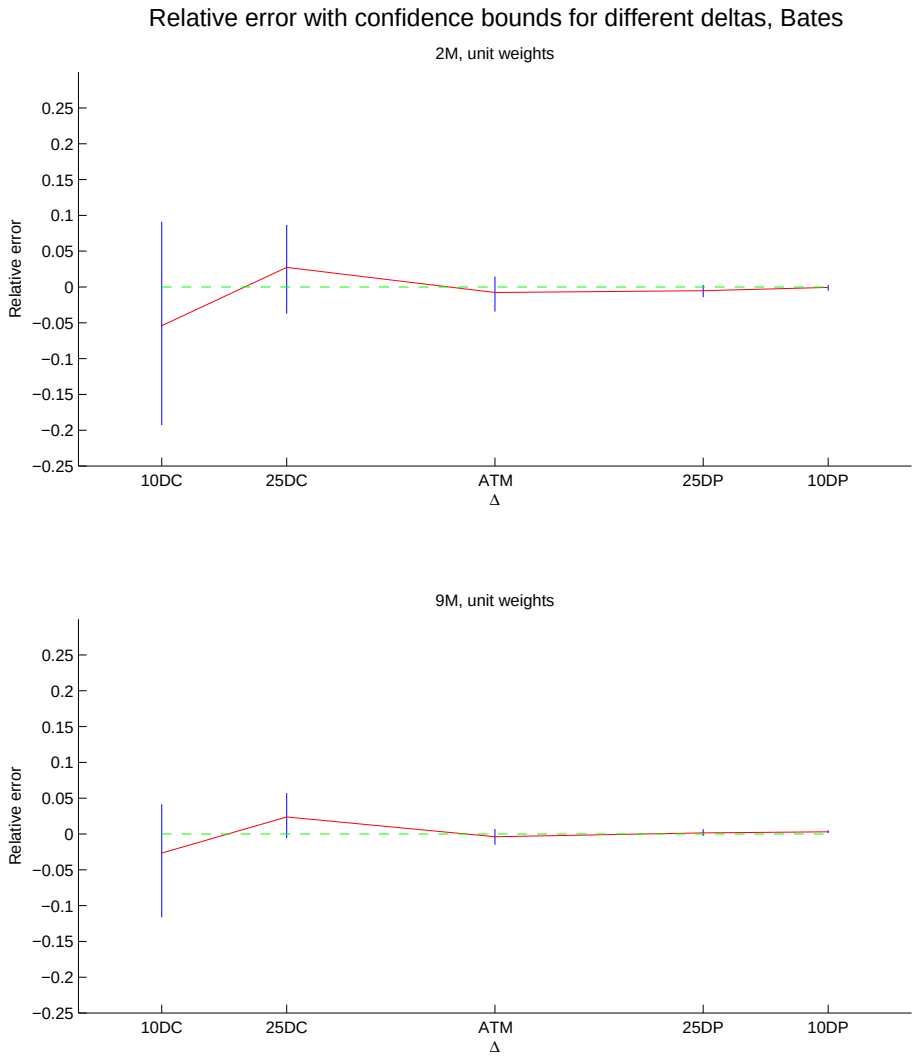

Figure 13: Unit weights, Bates: 95% confidence interval for the relative price errors for different moneyness levels.

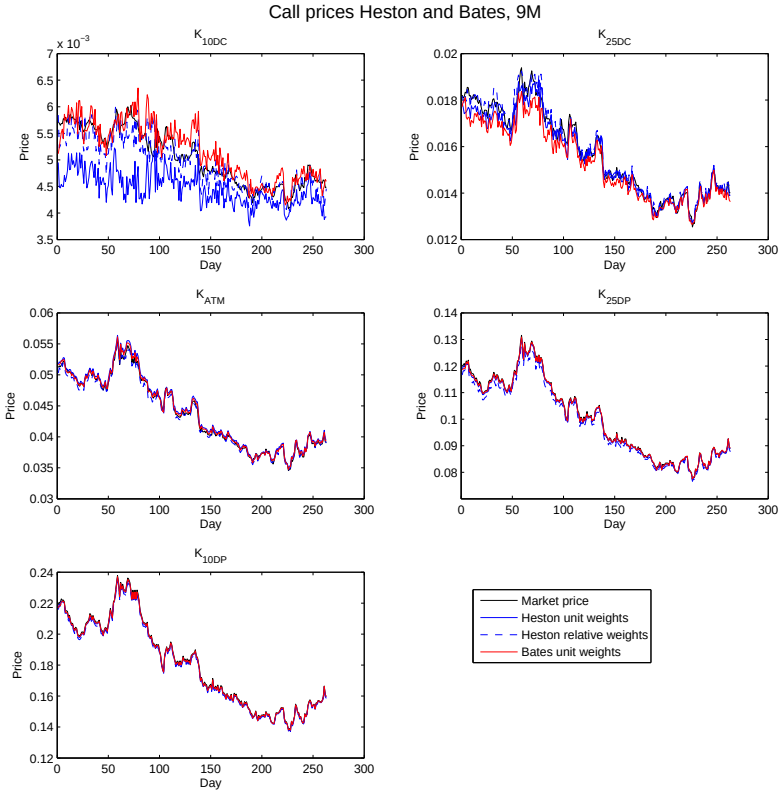

Figure 14: Call prices for Heston and Bates with 9M maturity together with market prices.

11 Implied volatility surface

11.1 Market implied volatility surface

The concept of implied volatility is rather peculiar really. We know that the Black-Scholes assumptions about constant (possibly time dependent) volatility is incorrect, since the volatility ought to be flat in strike/volatility dimension for a fixed maturity. Despite this traders are very interested in how the volatility surface looks. In this section the volatility surface will be constructed from market quotes. The implied volatility surface will also be derived from our stochastic volatility models. The difference in how the volatility surfaces look tells us something about the behaviour of the calibration.

The reader is reminded that for each maturity there are quotes on five implied volatility/delta pairs. In figure 15 this is depicted for all maturities on February 28 2012.

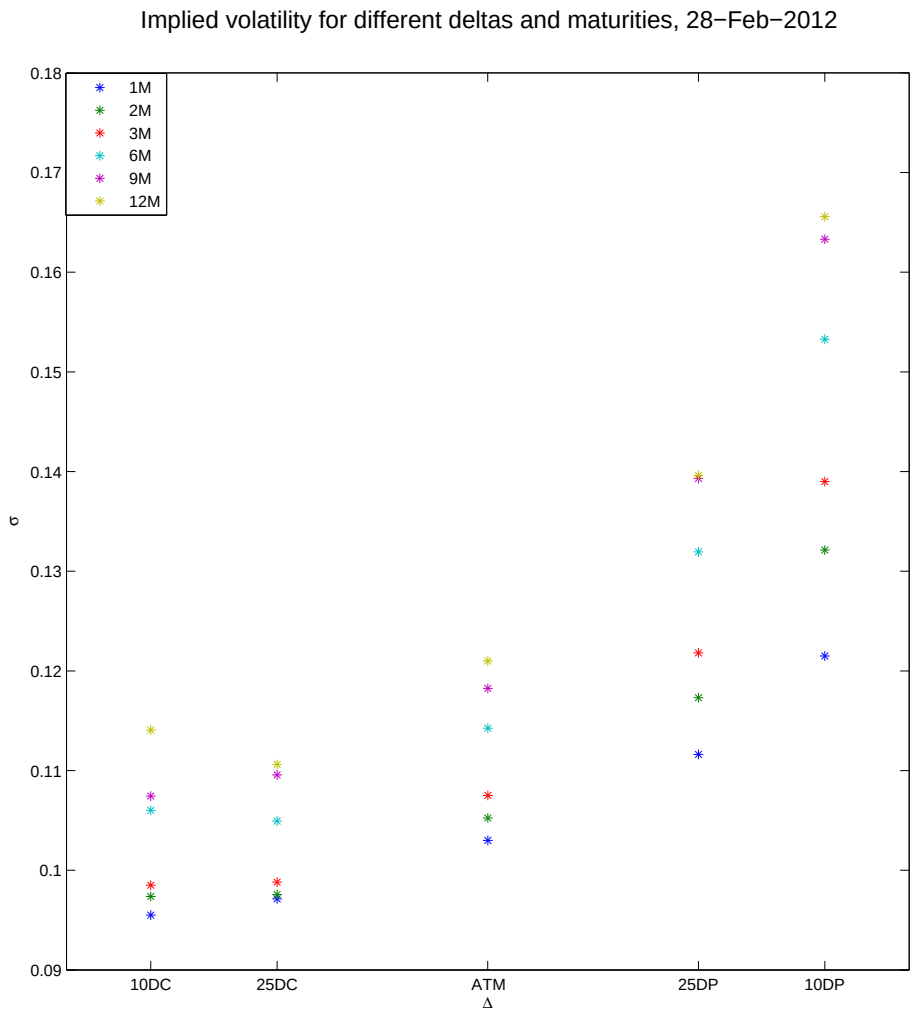

Figure 15: The volatility smile in delta/vol space for different maturities.

