



Munich Personal RePEc Archive

FX Smile in the Heston Model

Janek, Agnieszka and Kluge, Tino and Weron, Rafal and
Wystup, Uwe

Wrocław University of Technology, Poland, MathFinance AG,
Germany

September 2010

Online at <https://mpra.ub.uni-muenchen.de/25491/>
MPRA Paper No. 25491, posted 28 Sep 2010 20:17 UTC

FX Smile in the Heston Model¹

Agnieszka Janek², Tino Kluge³, Rafał Weron⁴, and Uwe Wystup³

Abstract: The Heston model stands out from the class of stochastic volatility (SV) models mainly for two reasons. Firstly, the process for the volatility is non-negative and mean-reverting, which is what we observe in the markets. Secondly, there exists a fast and easily implemented semi-analytical solution for European options. In this article we adapt the original work of Heston (1993) to a foreign exchange (FX) setting. We discuss the computational aspects of using the semi-analytical formulas, performing Monte Carlo simulations, checking the Feller condition, and option pricing with FFT. In an empirical study we show that the smile of vanilla options can be reproduced by suitably calibrating three out of five model parameters.

Keywords: Heston model; vanilla option; stochastic volatility; Monte Carlo simulation; Feller condition; option pricing with FFT

JEL: C5, C63, G13

¹ Chapter prepared for the 2nd edition of *Statistical Tools for Finance and Insurance*, P.Cizek, W.Härdle, R.Weron (eds.), Springer-Verlag, *forthcoming in 2011*.

² Institute of Mathematics and Computer Science, Wrocław University of Technology, Poland

³ MathFinance AG, Waldems, Germany

⁴ Institute of Organization and Management, Wrocław University of Technology, Poland

1 FX Smile in the Heston Model

Agnieszka Janek, Tino Kluge, Rafał Weron, and Uwe Wystup

1.1 Introduction

The universal benchmark for option pricing is flawed. The Black-Scholes formula is based on the assumption of a geometric Brownian motion (GBM) dynamics with constant volatility. Yet, the model-implied volatilities for different strikes and maturities of options are not constant and tend to be smile shaped (or in some markets skewed). Over the last three decades researchers have tried to find extensions of the model in order to explain this empirical fact.

As suggested already by Merton (1973), the volatility can be made a deterministic function of time. While this approach explains the different implied volatility levels for different times of maturity, it still does not explain the smile shape for different strikes. Dupire (1994), Derman and Kani (1994), and Rubinstein (1994) came up with the idea of allowing not only time, but also state dependence of the volatility coefficient, for a concise review see e.g. Fengler (2005). This local (deterministic) volatility approach yields a complete market model. It lets the local volatility surface to be fitted, but it cannot explain the persistent smile shape which does not vanish as time passes. Moreover, exotics cannot be satisfactorily priced in this model.

The next step beyond the local volatility approach was to allow the volatility to be driven by a stochastic process. The pioneering work of Hull and White (1987), Stein and Stein (1991), and Heston (1993) led to the development of stochastic volatility (SV) models, for reviews see Fouque, Papanicolaou, and Sircar (2000) and Gatheral (2006). These are multi-factor models with one of the factors being responsible for the dynamics of the volatility coefficient. Different driving mechanisms for the volatility process have been proposed, including GBM and mean-reverting Ornstein-Uhlenbeck type processes.

The Heston model stands out from this class mainly for two reasons. Firstly, the process for the volatility is non-negative and mean-reverting, which is what we observe in the markets. Secondly, there exists a fast and easily implemented semi-analytical solution for European options. This computational efficiency becomes critical when calibrating the model to market prices and is the greatest advantage of the model over other (potentially more realistic) SV models. Its popularity also stems from the fact that it was one of the first models able to explain the smile and simultaneously allow a front-office implementation and a valuation of many exotics with values closer to the market than the Black-Scholes model. Finally, given that all SV models generate roughly the same shape of implied volatility surface and have roughly the same implications for the valuation of exotic derivatives (Gatheral, 2006), focusing on the Heston model is not a limitation, rather a good starting point.

