Bloomberg
FINANCIAL MARKETS
COMMODITIES
NEWS

Implementation of the Heston Model for the Pricing of FX Options

Iddo Yekutieli

Bloomberg LP - Bloomberg Financial Markets (BFM)

June 22, 2004

1 Foreign Exchange Options Markets and the Heston Model

The Garman-Kohlhagen[1] model, which is the application of the Black-Scholes model to a Foreign Exchange Rate setting, has been used as the basic pricing and hedging tool for FX options for many years. And yet, it is well known to be incorrect. If we use the Black-Scholes implied volatility of the At-The-Money option to compute theoretical prices for away-from-the-money options, we find that these prices undervalue actual option prices. This artifact in known as the volatility smile: a plot of the Black-Scholes implied volatility as a function of option strike price, for a fixed expiry time, resembles a smile.

The Black-Scholes model is based on the assumption that the volatility of relative changes in the price of the underlying is constant. While this is inconsistent with actual market data, the model is still used as a language, or a representation tool, for options and their prices. FX options are usually quoted in terms of the Black-Scholes delta and the Black-Scholes implied volatility (instead of strike price and premium).

One way in which the Black-Scholes model can be generalized to account for the complexity seen in the market is by allowing the volatility of the underlying asset to vary in a random manner. This class of models is named Stochastic Volatility models. The dynamics of the underlying price in stochastic volatility models resembles that in the Black-Scholes model: the magnitude of diffusion term is proportional to the underlying price, but the volatility is no longer constant. It is a random variable itself. The Heston model falls under this category. The model assumes that the dynamics of the variance (square of the volatility) of the underlying asset undergoes random changes, with magnitude proportional to the square-root of the variance. In addition there is a mean-reverting drift term which assures that, asymptotically, the mean and the variance of the variance tend to finite limits. This dynamics is essentially that of the Cox-Ingersoll-Ross model[3].

The main advantage of the Heston model, relative to other stochastic volatility models is that it permits closed-form solutions for European option prices, while guaranteeing non-negative exchange rates. The Heston model is capable of generating the smirks and smiles seen in implied volatility curves, for option expiries ranging from several months upward. The model's main weakness lies in its incapacity to generate smirks and smiles for short term options. This is a characteristic of all stochastic volatility models, and will be discussed below.

Hull and White [6] have shown that a model of the underlying price process with stochastic volatility will generate an implied volatility curve with a smile (curvature as a function of the strike price). This results holds if there is no correlation between movements of the underlying and of volatility. The option price can be computed as the Black-Scholes option price, averaged over the distribution of values of volatility. This result indicates that a multiplicity of values of volatility can generate the fat tails of probability distributions of the Risk-Neutral Measure (RNM) that are implied by market option prices.

Willard [7] and Lewis [8] have shown that the addition of correlation in the co-movements of the underlying and of volatility will bring about a skew in the implied volatility curve. Their approaches also give a method for pricing derivatives with general payoffs paid at a single date, by performing a Monte-Carlo path-wise average of Black-Scholes prices.

In the following sections we describe how we have implemented the Heston model for the pricing of Foreign Exchange options. We begin by detailing the mathematical formulation of the Heston model. We then describe how we price European options with the Heston model. In Section 4 we show how American options are priced. Section 5 details the methodology of calibrating the Heston model parameters to market implied volatility surfaces.

2 Mathematical Formulation of the Heston Model

The stochastic process of the Heston model, in the RNM, is specified by coupled Stochastic Differential Equations (SDEs) for the spot price $S(t)$ and the variance (square of the volatility) of the spot price $v(t)$

$$dS = (r_d - r_f)S\,dt + \sqrt{v}S\,dW_1 \tag{1}$$

$$dv = \kappa(\theta - v)\,dt + \xi\sqrt{v}\,dW_2 \tag{2}$$

r_d and r_f are the risk-free interest rates in the domestic and foreign markets, respectively. κ is the mean-reversion rate, controlling the speed at which the variance moves towards the long-run variance θ. We assume that the market price of volatility risk is integrated into the parameters κ and θ. The two driving Wiener processes $dW_1(t)$ and $dW_2(t)$ are correlated, $\mathrm{E}[dW_1 dW_2] = \rho dt$. The initial conditions are $S(0) = S_0$ and $v(0) = v_0$. The initial variance v_0 is not directly observable.