In order to obtain the so called volatility smile (or rather skew) one wants to perform interpolation in delta/volatility space for every fixed maturity. A straightforward approach is to assume a second order polynomial in spot delta. The apparent advantage of this approach is its very simplicity (note however that we are merely interested in *interpolation*, in order to *extrapolate* a more advanced approach ought to be taken). Hence, in delta/vol space we have:

$$\sigma(\Delta) = c_0 + c_1\Delta + c_2\Delta^2. \tag{11.1}$$

For each maturity one wants to find parameters c_0, c_1 and c_2 such that we minimize the squared discrepancy between the volatility and our interpolation curve. MATLABs backslash operator was used to quickly find these parameters (this operation finds an approximate solution to the equation $Ax = b$).

We are now interested to construct the volatility smile also in strike/volatility space. We are interested in finding $\sigma(K)$ - not an explicit expression, but in principle we want to find the implied volatility for every given strike. Recall that the spot delta for a call option is calculated as

$$\Delta_S(K,\sigma) = e^{-r_f\tau}N(d_+) = e^{-r_f\tau}N\left(\frac{\ln\left(\frac{f(t,T)}{K}\right) + \frac{1}{2}\sigma^2(\Delta_S)\tau}{\sigma\sqrt{\tau}}\right). \tag{11.2}$$

In order to find the implied volatility from a given strike we need to solve (11.2) for σ. Note however that we have assumed a parametric form of σ: $\sigma = \sigma(\Delta_S) = c_0 + c_1\Delta_S + c_2\Delta_S^2$. For a given strike we hence need to solve (11.2) for Δ_S and then plug this value into our parametric form of σ. (11.2) can be solved using an implicit solver. In this study MATLABs FSOLVE was used with an initial guess of 0.5. To summarize:

$$\begin{aligned} &1.\ \text{For a given strike } K\text{, solve (11.2) for } \Delta_S, \\ &2.\ \text{Use (11.1) to find the implied volatility.} \end{aligned} \tag{11.3}$$

The result for a typical day can be seen in figure 16. Note that in the upper plot the moneyness level is increasing and in the lower plot it is decreasing. A high delta corresponds to high moneyness and low strike, and vice versa.

In order to obtain a whole surface and not only a curve one needs to perform interpolation in the maturity dimension as well. A first approach would be to perform linear interpolation in implied volatility. According to Clark (2012) the drawback with simple linear interpolation is that the forward volatility might become negative. Therefore a so called flat forward interpolation is proposed:

$$\sigma_{\text{imp}}^{\text{flat fwd}}(t) = \begin{cases} \sigma_1, & t < t_1, \\ \sqrt{[\sigma_i^2 t_i + \sigma_{i,i+1}^2(t - t_i)]/t}, & t_i \le t_{i+1} \text{ for } i < N, \\ \sigma_N, & t \ge t_N, \end{cases} \tag{11.4}$$

where

$$\sigma_{i,i+1}^2 = \frac{\sigma_{i+1}^2 t_{i+1} - \sigma_i^2 t_i}{t_{i+1} - t_i}. \tag{11.5}$$

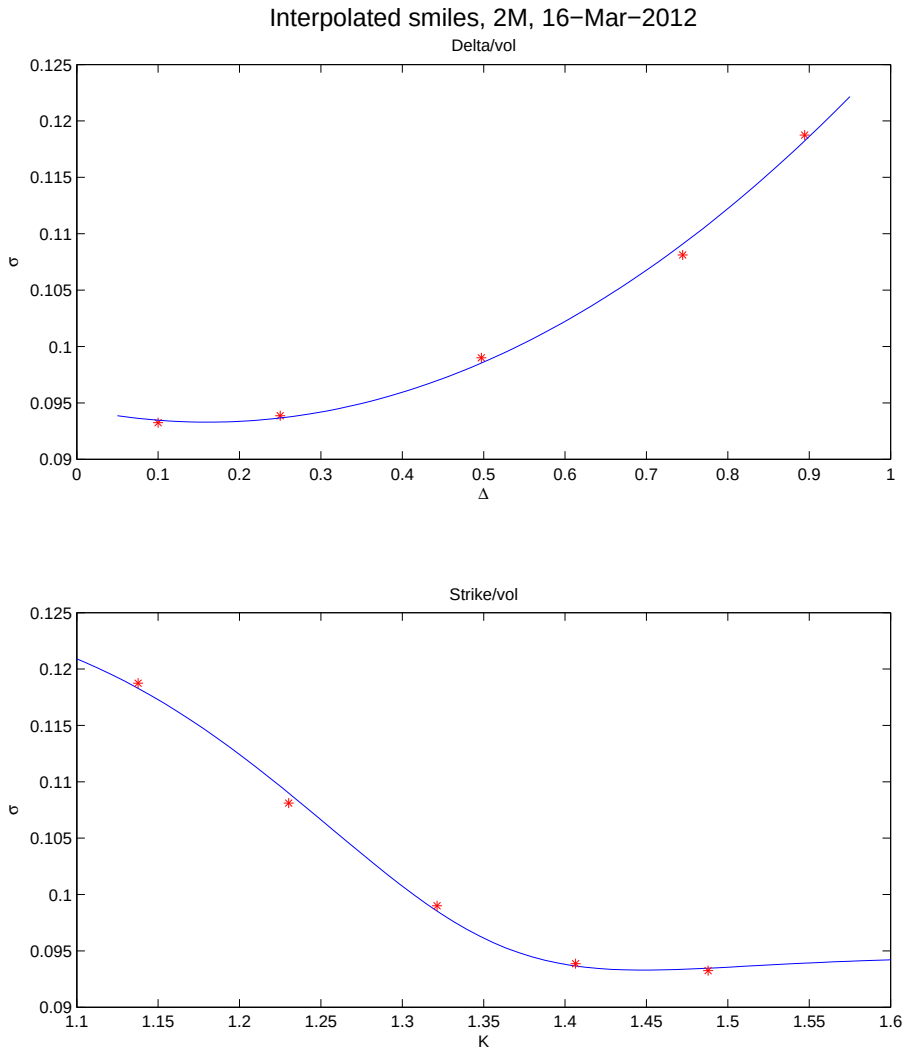

Figure 16: Upper plot: Interpolation in delta/vol space. Lower plot: Interpolation in strike/vol space. NB: In the upper plot the moneyness level is increasing and in the lower plot it is decreasing.

This clearly needs some explanation; the interested reader is referred to Clark (2012), but a heuristic explanation will be presented here. The effective variance to time t_2 is $\sigma_2^2 t_2$, and analogously for a time point $t_1 < t_2$. By additivity of variance we have:

$$\sigma_2^2 t_2 = \sigma_1^2 t_1 + \sigma_{12}^2 (t_2 - t_1) \leftrightarrow \sigma_{12} = \sqrt{\frac{\sigma_2^2 t_2 - \sigma_1^2 t_1}{t_2 - t_1}}, \tag{11.6}$$

where σ_{12} is the forward volatility between times t_1 and t_2. One can hence see σ_{12} as the market's view on how the volatility is going to behave between time points t_1 and t_2. Since we have no information in between the time points

t_1 and t_2 we assume that the instantaneous forward volatility is constant for all $t^* \in [t_1, t_2]$ (the instantaneous forward volatility is the volatility between two time points lying infinitesimally close). With this assumption the effective variance to time t^* must fulfill

$$\sigma(t^*)^2 t^* = \sigma_1^2 t_1 + \sigma_{12}^2 (t^* - t_1), \tag{11.7}$$

from which (11.4) follows. The result on a typical day is shown in figure 17. The spot for this day is equal to 1.3245, and the ATM strike level for 9M is equal to 1.3284. The observant reader will notice that the maturity is only plotted from six months to one year. If one were to plot the result from one month to a year in strike/volatility space, one would either have to depict the surface for only a narrow interval of strikes, or one would have to extrapolate in strike dimension. The reason is that since the data we have are implied volatilities for fixed *deltas*, it means that the volatilities are more scattered for higher maturities and more collected for lower maturities in strike dimension.