This chapter is structured as follows. In Section 1.2 we define the model and discuss its properties, including marginal distributions and tail behavior. Next, in Section 1.3 we adapt the original work of Heston (1993) to a foreign exchange (FX) setting. We do this because the model is particularly useful in explaining the volatility smile found in FX markets; in equity markets the typical volatility structure is a strongly asymmetric skew (also called a smirk or grimace). In Section 1.4 we show that the smile of vanilla options can be reproduced by suitably calibrating the model parameters. Finally, in Section 1.5 we briefly discuss the alternatives to the Heston model.

1.2 The Model

Following Heston (1993) consider a stochastic volatility model with GBM-like dynamics for the spot price:

$$dS_t = S_t \left(\mu dt + \sqrt{v_t} dW_t^{(1)} \right) \quad (1.1)$$

and a non-constant instantaneous variance v_t driven by a mean-reverting square root (or CIR) process:

$$dv_t = \kappa(\theta - v_t) dt + \sigma \sqrt{v_t} dW_t^{(2)}. \quad (1.2)$$

The stochastic increments of the two processes are correlated with parameter ρ , i.e. $dW_t^{(1)} dW_t^{(2)} = \rho dt$. The remaining parameters – μ , θ , κ , and σ – can be interpreted as the drift, the long-run variance, the rate of mean reversion





Figure 1.1: Sample trajectories (*left panel*) and volatilities (*right panel*) of the GBM and the Heston spot price process (1.1) obtained with the same set of random numbers.

 STF2hes01.m

to the long-run variance, and the volatility of variance (often called the *vol of vol*), respectively. Sample paths and volatilities of the GBM and the Heston spot price process are plotted in Figure 1.1.

By setting $x_t = \log(S_t/S_0) - \mu t$, we can express the Heston model (1.1)-(1.2) in terms of the centered (log-)return x_t and v_t . The process is then characterized by the transition probability $P_t(x, v | v_0)$ to have (log-)return x and variance v at time t given the initial return $x = 0$ and variance v_0 at time $t = 0$. The time evolution of $P_t(x, v | v_0)$ is governed by the following Fokker-Planck (or forward Kolmogorov) equation:

$$\begin{aligned} \frac{\partial}{\partial t} P &= \kappa \frac{\partial}{\partial v} \{ (v - \theta) P \} + \frac{1}{2} \frac{\partial}{\partial x} (v P) + \\ &+ \rho \sigma \frac{\partial^2}{\partial x \partial v} (v P) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (v P) + \frac{\sigma^2}{2} \frac{\partial^2}{\partial v^2} (v P). \end{aligned} \quad (1.3)$$

Solving this equation yields the following semi-analytical formula for the density of centered returns x , given a time lag t of the price changes (Dragulescu and Yakovenko, 2002):

$$P_t(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} e^{i\xi x + F_t(\xi)} d\xi, \quad (1.4)$$

with

$$F_t(\xi) = \frac{\kappa\theta}{\sigma^2} \gamma t - \frac{2\kappa\theta}{\sigma^2} \log \left(\cosh \frac{\Omega t}{2} + \frac{\Omega^2 - \gamma^2 + 2\kappa\gamma}{2\kappa\Omega} \sinh \frac{\Omega t}{2} \right),$$

$$\gamma = \kappa + i\rho\sigma\xi, \quad \text{and} \quad \Omega = \sqrt{\gamma^2 + \sigma^2(\xi^2 - i\xi)}.$$

Somewhat surprisingly, the introduction of SV does not change the properties of the spot price process in a way that could be noticed just by a visual inspection of its realizations, see Figure 1.1 where sample paths of a GBM and the spot process (1.1) in the Heston model are plotted. To make the comparison more objective both trajectories were obtained with the same set of random numbers.

In both cases the initial spot rate $S_0 = 4$ and the domestic and foreign interest rates are $r_d = 5\%$ and $r_f = 3\%$, respectively, yielding a drift of $\mu = r_d - r_f = 2\%$. The volatility in the GBM is constant $\sqrt{v_t} = \sqrt{4\%} = 20\%$, while in the Heston model it is driven by the mean-reverting process (1.2) with the initial variance $v_0 = 4\%$, the long-run variance $\theta = 4\%$, the speed of mean reversion $\kappa = 2$, and the vol of vol $\sigma = 30\%$. The correlation is set to $\rho = -0.05$.