Equation (1) resembles the SDE of the Black-Scholes model, except that the constant volatility σ of Black-Scholes is replaced in the Heston model by a time-dependent stochastic variable $\sqrt{v(t)}$.

Equation (2) on its own is precisely the SDE of the Cox-Ingersoll-Ross model[3]. This process ensures that the variance remains positive, provided that the model parameters obey the relation $4\kappa\theta > \xi^2$. This process allows an analytical form for the probability distribution of the future variance. This in turn permits obtaining a closed-form solution for the price of a European option.

The characteristic function for the Heston model, $H(z,T) = \mathrm{E}_{\mathrm{Q}}\left[\mathrm{e}^{izS(T)}\right]$, with expectation taken over the Risk-Neutral Measure (RNM), is given by

$$H(z,T) = \mathrm{e}^{iz\left[\ln S_0 + (r_d - r_f)T\right] + Dv_0 + C} \tag{3}$$

where

$$\begin{aligned} b &= \kappa - iz\rho\xi \\ d &= \sqrt{b^2 + z(z+i)\xi^2} \\ g &= \frac{b+d}{b-d} \\ D &= \frac{b+d}{\xi^2}\frac{1-\mathrm{e}^{dT}}{1-g\mathrm{e}^{dT}} \\ C &= \frac{\kappa\theta}{\xi^2}\left[(b+d)T - 2\ln\frac{1-g\mathrm{e}^{dT}}{1-g}\right] \end{aligned}$$

The Heston model Partial Differential Equation (PDE) for the price of a derivative security $P(S,v,t)$ is

$$\begin{aligned} \frac{\partial P}{\partial t} &+ \frac{1}{2}vS^2\frac{\partial^2 P}{\partial S^2} + \rho\xi Sv\frac{\partial^2 P}{\partial S\partial v} + \frac{1}{2}\xi^2 v\frac{\partial^2 P}{\partial v^2} \\ &+ (r_d - r_f)S\frac{\partial P}{\partial S} + \kappa(\theta - v)\frac{\partial P}{\partial v} - r_d P = 0 \end{aligned} \tag{4}$$

This equation is assumed to describe prices in the RNM, its parameters already incorporating the market price of volatility risk.

The Heston model is a generalization of the Black-Scholes model. Taking the limits $\xi \to 0$ and $\kappa \to 0$, we recover the Black-Scholes PDE, with $\sigma = \sqrt{v_0}$.

3 Pricing European Options

Given the characteristic function for the Heston model in Equation (3), the price of a European option can be expressed using the formulas

$$\Pi_1 = \frac{1}{2} + \frac{\phi}{\pi}\int_0^\infty \mathcal{R}\left[\frac{\mathrm{e}^{-iu\ln K}H(u-i,T)}{iuH(-i,T)}\right]du \tag{5}$$

$$\Pi_2 = \frac{1}{2} + \frac{\phi}{\pi}\int_0^\infty \mathcal{R}\left[\frac{\mathrm{e}^{-iu\ln K}H(u,T)}{iu}\right]du \tag{6}$$

$$P = \phi\left[\mathrm{e}^{-r_f t}S_0\Pi_1 - \mathrm{e}^{-r_d t}K\Pi_2\right] \tag{7}$$

where $\phi = 1(-1)$ for a call (put). This is the closed-form solution of Heston [2] (see also [4] and [5]). Equation (7) is a generalization of the Black-Scholes option pricing formula. Equations (5-7) can be used whenever the characteristic function of the probability distribution of the future underlying price for a given model is known.

The numerical integration of Equations (5-6) is performed in our implementation using an adaptive quadrature. With this method we can price options rapidly with a level of 10^{-6} relative accuracy.

If the strike price is very far from the money, such that the option delta is close to 0 or to 1 (within 10^{-5}), this method fails to compute the option price accurately under the Heston model. In such cases we revert to pricing the option using 0% volatility of volatility. This corresponds to a Black-Scholes model with time-dependent (deterministic) volatility. There is little loss of accuracy in this case, as the option price is very close to its intrinsic value.