11.2 Model implied volatility surface

We now want to construct the implied volatility from our models. The volatility surface in the Black-Scholes case is rather uninteresting, since it will merely be a flat surface. It is more interesting to study the volatility implied by our stochastic volatility models, both with unit and relative weight (only unit obviously for Bates). Since a calibration has been made for all days we can for any strike and maturity calculate the price of the corresponding option. We do not have the interest rates for intermediate maturities; a simple linear interpolation in time was used here. When the price has been calculated one has to find out what volatility this price *implies*. The price is a non-decreasing function of σ (this can for example be seen from the fact that the vega - $\frac{\partial \Pi}{\partial \sigma}$ - is always positve).

Because of this we can use a bisection method to find the implied volatility. For a given price target V_{target} we want to find the volatility that implies this price. The algorithm starts with an interval in which the volatility must be, (an interval of $[0.01, 1]$ was used here) and for each iteration the interval in which the volatility must be is halved. For a given interval we take the volatility in the middle of the interval and check whether this price is higher or lower than V_{target}. If higher, we know the price is between the lower limit and the middle limit, and vice versa if the price is lower. A new volatility is chosen in the middle of the interval. Since the interval is halved each iteration, convergence is quite fast. Choosing 30 iterations gives a final interval of $(\frac{1}{2})^{30} \approx 10^{-9}$, which is more than sufficient for our purposes.

After the construction of the volatility surface it is interesting to study the difference between the market implied volatilities and the volatility implied by the models. In figures 18 - 20 the result is shown for a typical day.

Note the qualitative difference between the unit and relative weights for the Heston model. In the former the difference is relatively smooth over the whole surface, whereas in the latter the error is significantly higher for lower strikes, and lower for the highest strikes. The surface for Bates is even flatter and the error is overall closer to zero.

For a given maturity and day we can choose to look at the volatility skew in one dimension. In figure 21 we see the difference between the models.

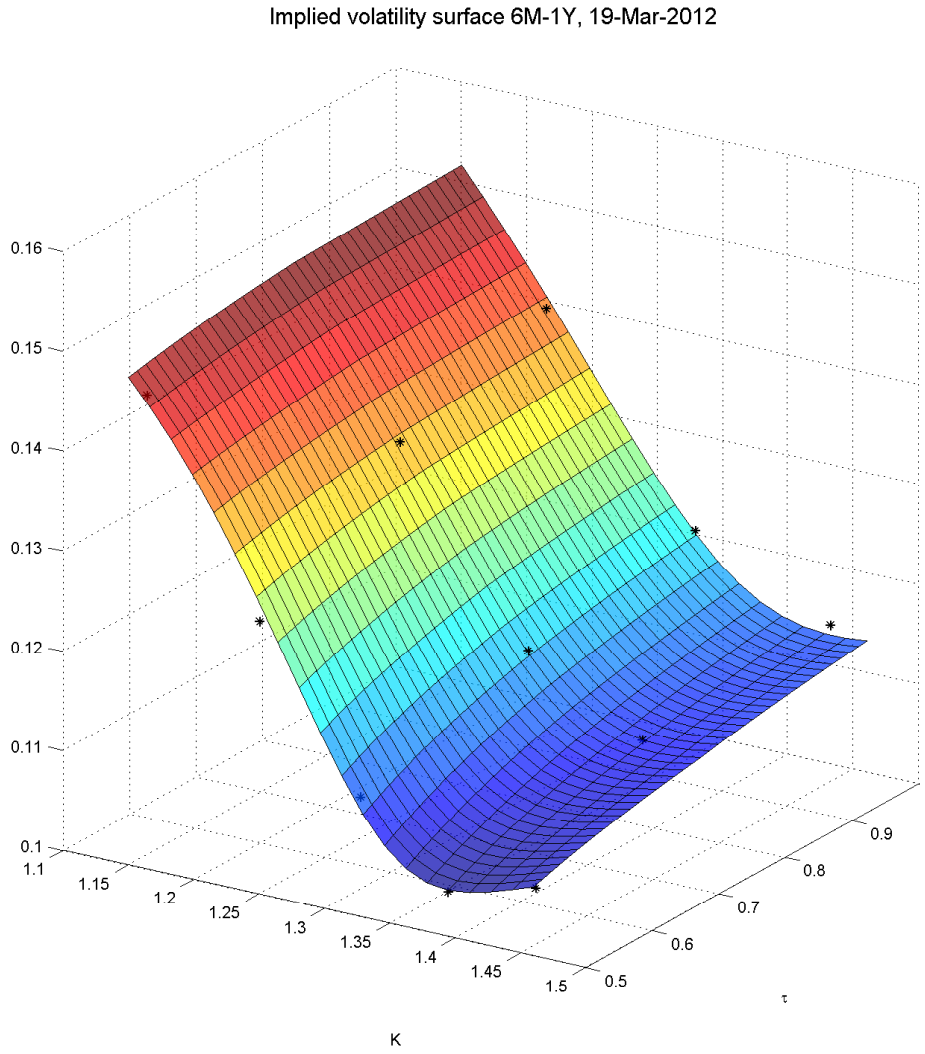

Figure 17: The market volatility surface in strike/maturity space.

To begin with, note the difference for the Heston model between unit and relative weights. The slope of the skew is steeper in the case of unit weights, compared to relative weights. Also note that the skew for the Bates model is similar to Heston unit weights for lower strikes, but behaves very differently for higher strikes. It seems like the Bates model has some more degree of freedom, stemming from the fact that it has three extra parameters.

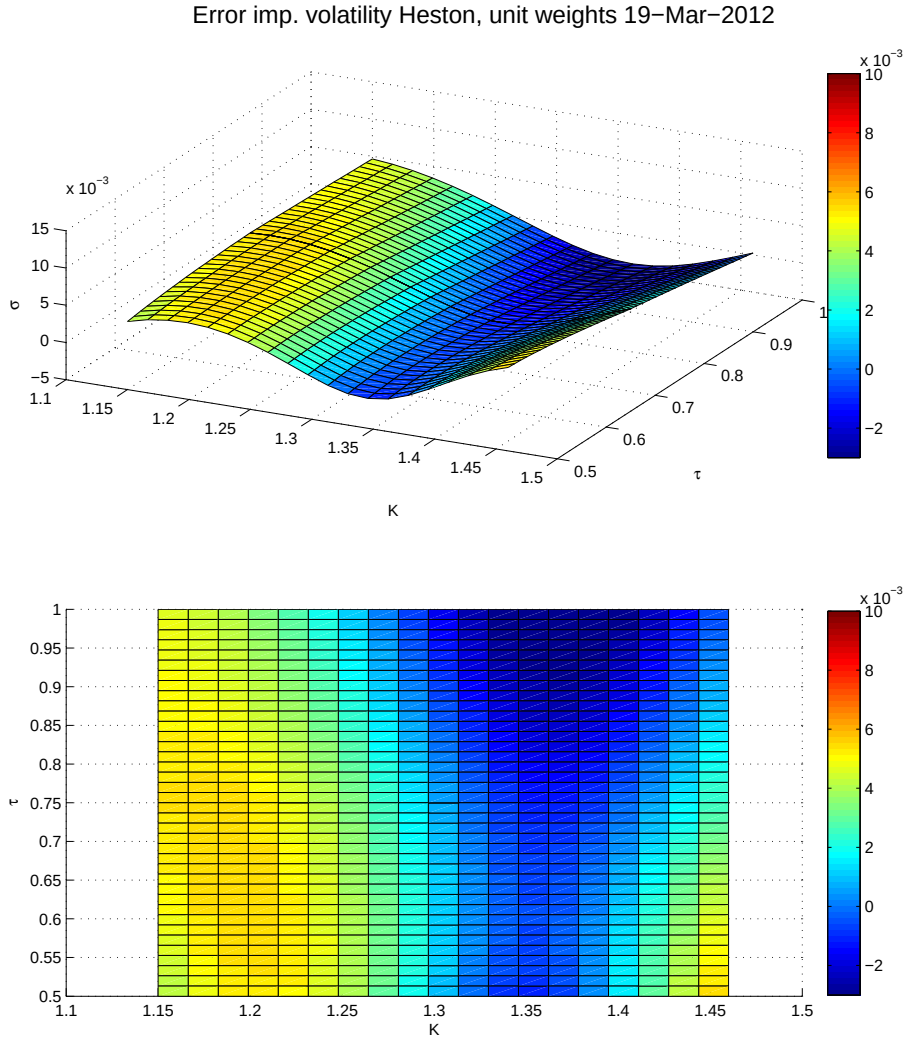

Figure 18: The difference between the market implied volatility surface and the volatility surface implied by the Heston model using unit weights.