A closer inspection of the Heston model does, however, reveal some important differences with respect to GBM. For instance, the probability density functions (pdfs) of (log-)returns have heavier tails – exponential compared to Gaussian, see Figure 1.2. In this respect they are similar to hyperbolic distributions (Weron, 2004, see also Chapter ??), i.e. in the log-linear scale they resemble hyperbolas, rather than parabolas of the Gaussian pdfs.

1.3 Option Pricing

Consider the value function of a general contingent claim $U(t, v, S)$ paying $g(S) = U(T, v, S)$ at time T . We want to replicate it with a self-financing portfolio. Due to the fact that in the Heston model we have two sources of uncertainty (the Wiener processes $W^{(1)}$ and $W^{(2)}$) the portfolio must include the possibility to trade in the money market, the underlying and another derivative security with value function $V(t, v, S)$.

We start with an initial wealth X_0 which evolves according to:

$$dX = \Delta dS + \Gamma dV + r_d(X - \Gamma V) dt - (r_d - r_f)\Delta S dt, \quad (1.5)$$

where Δ is the number of units of the underlying held at time t and Γ is the number of derivative securities V held at time t .





Figure 1.2: The marginal pdfs in the Black-Scholes (GBM) and Heston models for the same set of parameters as in Figure 1.1 (*left panel*). The tails of the Heston marginal pdfs are exponential, which is clearly visible in the log-linear scale (*right panel*).

 STF2hes02.m

The goal is to find Δ and Γ such that $X_t = U(t, v_t, S_t)$ for all $t \in [0, T]$. The standard approach to achieve this is to compare the differentials of U and X obtained via Itô's formula. After some algebra we arrive at the partial differential equation (PDE) which U must satisfy (for details on the derivation in the foreign exchange setting see Hakala and Wystup, 2002):

$$\begin{aligned} \frac{1}{2}vS^2\frac{\partial^2 U}{\partial S^2} + \rho\sigma vS\frac{\partial^2 U}{\partial S\partial v} + \frac{1}{2}\sigma^2 v\frac{\partial^2 U}{\partial v^2} + (r_d - r_f)S\frac{\partial U}{\partial S} + \\ + \left\{ \kappa(\theta - v) - \lambda(t, v, S) \right\} \frac{\partial U}{\partial v} - r_d U + \frac{\partial U}{\partial t} = 0. \end{aligned} \quad (1.6)$$

The term $\lambda(t, v, S)$ is called the *market price of volatility risk*. Heston (1993) assumed it to be linear in the instantaneous variance v_t , i.e. $\lambda(t, v, S) = \lambda v_t$, in order to retain the form of the equation under the transformation from the statistical (or risky) measure to the risk-neutral measure.

1.3.1 European Vanilla FX Option Prices and Greeks

We can solve (1.6) by specifying appropriate boundary conditions. For a European vanilla FX option these are:

$$U(T, v, S) = \max\{\phi(S - K), 0\}, \quad (1.7)$$

$$U(t, v, 0) = \frac{1 - \phi}{2} K e^{-r_d \tau}, \quad (1.8)$$

$$\frac{\partial U}{\partial S}(t, v, \infty) = \frac{1 + \phi}{2} e^{-r_f \tau}, \quad (1.9)$$

$$\begin{aligned} r_d U(t, 0, S) &= (r_d - r_f) S \frac{\partial U}{\partial S}(t, 0, S) + \\ &+ \kappa \theta \frac{\partial U}{\partial v}(t, 0, S) + \frac{\partial U}{\partial t}(t, 0, S), \end{aligned} \quad (1.10)$$

$$U(t, \infty, S) = \begin{cases} S e^{-r_f \tau}, & \text{for } \phi = +1, \\ K e^{-r_d \tau}, & \text{for } \phi = -1, \end{cases} \quad (1.11)$$

where $\phi = \pm 1$ for call and put options, respectively. The strike K is in units of the domestic currency and $\tau = T - t$ is the time to maturity (i.e. T is the expiration time in years and t is the current time).