4 Pricing American Options

To price American options with the Heston model, we use an implicit finite difference scheme on a regular two-dimensional lattice to solve the PDE (4). We found that it is advantageous to perform a double change of variables. We replace the spot price S by its logarithm $x = \log S$, and the variance v by volatility $y = \sqrt{v}$ (we distinguish between $y(t)$ and the Black-Scholes constant volatility σ)

$$\begin{aligned}\frac{\partial P}{\partial t} &+ \frac{1}{2}y^2\frac{\partial^2 P}{\partial x^2} + \frac{\rho\xi}{2}\frac{1}{y}\frac{\partial^2 P}{\partial x \partial y} + \frac{\xi^2}{8}\frac{\partial^2 P}{\partial y^2} \\ &+ \left(r_d - r_f - \frac{y^2}{2}\right)\frac{\partial P}{\partial x} + \frac{1}{2}\left(\kappa(\theta - y^2) - \frac{\xi^2}{4}\right)\frac{1}{y}\frac{\partial P}{\partial y} - r_d P = 0 \quad (8)\end{aligned}$$

With this representation there is a finer resolution in volatility space of values $y < 1$, which is the typical region of interest. In addition, the boundary condition at $y = 0$ becomes $\frac{\partial P}{\partial y} = 0$, which is easy to enforce accurately. For boundary conditions at the three other boundaries, $x = x_{\min}, x = x_{\max}, y = y_{\max}$, we use the option price computed with $\xi = 0$: this is a Black-Scholes price with an effective volatility σ_{eff} [1]. This is a relatively good approximation of the actual option price at these boundaries. We choose the size of the region covered by the lattice, and the increment sizes, in a way to ensure that the effect of the boundaries does not propagate to the region of interest for the pricing of an option.

Time stepping is done with one of two schemes: an Alternating Direction Integration scheme when the model parameters permit it, and a Stabilizing Correction scheme otherwise. The former scheme converges more quickly, while the latter is more stable. The decision on which scheme to use depends on the various model parameters, but is mostly affected by the volatility of volatility ξ and the correlation ρ. This assures us, together with an optimal choice of grid size and number of time steps, good convergence and acurate option prices.

The accuracy of option prices depends on the Heston model parameters and time to expiry; it does not depend much on the option strike price. For reasonable values of the model parameters, option prices are assured to have an accuracy of 0.2%, in relative terms. For extreme model parameter values, especially if $|\rho|$ is close to 1, or if volatility of volatility ξ is high (80% or more), accuracy may degrade.

5 Calibration of Heston Model Parameters

The FX options market consists essentially of OTC options, as opposed to equity options, for example, which are mainly exchange traded. Accordingly, FX option prices are not quoted on specific contracts, with sets of fixed expiry dates and strike prices. Instead, quotes are given in terms of fixed times to expiry (1 week, 1 month, 2 months, etc.) and fixed values of option delta (At-The-Money, 25 delta and 10 delta). Given the Black-Scholes option delta and implied volatility, one can deduce its strike price and premium.

The At-The-Money (ATM) implied volatility quote sets the general level

[1] In the limit $\xi \to 0$, the Heston option price equals the Black-Scholes price computed with a volatility σ_{eff}

$$(\sigma_{\text{eff}})^2 = \frac{1}{T-t}\int_t^T y^2(s)ds = \frac{1-\mathrm{e}^{-\kappa\tau}}{\kappa\tau}y^2(t) + \left(1 - \frac{1-\mathrm{e}^{-\kappa\tau}}{\kappa\tau}\right)\theta \text{ , where } \tau = T-t$$

Even for $\xi > 0$, this is seen in practice to be a good Dirichlet boundary condition.

of volatilities for a specific time-to-expiry. The out-of-the-money data is not quoted in terms of single options (calls or puts), but rather in terms of specific strategies. The 25-delta Risk-Reversal is a quote of the spread between the 25-delta call and the 25-delta put implied volatilities. It is indicative of the slope of the volatility curve. The 25-delta Butterfly (or Straddle) is the spread between the average of the 25-delta call and the 25-delta put implied volatilities, and the ATM implied volatility. It indicates the curvature of the volatility curve. These can be written as equations representing implied volatilities

$$\begin{aligned} \mathrm{RR}_{25} &= \mathrm{Call}_{25} - \mathrm{Put}_{25} \\ \mathrm{BF}_{25} &= \frac{1}{2}\left[\mathrm{Call}_{25} + \mathrm{Put}_{25}\right] - \mathrm{ATM} \end{aligned}$$

The same convention is used for 10-delta quotes.

From the quotes of Black-Scholes implied volatilities, we generate a set of specific option contracts. These are then used as inputs to a calibration routine which finds the best-fit values of the Heston model parameters. The cost function for the fit is the sum of squares of errors between market quoted implied volatilities and the model implied values.