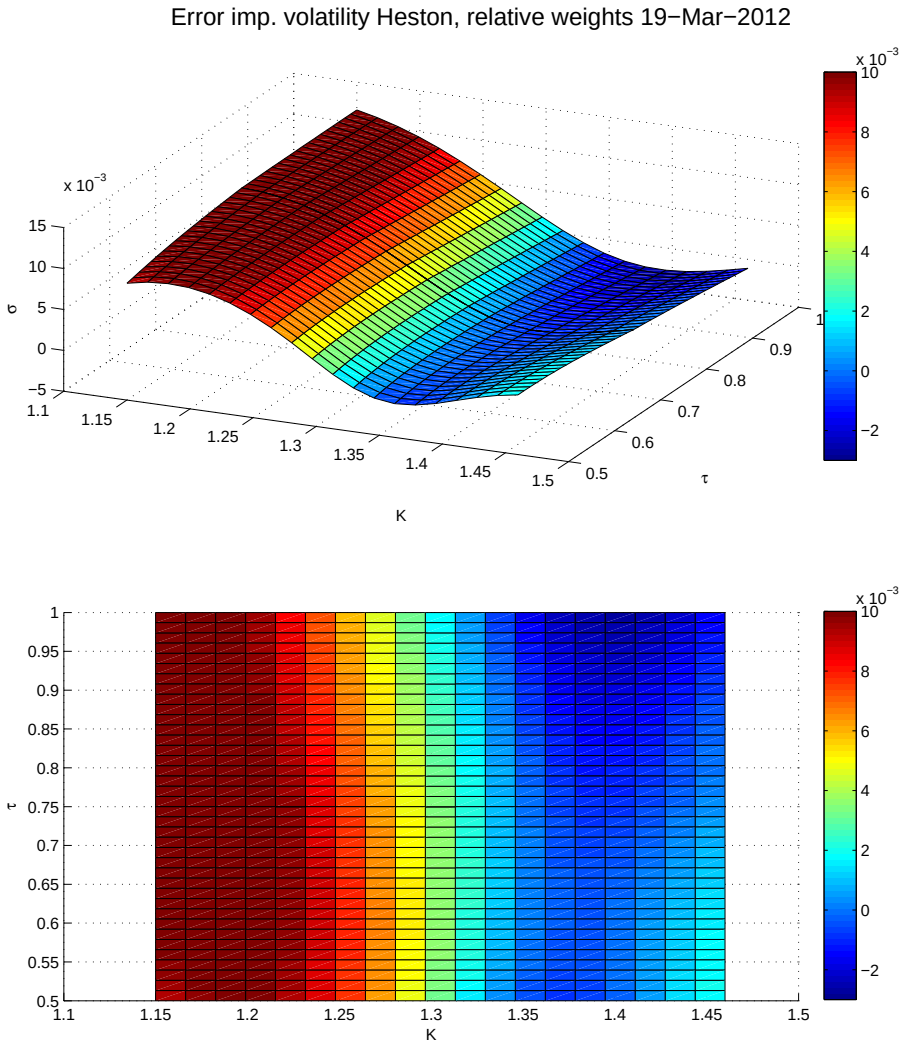

Figure 19: The difference between the market implied volatility surface and the volatility surface implied by the Heston model using relative weights.

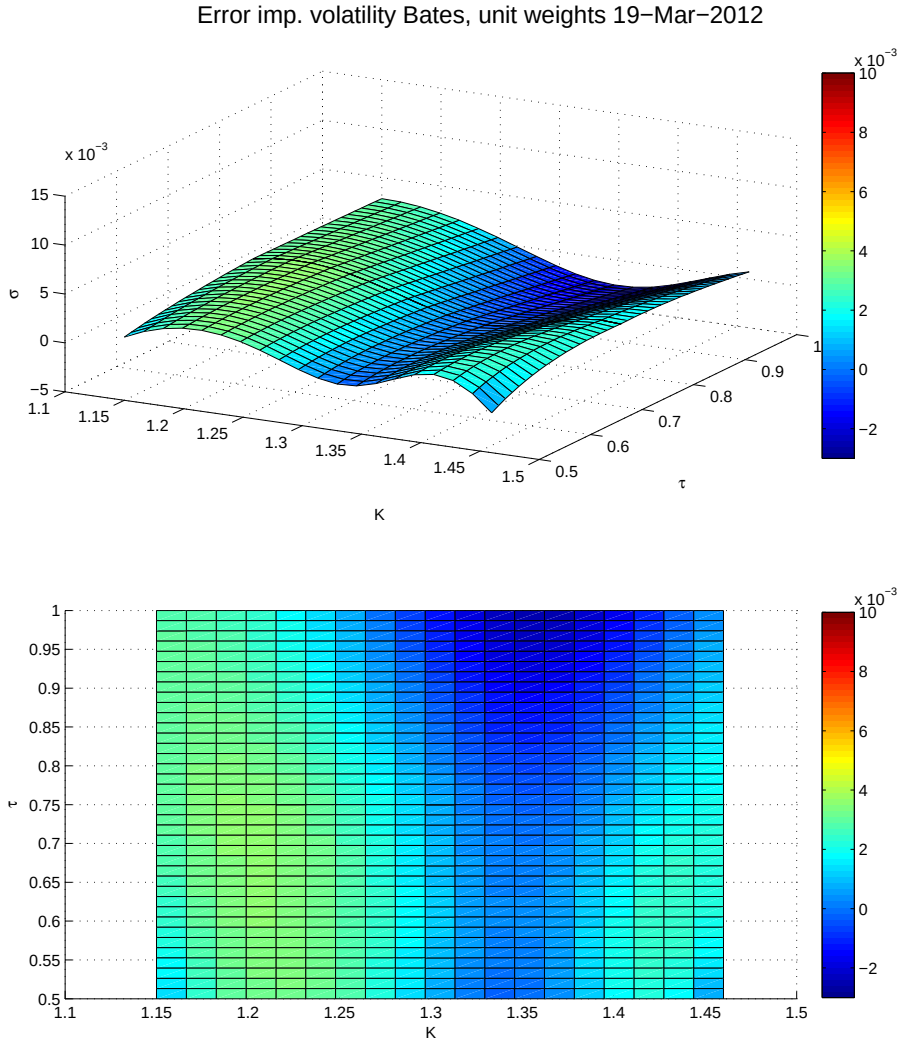

Figure 20: The difference between the market implied volatility surface and the volatility surface implied by the Bates model using unit weights.

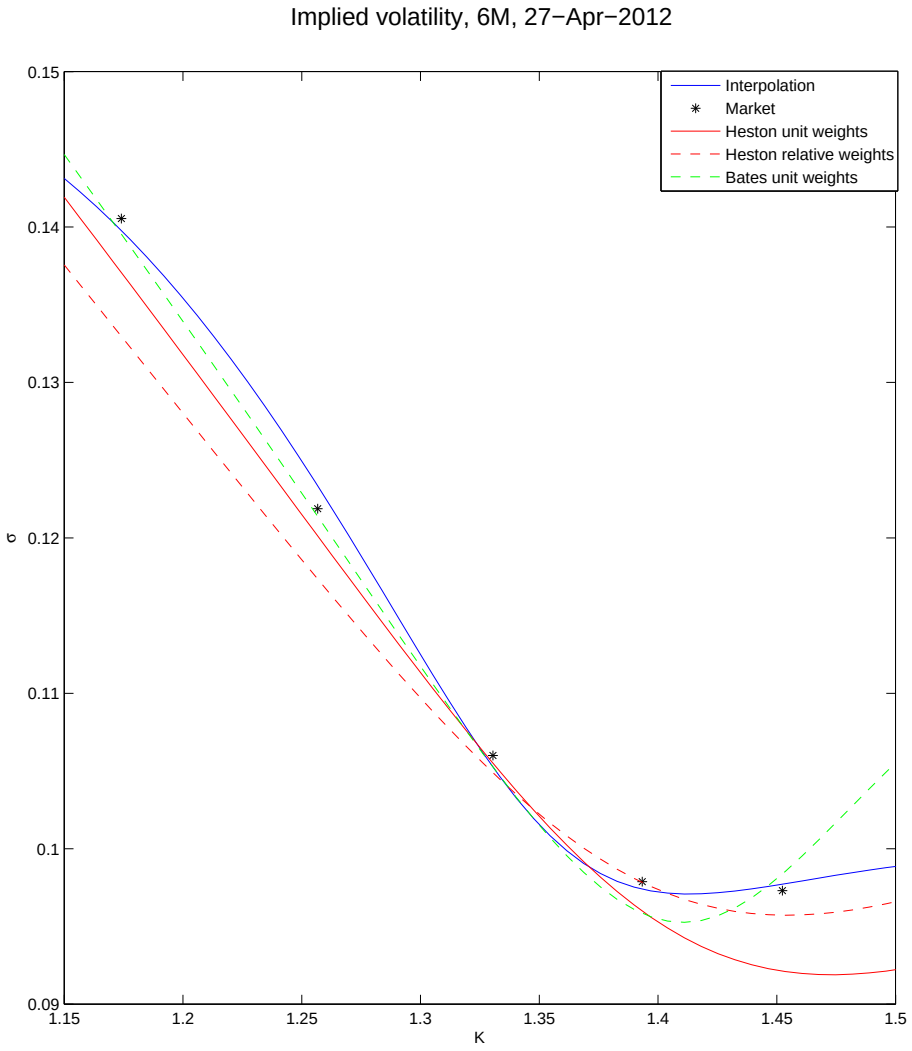

Figure 21: The volatility skew depicted by interpolation, Heston unit/relative weights and Bates unit weights.