Heston (1993) solved the PDE analytically using the method of characteristic functions. For European vanilla FX options the price is given by:

$$\begin{aligned} h(\tau) &= \text{HestonVanilla}(\kappa, \theta, \sigma, \rho, \lambda, r_d, r_f, v_t, S_t, K, \tau, \phi) \\ &= \phi \{ e^{-r_f \tau} S_t P_+(\phi) - K e^{-r_d \tau} P_-(\phi) \}, \end{aligned} \quad (1.12)$$

where $u_{1,2} = \pm \frac{1}{2}$, $b_1 = \kappa + \lambda - \sigma \rho$, $b_2 = \kappa + \lambda$, and

$$d_j = \sqrt{(\rho \sigma \varphi i - b_j)^2 - \sigma^2 (2u_j \varphi i - \varphi^2)}, \quad (1.13)$$

$$g_j = \frac{b_j - \rho \sigma \varphi i + d_j}{b_j - \rho \sigma \varphi i - d_j}, \quad (1.14)$$

$$\begin{aligned} C_j(\tau, \varphi) &= (r_d - r_f) \varphi i \tau + \\ &+ \frac{\kappa \theta}{\sigma^2} \left\{ (b_j - \rho \sigma \varphi i + d_j) \tau - 2 \log \left(\frac{1 - g_j e^{d_j \tau}}{1 - g_j} \right) \right\}, \end{aligned} \quad (1.15)$$

$$D_j(\tau, \varphi) = \frac{b_j - \rho \sigma \varphi i + d_j}{\sigma^2} \left(\frac{1 - e^{d_j \tau}}{1 - g_j e^{d_j \tau}} \right), \quad (1.16)$$

$$f_j(x, v_t, \tau, \varphi) = \exp\{C_j(\tau, \varphi) + D_j(\tau, \varphi)v_t + i\varphi x\}, \quad (1.17)$$

$$P_j(x, v_t, \tau, y) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left\{ \frac{e^{-i\varphi y} f_j(x, v_t, \tau, \varphi)}{i\varphi} \right\} d\varphi. \quad (1.18)$$

Note that the functions P_j are the cumulative distribution functions (in the variable $y = \log K$) of the log-spot price after time $\tau = T - t$ starting at $x = \log S_t$ for some drift μ . Finally:

$$P_+(\phi) = \frac{1 - \phi}{2} + \phi P_1(x, v_t, \tau, y), \quad (1.19)$$

$$P_-(\phi) = \frac{1 - \phi}{2} + \phi P_2(x, v_t, \tau, y). \quad (1.20)$$

The Greeks can be evaluated by taking the appropriate derivatives or by exploiting homogeneity properties of financial markets (Reiss and Wystup, 2001). In the Heston model the *spot delta* and the so-called *dual delta* are given by:

$$\Delta = \frac{\partial h(t)}{\partial S_t} = \phi e^{-r_f \tau} P_+(\phi) \quad \text{and} \quad \frac{\partial h(t)}{\partial K} = -\phi e^{-r_d \tau} P_-(\phi), \quad (1.21)$$

respectively. *Gamma*, which measures the sensitivity of delta to the underlying has the form:

$$\Gamma = \frac{\partial \Delta}{\partial S_t} = \frac{e^{-r_f \tau}}{S_t} p_1(\log S_t, v_t, \tau, \log K), \quad (1.22)$$

where

$$p_j(x, v, \tau, y) = \frac{1}{\pi} \int_0^\infty \Re \{ e^{-i\varphi y} f_j(x, v, \tau, \varphi) \} d\varphi, \quad j = 1, 2, \quad (1.23)$$

are the densities corresponding to the cumulative distribution functions P_j (1.18).