There are three main problems when trying to fit the Heston model to an actual volatility surface:

- Being a 5-parameter model, the Heston model cannot exactly fit a general volatility surface, which may contain 20-30 data points.
- The Heston model generates skew and curvature dynamically, from the stochastic nature of volatility. Over short horizons it does not generate the amount of skew and curvature quoted in actual volatility surfaces [2].
- The time dependence of implied volatilities of at-the-money options is not in general of the form of an exponential transition from short-term to long-term levels as prescribed by the form of mean-reversion of the Heston model.

In view of these limitations, we exclude options with expiry times up to 2 months in the calibration set. We also exclude 10-delta (put and call) options, limiting the dataset to ATM and 25-delta options.

[2] This shortcoming of stochastic volatility models can be alleviated by adding jump components to the model. Moreover, it is taken as evidence to the necessity of using a model with jumps.

The following figures detail statistics of the goodness-of-fit of the Heston model to actual volatility surfaces. The input data contains volatility surfaces at 28 dates for each of four currency pairs: EUR/JPY, EUR/USD, GBP/USD and USD/JPY. We calibrate a set of Heston model parameters to each of the surfaces. We then measure the relative error of the volatility implied by the Heston model option price, and the input implied volatility. This is done for ATM, 25-delta and 10-delta options, with expiries of 1 week, 1 month, 6 months and 1 year. Figure 1 shows mean errors, and Figure 2 shows the root-mean-square of the errors . For 6 month and 1 year options the errors are similar, with mean and rms at around 1-2%. Errors are smaller for ATM and 25-delta options, increasing somewhat for the 10-delta options. For 1 week and 1 month options the fit is not very good, as is expected for the Heston model. Errors are around 5-10% for all delta values.

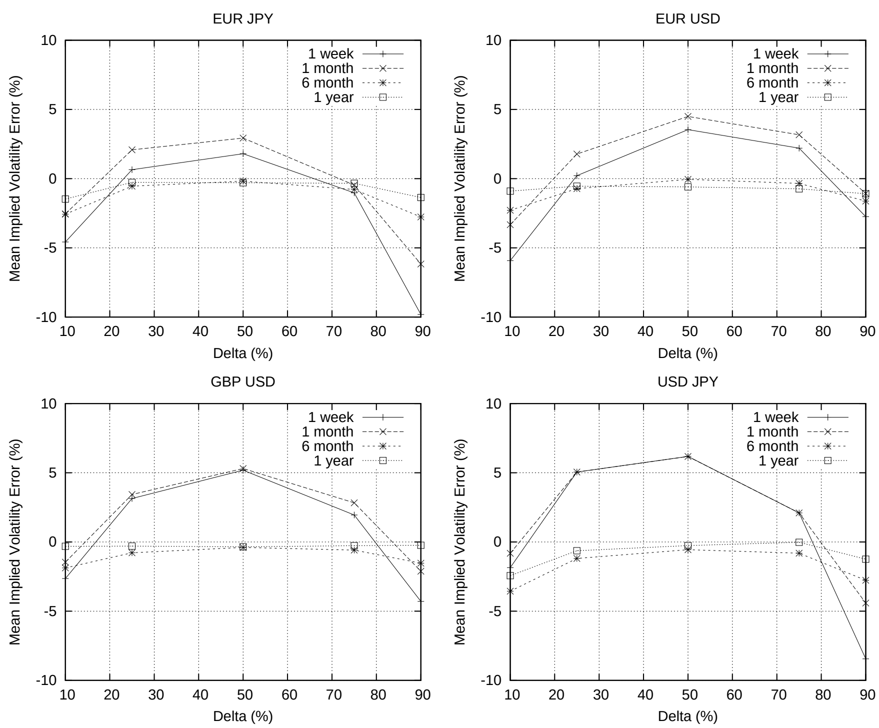

Figure 1: Mean of the relative errors of fit of the Heston model to real implied volatility surfaces. Datasets contain surfaces from 28 dates for each of four currency pairs. These are relative errors in implied volatility, in percent. We measure goodness-of-fit for ATM, 25-delta and 10-delta options, with expiries of 1 week, 1 month, 6 months and 1 year.