12 Pricing barrier options

An investor might want to buy a call option, but finds it too expensive. The investor might speculate that the spot rate is not going to cross a barrier, either during the life time of the option or at the end. This is a so called barrier option, and we will start with the easiest example - a European barrier.

12.1 European barriers

The name European in this case stems from the fact that the payoff is entirely decided on the maturity day. A European up-and-out call with a strike K and barrier U has the following payoff:

$$V_T^{EUOC} = (S_T - K)^+ \mathbb{1}_{\{S_T < U\}} = (S_T - K)\mathbb{1}_{\{S_T \geq K \cap S_T < U\}}. \tag{12.1}$$

Analogously, an up-and-in call has the following payoff:

$$V_T^{EUIC} = (S_T - K)^+ \mathbb{1}_{\{S_T > U\}} = (S_T - K)\mathbb{1}_{\{S_T > U\}}. \tag{12.2}$$

The payoffs are depicted in figure 22. Note (and this can most easily be seen by studying the payoff functions) that the sum of these contracts equals a standard European call option. Hence this is a parity relation that should hold regardless of model choice.

In order to see how the different models price these contracts we want to find a market consistent price for these contracts. Thus we need to replicate the payoffs using solely call options. This cannot be done perfectly in reality since we have a discontinuity in the barrier U. We can however in theory come arbitrarily close using the following strategy for the UO contract: we go long a call with strike K, and then we short N options with strike U and go long $N-1$ options with strike $U+\xi$, where $\xi > 0$ should be close to zero. For a given ξ we now need to figure out how many options we should go short:

$$\begin{aligned} \phi^{UO}(S_T) &= \max(S_T - K, 0) - N\max(S_T - U, 0) + (N-1)\max(S_T - (U+\xi), 0), \\ \phi^{UO}(U+\xi) = 0 &\Rightarrow \max(U + \xi - K, 0) - N\max(\xi, 0) = 0 \Rightarrow \\ &\Rightarrow N = \frac{U - K + \xi}{\xi} \end{aligned} \tag{12.3}$$

For the UI contract we can instead go long $\tilde{N}$ calls with strike $U - \xi$, and go short $\tilde{N} - 1$ calls with strike U. One gets:

$$\begin{aligned} \phi^{UI}(S_T) &= \tilde{N}\max(S_T - (U-\xi), 0) - (\tilde{N}-1)\max(S_T - U, 0), \\ \phi^{UI}(U) = U - K &\Rightarrow \tilde{N}\max(\xi, 0) = U - K, \\ &\Rightarrow \tilde{N} = \frac{U-K}{\xi}. \end{aligned} \tag{12.4}$$

Since the implied volatility surface has been constructed, we could on this basis form market consistent prices for different combinations of strikes and

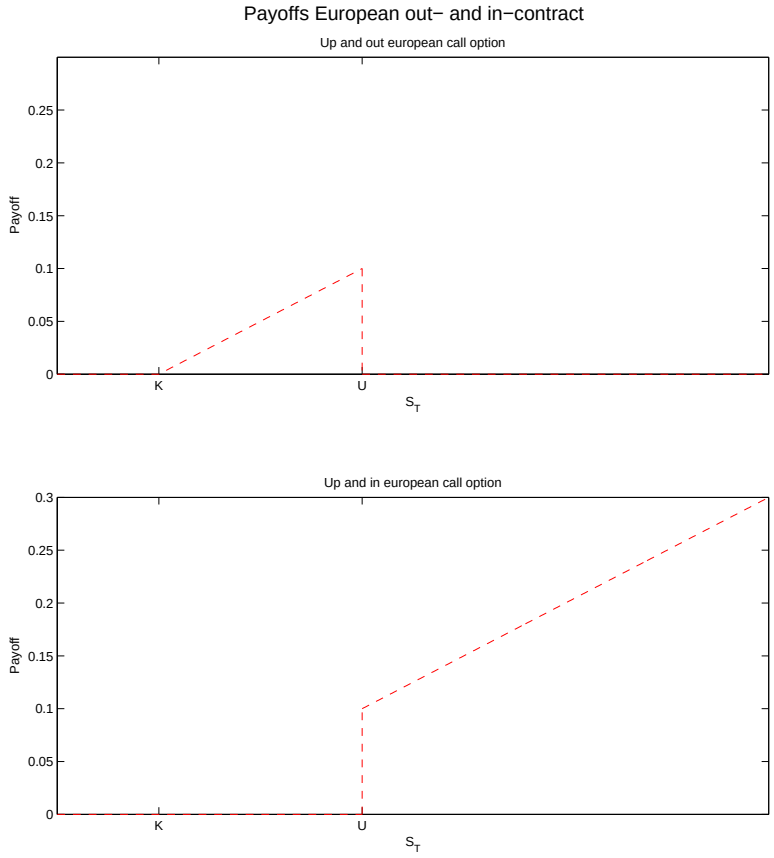

Figure 22: Payoff for an up-and-out European call and an up-and-in European call.

barriers. However, there is clearly an uncertainty here since these prices are merely interpolated. Instead one could choose to look at strikes and barriers which correspond to the known strikes. The following scheme summarizes this:

$$\begin{array}{ll} \underline{\text{Strike}} & \underline{\textit{Barrier}} \\ K_{10DP} & K_{25DP}, K_{ATM}, K_{25DC}, K_{10DC} \\ K_{25DP} & K_{ATM}, K_{25DC}, K_{10DC} \\ K_{ATM} & K_{25DC}, K_{10DC} \\ K_{25DC} & K_{10DC} \end{array} \tag{12.5}$$

There is still one obstacle one needs to address when forming the market consistent prices. When pricing the up-and-out contract one needs to price a contract with strike $U + \xi$, and $U - \xi$ in the case of an up-and-in contract. One hence needs to approximate the implied volatility in these points. At first glance one might think that one could casually approximate the volatility with

the value in the barrier point. However, since so many options are short and long, the difference could be very important.

One approach is to approximate the derivative in the smile at the barrier level. Hence one wants to approximate $\frac{\partial \sigma}{\partial K}|_{K=U}$. The implied volatility of interest is then obtained by adding respectively subtracting the derivative times ξ. The derivative is formed in the following way:

$$\begin{aligned}
&1.\ \text{The strikes of interest are } U+\xi \text{ and } U-\xi,\\
&2.\ \text{Given these strikes, find out the corresponding deltas: } \Delta_+, \Delta_-,\\
&3.\ \text{Use the parametrization } \sigma(\Delta) \text{ to find the volatilities: } \sigma_+, \sigma_-,\\
&4.\ \text{The derivative is approximated as } \frac{\sigma_+ - \sigma_-}{2\xi}.
\end{aligned} \tag{12.6}$$

For Black-Scholes, Heston and Bates, the pricing is trivial in that one simply forms a portfolio consisting of some long and some short options with different strikes. For a given date, maturity and strike one can study how the pricing for different barrier levels looks like. We can then study the difference between the models, and also between the two weighting schemes. In figure 23 the result is plotted for July 4 2012 for 9 months maturity with Black-Scholes. In figure 24 the result is shown for Heston and Bates. Both Heston and Bates seem to price the barriers quite well. One notices that the difference between the two weighting schemes becomes more pronounced for higher barriers. The result is only shown for up-and-out contracts.

12.2 Path-dependent contracts

A path-dependent contract is a contract where the payoff depends on the path the underlying has taken. The contract of interest here will be the up-and-out and up-and-in call options. These are similar to the European barriers, but with the difference that one continuosly must monitor if the barrier is ever breached during the option's life time. The price of an up-and-out and up-and-in call option is:

$$\begin{aligned}
\Pi^{UO} &= e^{-r_d \tau}\mathbb{E}[(S_T - K)^+ \mathbb{1}_{\tau_B > T}],\\
\Pi^{UI} &= e^{-r_d \tau}\mathbb{E}[(S_T - K)^+ \mathbb{1}_{\tau_B < T}].
\end{aligned} \tag{12.7}$$

The indicator function is one when the barrier has not been breached until time T, and zero otherwise.