The time sensitivity parameter *theta* = $\partial h(t)/\partial t$ can be easily computed from (1.6), while the formulas for *rho* are the following:

$$\frac{\partial h(t)}{\partial r_d} = \phi K e^{-r_d \tau} \tau P_-(\phi), \quad \frac{\partial h(t)}{\partial r_f} = -\phi S_t e^{-r_f \tau} \tau P_+(\phi). \quad (1.24)$$

Note, that in the foreign exchange setting there are two *rho*'s – one is a derivative of the option price with respect to the domestic interest rate and the other is a derivative with respect to the foreign interest rate.

The notions of *vega* and *volga* usually refer to the first and second derivative with respect to volatility. In the Heston model we use them for the first and second derivative with respect to the *initial variance*:

$$\begin{aligned} \frac{\partial h(t)}{\partial v_t} &= e^{-r_f \tau} S_t \frac{\partial}{\partial v_t} P_1(\log S_t, v_t, \tau, \log K) - \\ &\quad - K e^{-r_d \tau} \frac{\partial}{\partial v_t} P_2(\log S_t, v_t, \tau, \log K), \end{aligned} \quad (1.25)$$

$$\begin{aligned} \frac{\partial^2 h(t)}{\partial v_t^2} &= e^{-r_f \tau} S_t \frac{\partial^2}{\partial v_t^2} P_1(\log S_t, v_t, \tau, \log K) - \\ &\quad - K e^{-r_d \tau} \frac{\partial^2}{\partial v_t^2} P_2(\log S_t, v_t, \tau, \log K), \end{aligned} \quad (1.26)$$

where

$$\frac{\partial}{\partial v_t} P_j(x, v_t, \tau, y) = \frac{1}{\pi} \int_0^\infty \Re \left[\frac{D(\tau, \varphi) e^{-i\varphi y} f_j(x, v_t, \tau, \varphi)}{i\varphi} \right] d\varphi, \quad (1.27)$$

$$\frac{\partial^2}{\partial v_t^2} P_j(x, v_t, \tau, y) = \frac{1}{\pi} \int_0^\infty \Re \left[\frac{D^2(\tau, \varphi) e^{-i\varphi y} f_j(x, v_t, \tau, \varphi)}{i\varphi} \right] d\varphi. \quad (1.28)$$

1.3.2 Computational Issues

Heston's solution is semi-analytical. Formulas (1.19-1.20) require to integrate functions f_j , which are typically of oscillatory nature. Hakala and Wystup (2002) propose to perform the integration in (1.18) with the Gauss-Laguerre quadrature using 100 for ∞ and 100 abscissas. Jäckel and Kahl (2005) suggest using the Gauss-Lobatto quadrature (e.g. Matlab's *quadl.m* function) and transform the original integral boundaries $[0, +\infty)$ to the finite interval $[0, 1]$.

As a number of authors have recently reported (Albrecher et al., 2006; Gatheral, 2006; Jäckel and Kahl, 2005), the real problem starts when the functions f_j are evaluated as part of the quadrature scheme. In particular, the calculation of the complex logarithm in eqn. (1.15) is prone to numerical instabilities. It turns out that taking the principal value of the logarithm causes C_j to jump discontinuously each time the imaginary part of the argument of the logarithm crosses the negative real axis. One solution is to keep track of the winding number in the integration (1.18), but is difficult to implement because standard numerical integration routines cannot be used. Jäckel and Kahl (2005) provide a practical solution to this problem.

A different approach was taken by Albrecher et al. (2006), see also Gatheral (2006). They proposed to make a simple transformation when computing the characteristic function and proved that the numerical stability of the resulting formulas is guaranteed under the full dimensional and unrestricted parameter space. Namely, their idea is to switch from g_j in (1.14) to

$$\tilde{g}_j = \frac{1}{g_j} = \frac{b_j - \rho\sigma\varphi i - d_j}{b_j - \rho\sigma\varphi i + d_j}, \quad (1.29)$$

which leads to new formulas for C_j and D_j :

$$C_j(\tau, \varphi) = (r_d - r_f)\varphi i\tau + \frac{\kappa\theta}{\sigma^2} \left\{ (b_j - \rho\sigma\varphi i - d_j)\tau - 2 \log \left(\frac{1 - \tilde{g}_j e^{-d_j\tau}}{1 - \tilde{g}_j} \right) \right\}, \quad (1.30)$$

$$D_j(\tau, \varphi) = \frac{b_j - \rho\sigma\varphi i - d_j}{\sigma^2} \left(\frac{1 - e^{-d_j\tau}}{1 - \tilde{g}_j e^{-d_j\tau}} \right), \quad (1.31)$$

in (1.17). Note, that the only differences between eqns. (1.14)-(1.16) and (1.29)-(1.31), respectively, are the flipped minus and plus signs in front of the d_j 's.