The data used in the plots above results from a calibration of all Heston

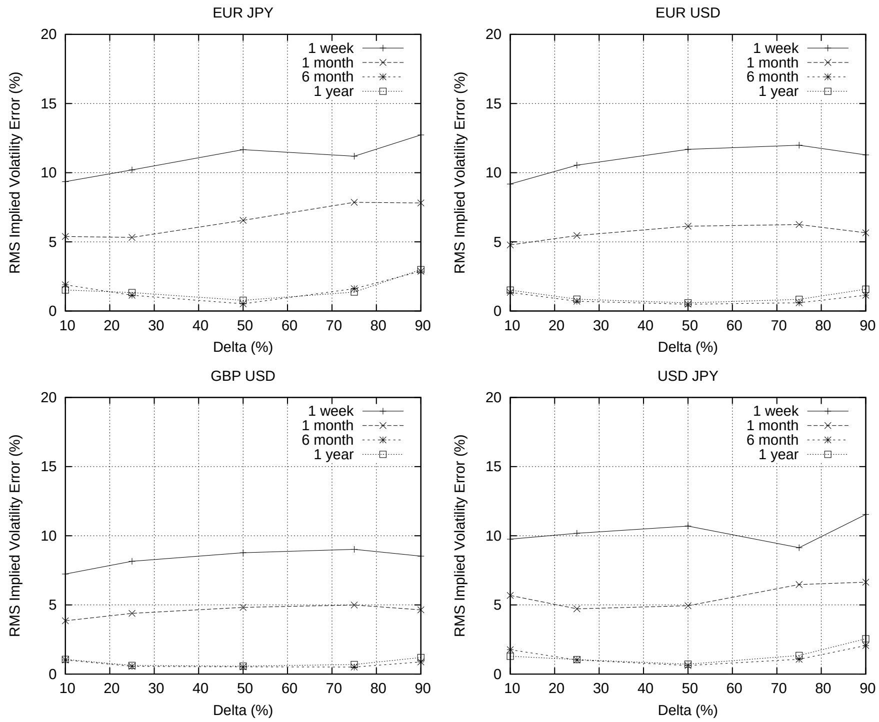

Figure 2: Root mean square of the relative errors of fit of the Heston model to real implied volatility surfaces. Datasets contain surfaces from 28 dates for each of four currency pairs. These are relative errors in implied volatility, in percent. We measure goodness-of-fit for ATM, 25-delta and 10-delta options, with expiries of 1 week, 1 month, 6 months and 1 year.

model parameters, except for the mean reversion parameter κ, which is fixed at a typical value of 1. The reason for this is that there is strong interplay in the dynamics of the volatility between mean reversion and volatility of volatility. Volatility of volatility tends to widen the distribution of future volatility values, while mean reversion shrinks this distribution. A high value of mean reversion can be compensated for by a high value of volatility of volatility. Due to this, there is a certain degeneracy in the determination of the best-fit values of these two parameters. Since in general, the mean reversion parameter is harder to pinpoint in model parameter estimation, we have decided to clamp its value in this test.

References

[1] M.B Garman and S.W. Kohlhagen, "Foreign Currency Option Values," Journal of International Money and Finance, Vol. 2 (1983).

[2] S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies, Vol. 6, No. 2, 1993.

[3] J. C. Cox, J. E. Ingersoll and S. A. Ross, "A Theory of the Term Structure of Interest Rates," Econometrica, Vol. 53, No. 2, March 1985.

[4] Foreign exchange risk: models, instruments and strategies. Ed. J. Hakala, U. Wystup. em Risk Books, 2002.

[5] P. Carr and D. Madan, "Option Valuation using the Fast Fourier Transform," Journal of Computational Finance, Vol. 2 No. 4, Summer 1999.

[6] J. C. Hull and A. White, "The Pricing of Options on Assets with Stochastic Volatilities," Journal of Finance, Vol. 42 (1987).

[7] G. A. Willard, "Calculating Prices and Sensitivities for Path-Independent Derivative Securities in Multifactor Models," The Journal of Derivatives, Vol. 5, No. 1, Fall 1997.

[8] A. L. Lewis "Option Valuation Under Stochastic Volatility with Mathematica Code," Finance Press, 2000.

Help x2

Press the <HELP> key twice for instant live assistance.

bloomberg.com

Frankfurt
+49 69 92041 0

Hong Kong
+852 2977 6000

London
+44 20 7330 7500

New York
+1 212 318 2000

San Francisco
+1 415 912 2960

São Paulo
+55 11 3048 4500

Singapore
+65 6212 1000

Sydney
+61 2 9777 8600

Tokyo
+81 3 3201 8900