These contracts have in the Black-Scholes case an explicit pricing formula. It is quite lengthy to append here, and the interested reader is referred to Wilmott (2006), page 408 to find the explicit prices.

The simulation techniques described in the simulation chapter will now be useful to price barriers with our stochastic volatility models. In order to price an up-and-out contract one can price an up-and-in contract, and the price of the out-contract is obtained from the parity relation between up- and in contracts. Since the pricing is done with Monte Carlo techniques there is clearly an error connected to the variance of the Monte Carlo sampler. If one were to simulate paths for different barriers, the plot would show a zig-zag pattern, and

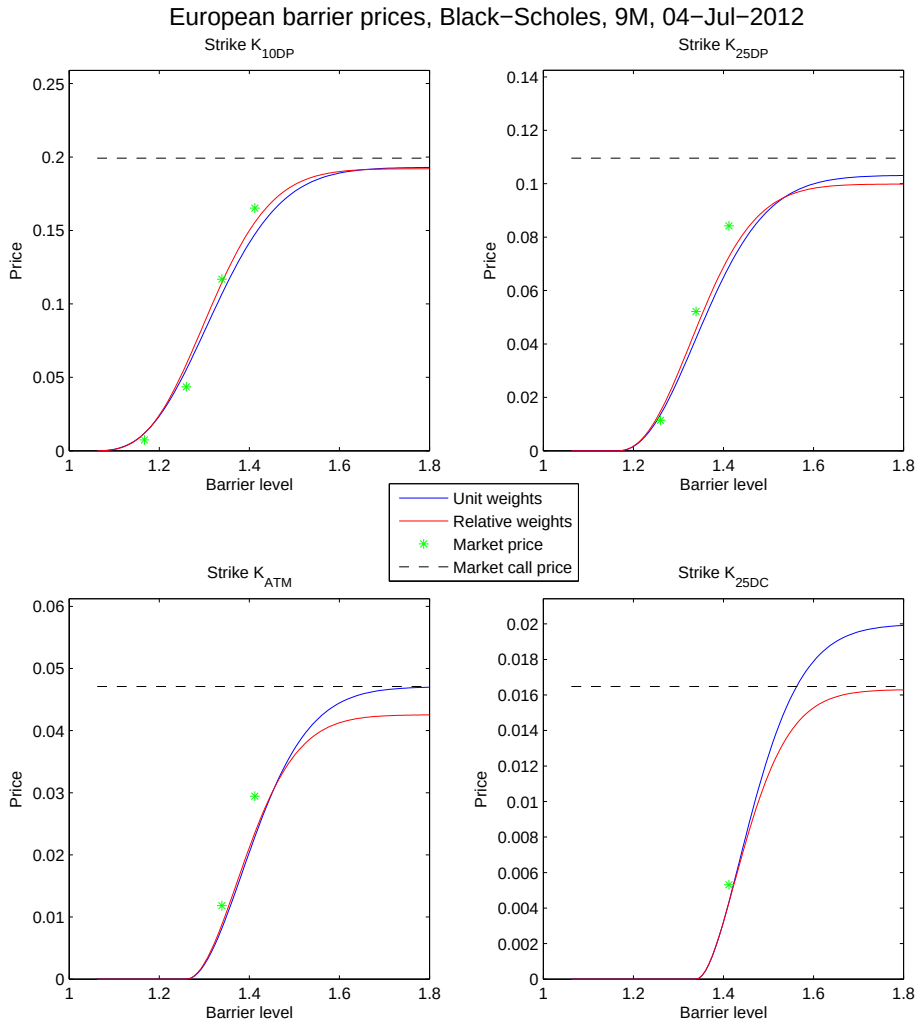

Figure 23: Prices for an up-and-out European call with different strike levels, Black Scholes.

it would be somewhat cumbersome to discern the qualitative difference between the weighting schemes and models. For this reason a seed was set for the Monte Carlo - engine, meaning that for each barrier level the same set of random numbers were used in the simulation (this is the reason why it is more convenient to simulate the up-and-in contract instead of the up-and-out contract). For the Heston model we will use exactly the same set of numbers; for the Bates model however this will not be the case. The reason for this is that if there is at least one jump occuring, an extra variance calculation is needed. The result for up-and-out is shown in figure 25 where all model prices and weighting schemes are depicted for 9M, July 4 2012.

Note that the pricing structure looks very similar as for the European barriers. A very interesting fact to note is that the pricing differences between Heston unit weights and relative weights are not as big as that of Heston unit

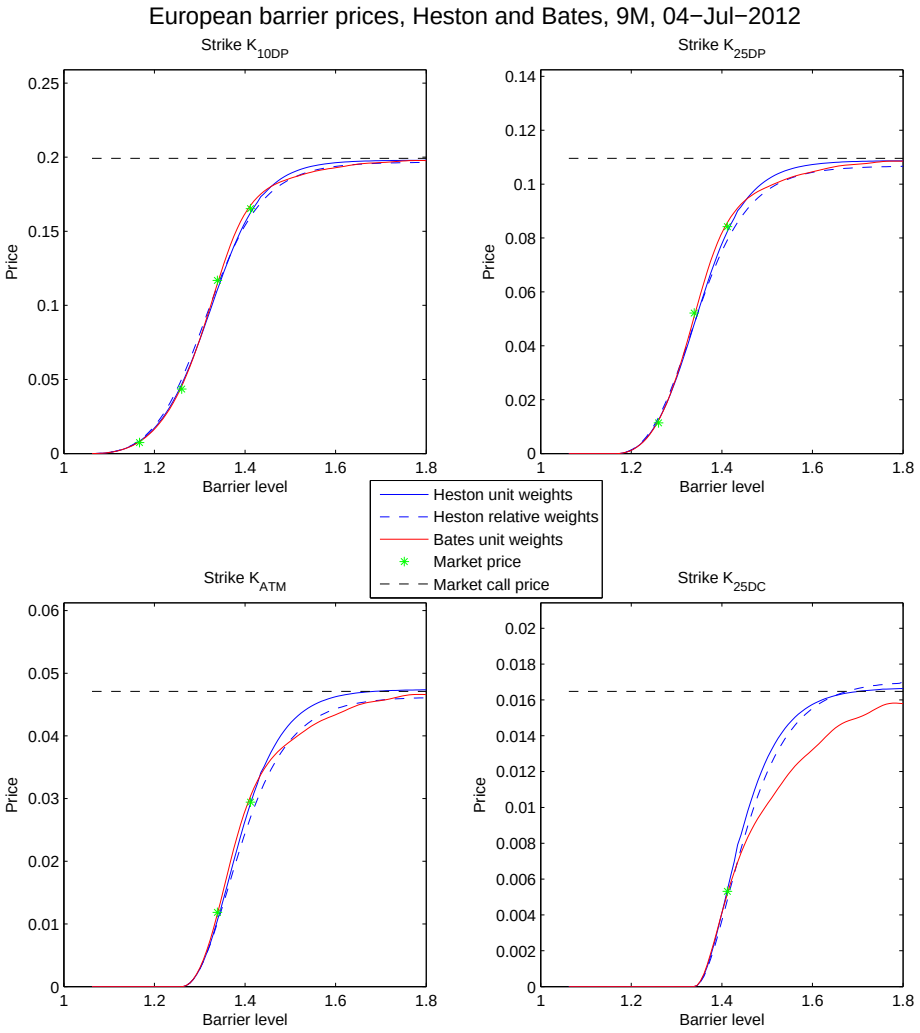

Figure 24: Prices for an up-and-out European call with different strike levels, Heston and Bates.

weights and Bates unit weights. The Barrier prices for Bates are overestimated compared to Heston for lower barriers, and vice versa for higher barriers. This qualitative conclusion seems to hold for any date, and not just specifically for this date. Studying figure 21 again one notices that Bates has a similar skew for low maturities but increases much in implied volatility for higher strikes compared to Heston. When replicating the European out contract we recall that we subtracted N call options and added $N-1$ call options. Hence, if the price of an OTM option is quite big we may subtract relatively more compared to Heston.