The mispricings resulting from using (1.14)-(1.16) are not that obvious to notice. In fact, if we price and backtest on short or middle term maturities only, we might not detect the problem at all. However, the deviations can become extreme for longer maturities (typically above 3-5 years; the exact threshold is parameter dependent, see Albrecher et al., 2006).

Apart from the above semi-analytical solution for vanilla options, alternative approaches can be utilized. These include the general Fast Fourier Transform (FFT) approach of Carr and Madan (1999), finite differences, finite element methods and Monte Carlo simulations. The FFT-based pricing technique is discussed in Section 1.3.4. As for the other methods, finite differences must be used with care since high precision is required to invert scarce matrices. The Crank-Nicholson, ADI (Alternate Direction Implicit), and Hopscotch schemes can be used, however, ADI is not suitable to handle nonzero correlation. Boundary conditions must be set appropriately, for details see Kluge (2002). On the other hand, finite element methods can be applied to price both the vanillas and exotics, as explained for example in Apel, Winkler, and Wystup (2002).

Finally, the Monte Carlo approach requires attention as the simple Euler discretization of the CIR process (1.2) may give rise to a negative variance. To deal with this problem, practitioners generally either adopt (i) the absorbing $v_t = \max(v_t, 0)$, or (ii) the reflecting principle $v_t = \max(v_t, -v_t)$. More elegant,

but computationally more demanding solutions include sampling from the exact transition law (Glasserman, 2004) or application of the quadratic-exponential (QE) scheme (Andersen, 2008). For a recent survey see the latter reference, where several new algorithms for time-discretization and Monte Carlo simulation of Heston-type SV models are introduced and compared in a thorough simulation study. Both the absorbing/reflecting patches and the QE scheme are implemented in the `simHeston.m` function used to plot Figure 1.1.

1.3.3 Behavior of the Variance Process and the Feller Condition

The CIR process for the variance, as defined by (1.2), always remains non-negative. This is an important property which, for instance, is not satisfied by the Ornstein-Uhlenbeck process. However, ideally one would like a variance process which is strictly positive, because otherwise it degenerates to a deterministic function for the time it stays at zero. As it turns out, the CIR process remains strictly positive under the condition that

$$\alpha := \frac{4\kappa\theta}{\sigma^2} \geq 2, \quad (1.32)$$

which is often referred to as the *Feller condition*. We call α the *dimensionality* of the corresponding Bessel process (see below). If the condition is not satisfied, i.e. for $0 < \alpha < 2$, the CIR process will visit 0 recurrently but will not stay at zero, i.e. the 0-boundary is strongly reflecting.

Unfortunately, when calibrating the Heston model to market option prices it is not uncommon for the parameters to violate the Feller condition (1.32). This is not a complete disaster, as the variance process can only hit zero for an infinitesimally small amount of time, but it is still worrying as very low levels of volatility (e.g. say below 1%) are repeatedly reached for short amounts of time and that is not something observed in the market.

Besides being important from a modeling point of view, the Feller condition also plays a role in computational accuracy. For Monte Carlo simulations special care has to be taken so that the simulated paths do not go below zero if (1.32) is not satisfied. On the PDE side, the Feller condition determines whether the zero-variance boundary is in- or out-flowing, that is to say whether the convection vector at the boundary points in- or outwards. To see this, we need

to write the log-spot transformed Heston PDE in convection-diffusion form

$$\frac{\partial}{\partial t}U = \operatorname{div}(A \operatorname{grad} U) - \operatorname{div}(Ub) + f, \quad (1.33)$$

and obtain

$$b(x, v) = v \left(\frac{\frac{1}{2}}{\kappa + \lambda} \right) + \left(\frac{\frac{1}{2}\rho\sigma + r_f - r_d}{\frac{1}{2}\sigma^2 - \kappa\theta} \right), \quad (1.34)$$

which is out-flowing at the $v = 0$ boundary if

$$\frac{1}{2}\sigma^2 - \kappa\theta < 0, \quad (1.35)$$

in which case the Feller condition (1.32) is satisfied.