In figure 26 the pricing for all days using Black-Scholes is shown. The strike has here been chosen as $K = K_{ATM}$ and the barrier $U = K_{10DC}$. Note the distinct difference that unit weights overestimate in-contracts and relative weights

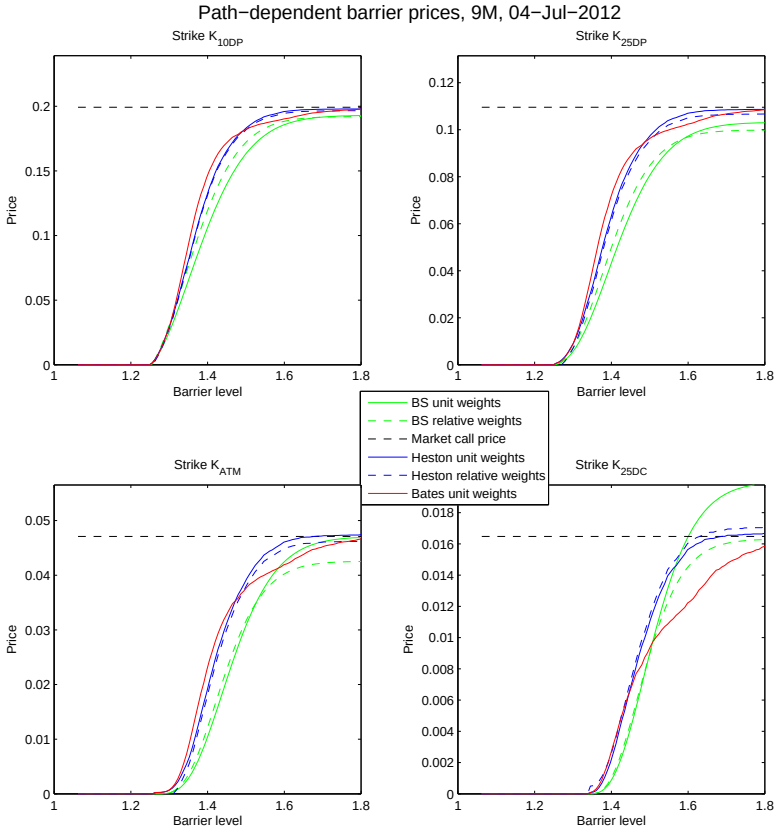

Figure 25: Price up-and-out contracts July 4 2012, 9M for different barrier levels.

overestimate out-contracts.

In figure 27 the pricing for all days using Heston and Bates is shown. Note again that the price differences between unit and relative weights are not as prominent as the difference between Heston and Bates. Heston overestimates in-contracts and Bates overestimates out-contracts. As for the weights, the relative weights overestimate the in-prices and the unit weights overestimate the out-prices.

In this simulation 20000 paths were generated (with a grid distance of one day) for Heston unit and relative weights and Bates unit weights. The standard deviation for the estimates were around $1 \cdot 10^{-4}$ for 2M and $2 \cdot 10^{-4}$ for 9M.

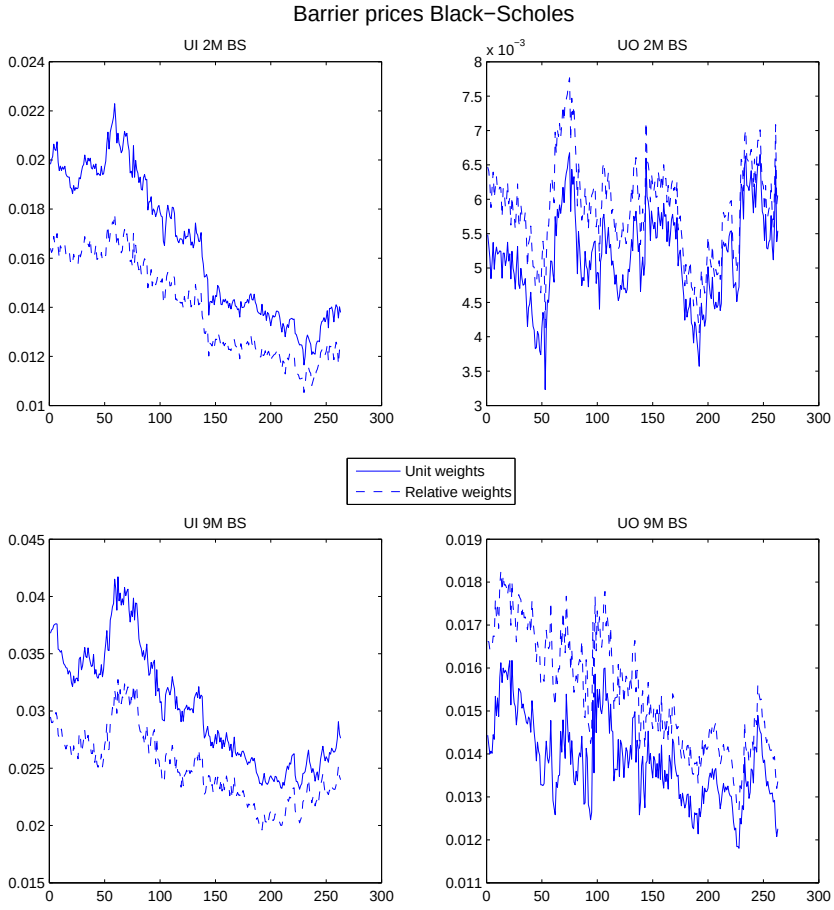

Figure 26: Black-Scholes barrier prices during February 2012 - February 2013.

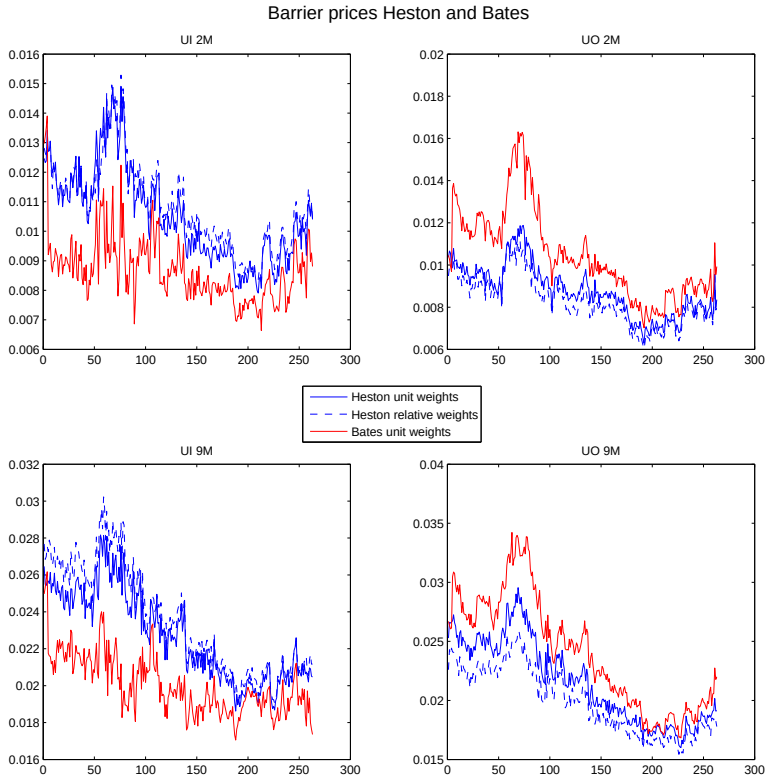

Figure 27: Heston and Bates barrier prices during February 2012 - February 2013.

13 Discussion

13.1 Conclusion

- The quotes in this thesis were ATM volatilities, risk-reversals and butterflies. These quotes and for these delta levels are not the only way to quote prices, (strangles are for instance sometimes quoted, see Wystrup (2010)) and one must know the exact quotes, ATM definitions and delta definitions in order to get things correct. The quintessence is that translating quotes on the FX market into prices for vanilla options is not as straight forward as one would first expect.

- It was clear - for instance when studying the skew in the surface - that Black-Scholes wasn't able to capture the behaviour of the exchange rate. For Heston we noted that the calibration using unit and relative weights were rather similar, with the difference that relative weights gave a more erratic parameter estimation. Also, the unit weights predicted a more negative correlation parameter ρ giving a more prominent skew in the volatility smile. One could note that relative weights took the OTM call options more into account compared to the unit weights. This could for instance be seen when studying the error in the implied volatility surface.

  The calibration for Bates showed that it was very hard to find a good calibration using relative weights. The calibration for unit weights were similar to that of Heston using unit weights.

- The market implied volatility surface was constructed by an interpolation using a second degree polynomial. There is not much litterature on the subject, and most likely financial institutions use different interpolation schemes. The problem of extrapolation was not mentioned in this thesis, but clearly this is an interesting topic as well. The interpolation in maturity dimension used so called flat forward interpolation, and not linear interpolation.

  What could be noted when comparing the implied volatility surfaces from the models with the market implied volatility was that for relative weights the error in volatility became more prominent for lower maturities compared to unit weights. The reverse could be seen for higher strikes. The error for Bates was relatively small and flat, and similar to Heston unit weights.

  When studying the skew implied by the market and models we noted that the skew was more prominent for Heston unit weights than relative weights. The behaviour for Bates was similar to that of Heston unit weights for lower maturities (higher moneyness) but differed quite a bit for higher strikes (lower moneyness).