Having introduced and discussed the importance of this condition, the question remains how it is derived. We only give an overview here and refer the interested reader to Chapter 6.3 in Jeanblanc, Yor, and Chesney (2009) for a more thorough treatment.

The main idea is to relate the CIR process to the family of squared Bessel processes which have well known properties. We call X_t an α -dimensional squared Bessel process and denote it by $\operatorname{BES}^2(\alpha)$ if it follows the dynamics of

$$dX_t = \alpha dt + 2\sqrt{X_t^+} dW_t, \quad (1.36)$$

with $X_t^+ := \max\{0, X_t\}$. This definition makes sense for any real valued α . However, in the case of integer valued α we have an interesting interpretation: a $\operatorname{BES}^2(\alpha)$ process X_t , with $X_0 = 0$, follows the same dynamics as the squared distance to the origin of an α -dimensional Brownian motion, i.e.

$$X_t = \sum_{i=1}^{\alpha} B_t^{(i)}, \quad (1.37)$$

with $B^{(i)}$ being independent Brownian motions. From this we can already conclude that for $\alpha = 2$, X_t will never reach zero and, using the stochastic comparison theorem (Rogers and Williams, 2000), this property remains true for any $\alpha \geq 2$. Similarly for $0 < \alpha \leq 1$ the value zero will be repeatedly hit (for $\alpha = 1$ this happens as often as a one-dimensional Brownian motion crosses zero).

The stochastic comparison theorem also immediately tells us that $\text{BES}^2(\alpha)$ processes are non-negative for non-negative α : for $\alpha = 0$ we get the trivial solution $X_t = 0$ and increasing the drift term to $\alpha > 0$ cannot make the paths any smaller. Different methods are needed, however, to examine the behavior for $1 < \alpha < 2$ (which turns out to be similar to the case of $0 < \alpha \leq 1$). For a proof we refer to Chapter V.48 in Rogers and Williams (2000), from which we also state the following additional properties. If X_t is an α -dimensional squared Bessel process $\text{BES}^2(\alpha)$ and $X_0 \geq 0$ then it is always non-negative and:

1. for $0 < \alpha < 2$ the process hits zero and this is recurrent, but the time spent at zero is zero,
2. for $\alpha = 2$ the process is strictly positive but gets arbitrarily close to 0 and ∞ ,
3. for $\alpha > 2$ the process is strictly positive and tends to infinity as time approaches infinity.

To translate these properties to the class of CIR processes, we only need to apply the following space-time transformation to the squared Bessel process. Define $dY_t := e^{-\kappa t} X_{\beta(e^{\kappa t} - 1)}$. Then Y_t follows the dynamics of

$$dY_t = \kappa(\alpha\beta - Y_t) dt + 2\sqrt{\kappa\beta Y_t} dW_t, \quad (1.38)$$

which is the same as the dynamics of the CIR process (1.2) if we set $\beta = \frac{\sigma^2}{4\kappa}$ and $\alpha = \frac{\theta}{\beta} = \frac{4\kappa\theta}{\sigma^2}$.

1.3.4 Option Pricing with FFT

In this section, we briefly describe the numerical option pricing approach of Carr and Madan (1999), which utilizes the characteristic function (cf) of the underlying instrument's price process. The basic idea of the method is to compute the price by Fourier inversion given an analytic expression for the cf of the option price.

The rationale for using this approach is twofold. Firstly, the algorithm offers a speed advantage, including the possibility to calculate prices for a whole range of strikes in a single run. Secondly, the cf of the log-price is known and has a simple form for many models, while the pdf is often either unknown in closed-form or complicated from the numerical point of view. For instance, for the

Heston model the cf takes the form (Heston, 1993; Jäckel and Kahl, 2005):

$$E\{\exp(i\varphi \log S_T)\} = f_2(x, v_t, \tau, \varphi), \quad (1.39)$$

where f_2 is given by (1.17).