- We noted that for Black-Scholes the European barriers were not priced very well, since the pricing boils down to pricing different contracts with quite different implied volatilities. Heston unit weights tended to overestimate prices compared to relative weights, and Bates unit weights tended to overestimate for lower barrier levels, and underestimate for higher barrier levels.

The prices for different barriers looked very similar when studying path-dependent contracts. Again unit weights tended to overestimate the out-contracts, and Bates overestimated for lower barrier levels and underestimated for higher barrier levels. When choosing the strike as the ATM level and the barrier level as the $10DC$ level we noted that the difference between Heston unit and relative weights were not as big as the difference between Heston unit and Bates unit weights.

13.2 Further develoment

- The weighting schemes in this thesis considered minimizing squared price differences with two weighting schemes. Another natural choice would be to consider the squared difference in implied volatilities with suitable weights. This approach is more time consuming (although approximations of the implied volatility have been proposed in the litterature), but is in a way more consistent with the way prices are quoted on the FX market.

- In 'reality' traders have both a bid and an ask price (quoted in volatilities). The calibration procedure can be much more dynamic if one incorporates the uncertainty in the price represented by a big ask-bid spread. For instance, if the spread is big that price is considered uncertain, and if the spread is narrow that price is considered certain. The bid-ask spread is often associated with liquidity. ATM options for instance are generally the most liquid options, resulting in a narrow ask-bid spread.

- The interpolation scheme proposed in this thesis - using second order polynomial in delta - can clearly be improved. Above all, in order to extrapolate one would need to use economic arguments to suggest boundaries to ensure no arbitrage opportunities.

- In this thesis least-square methods were used in the calibration procedure. Instead of this one could clearly have used some filtering techniques - an unscented Kalman filter for instance. This could give both faster calibration and a more smooth parameter development.

- It could be interesting to study another currency pair where the skew looks different from the EURUSD case. One could for instance study a currency pair where the correlation parameter ρ is positive instead of negative, and see how this changes the calibration.

- In this thesis the Black-Scholes, Heston and Bates models were considered; other models would be interesting to consider, the SABR model for instance.

  It is clearly a matter of purpose what model to be used. A trader will presumably for instance prefer a simple and robust model, while a quant may suggest a more advanced model. It is vital to understand many aspects of a model to evaluate whether to use it or not.

14 Appendix

Calibration result

Black-Scholes

		10DC	25DC	ATM	25DP	10DP
1M	Unit weights	-0.81044	-0.43426	-0.16609	-0.03904	-0.00387
	Relative weights	-0.38672	-0.22415	-0.07573	-0.00554	0.006313
2M	Unit weights	-0.75801	-0.41044	-0.14087	-0.02089	0.004925
	Relative weights	-0.34174	-0.20218	-0.05213	0.010876	0.013487
3M	Unit weights	-0.66502	-0.37256	-0.11413	-0.00472	0.011544
	Relative weights	-0.26479	-0.16787	-0.02712	0.02549	0.018842
6M	Unit weights	-0.40943	-0.25091	-0.05012	0.026433	0.021883
	Relative weights	-0.05516	-0.05831	0.033005	0.05395	0.027312
9M	Unit weights	-0.25168	-0.14751	-0.00572	0.046901	0.02765
	Relative weights	0.073965	0.03478	0.075012	0.072908	0.032147
12M	Unit weights	-0.0964	-0.08226	0.026367	0.059124	0.031179
	Relative weights	0.200046	0.093417	0.105606	0.084576	0.035248

Heston

		10DC	25DC	ATM	25DP	10DP
1M	Unit weights	0.043596	0.046493	-0.0044	-0.01224	-0.00376
	Relative weights	-0.03565	0.033151	0.005388	-0.00609	-0.00193
2M	Unit weights	0.127368	0.020504	-0.017	-0.0069	0.000341
	Relative weights	0.032974	-0.00215	-0.00803	-0.0004	0.002989
3M	Unit weights	0.148049	0.015661	-0.01322	-0.0028	0.001625
	Relative weights	0.044446	-0.01108	-0.00453	0.004789	0.004728
6M	Unit weights	0.149052	0.01173	-0.00699	0.001706	0.003528
	Relative weights	0.041132	-0.01322	0.004861	0.012338	0.007756
9M	Unit weights	0.111002	0.008129	-0.00949	0.004225	0.004814
	Relative weights	0.016646	-0.00729	0.007713	0.017584	0.010054
12M	Unit weights	0.110504	-0.01372	-0.01578	0.002724	0.005038
	Relative weights	0.034862	-0.01745	0.007075	0.018713	0.011241

Bates

	Unit weights				
	10DC	25DC	ATM	25DP	10DP
1M	-0.004573425	0.036697	-0.0029	-0.00299	-0.00038
2M	-0.053492134	0.027279	-0.00778	-0.00536	-0.00074
3M	-0.069419072	0.023433	-0.00635	-0.00387	0.000351
6M	-0.049344691	0.021101	-0.00247	-0.00099	0.001822
9M	-0.026782877	0.023621	-0.00397	0.001487	0.003001
12M	0.026465477	0.010081	-0.00853	0.000294	0.003205

Figure 28: The calibration results for Black-Scholes, Heston and Bates for all maturities, using average relative error for all dates.

Comparison error: 2M and 9M

Unit weights

	2M				
	10DC	25DC	ATM	25DP	10DP
Black-Scholes	-0.75801154	-0.41044	-0.14087	-0.02089	0.004924646
Heston	0.127367954	**0.020504**	-0.017	-0.0069	**0.000341106**
Bates	**-0.053492134**	0.027279	**-0.00778**	**-0.00536**	-0.000738794

Relative price weights

	2M				
	10DC	25DC	ATM	25DP	10DP
Black-Scholes	-0.34173552	-0.20218	-0.05213	0.010876	0.013486554
Heston	**0.03297399**	**-0.00215**	**-0.00803**	**-0.0004**	**0.002989219**

Unit weights

	9M				
	10DC	25DC	ATM	25DP	10DP
Black-Scholes	-0.251683096	-0.14751	-0.00572	0.046901	0.027649868
Heston	0.111001913	**0.008129**	-0.00949	0.004225	0.004814065
Bates	**-0.026782877**	0.023621	**-0.00397**	**0.001487**	**0.003000542**

Relative price weights

	9M				
	10DC	25DC	ATM	25DP	10DP
Black-Scholes	0.073965298	0.03478	0.075012	0.072908	0.032146615
Heston	**0.016646427**	**-0.00729**	**0.007713**	**0.017584**	**0.010053738**

Figure 29: The calibration results for Black-Scholes, Heston and Bates for 2M and 9M. The model where the absolute relative error is the smallest has been marked as bold.

References

Andersen L., 2007. *Efficient Simulation of the Heston Stochastic Volatility Model.* Is available at:
http://papers.ssrn.com/sol3/papers.cfm?abstract_id=946405

Andersen L.B.G, Jäckel P., Kahl C., 2010. *Simulation of Square-root processes.* Encyclopedia of Quantitative finance Wiley.

Björk T., 2009. *Arbitrage Thory in Continuous Time, third edition.* Oxford university press.

Clark I. J., 2012. *Foreign exchange option pricing, A practioner's Guide.* Wiley Finance.

Gilli M., Schumann E., 2010. *Calibration Option Pricing Models with Heuristics*. COMISEF Working papers series 30, 1-7.

Gut A., 2009. *An Intermediate Course in Probability*. Springer Second edition.

Haastrecht A., Pelsser, A.A.J., 2008. *Efficient, Almost Exact Simulation of the Heston Stochastic volatility model*. Netspar 44, 7-9.

Kleist et al (2010) *Report on global foreign exchange market activity in 2010, Bank for International Settlements* Is available at:
http://www.bis.org/publ/rpfxf10t.pdf

Moodley N., 2005. *The Heston Model: A Practical Approach with Matlab Code*. Is available at:
http://math.nyu.edu/ atm262/fall06/compmethods/a1/nimalinmoodley.pdf

Reiswich D. and Wystrup U., 2010. *FX Volatility Smile Contruction*. CPQF Working Paper series No. 20.

Schmelzle M., 2010. *Option Pricing Formulae using Fourier Transform: Theory and Application*. Is available at:
http://pfadintegral.com/docs/Schmelzle2010%20Fourier%20Pricing.pdf

Shamah S., 2003. *A Foreign exchange primer*. Wiley Finance.

Wilmott P., 2006. *On Quantitative Finance*. Johan Wiley & Sons, Ltd.

Wystrup U., 2006. *FX Options and Structured Products*, Wiley Finance

Åberg S., 2010. *Lecture Notes to FMS170/MASM19*. Centre for mathematical sciences Lund university.