Let us now derive the formula for the price of a European vanilla call option. Derivation of the put option price follows the same lines, for details see Lee (2004) and Schmelzle (2010). Alternatively we can use the call-put parity for European vanilla FX options (see e.g. Wystup, 2006). Let $h^C(\tau; k)$ denote the price of the call option maturing in $\tau = T - t$ years with a strike of $K = \exp(k)$:

$$h^C(\tau; k) = \int_k^\infty e^{-rT} (e^s - e^k) q_T(s) ds, \quad (1.40)$$

where q_T is the risk-neutral density of $s_T = \log S_T$. The function $h^C(\tau; k)$ is not square-integrable (see e.g. Rudin, 1991) because it converges to S_0 for $k \rightarrow -\infty$. Hence, consider a modified function $H^C(\tau; k) = e^{\alpha k} h^C(\tau; k)$, which is square-integrable for a suitable constant $\alpha > 0$. A sufficient condition is given by:

$$E\{(S_T)^{\alpha+1}\} < \infty, \quad (1.41)$$

which is equivalent to $\psi_T(0)$, i.e. the Fourier transform of $H^C(\tau; k)$, see (1.42) below, being finite. In an empirical study Schoutens, Simons, and Tistaert (2004) found that $\alpha = 0.75$ leads to stable algorithms, i.e. the prices are well replicated for many model parameters. This value also fulfills condition (1.41) for the Heston model (Borak, Detlefsen, and Härdle, 2005). Note, that for put options the condition is different: it is sufficient to choose $\alpha > 0$ such that $E\{(S_T)^{-\alpha}\} < \infty$ (Lee, 2004).

Now, compute the Fourier transform of $H^C(\tau; k)$:

$$\begin{aligned} \psi_T(v) &= \int_{-\infty}^{\infty} e^{ivk} H^C(\tau; k) dk \\ &= \int_{-\infty}^{\infty} e^{ivk} \int_k^\infty e^{\alpha k} e^{-rT} (e^s - e^k) q_T(s) ds dk \\ &= \int_{-\infty}^{\infty} e^{-rT} q_T(s) \int_{-\infty}^s \left\{ e^{\alpha k+s} - e^{(\alpha+1)k} \right\} e^{ivk} dk ds \\ &= \int_{-\infty}^{\infty} e^{-rT} q_T(s) \left\{ \frac{e^{(\alpha+1+iv)s}}{\alpha+iv} - \frac{e^{(\alpha+1+iv)s}}{\alpha+1+iv} \right\} ds \\ &= \frac{e^{-rT} f_2\{v - (\alpha+1)i\}}{\alpha^2 + \alpha - v^2 + i(2\alpha+1)v}, \end{aligned} \quad (1.42)$$





Figure 1.3: European call (*top left*) and put (*bottom left*) FX option prices obtained using the FFT method and the ‘Analytical’ formula (1.12) for a sample set of parameters. *Right panels*: The percentage differences between the prices: $(\text{FFT} - \text{Analytical})/\text{Analytical} \times 100\%$.

 STF2hes03.m

where f_2 is the cf of q_T , see (1.39). We get the option price in terms of ψ_T using Fourier inversion:

$$h^C(\tau; k) = \frac{\exp(-\alpha k)}{\pi} \int_0^\infty e^{-ivk} \psi(v) dv. \quad (1.43)$$

This integral can be numerically approximated as (Carr and Madan, 1999):

$$h^C(\tau; k_u) \approx \frac{e^{-\alpha k_u}}{\pi} \sum_{j=1}^N e^{-\frac{2\pi i}{N}(j-1)(u-1)} e^{ibv_j} \psi(v_j) \frac{\eta}{3} \{3 + (-1)^j - \delta_{j-1}\}, \quad (1.44)$$

where $k_u = \frac{1}{\eta} \{-\pi + \frac{2\pi}{N}(u-1)\}$, $u = 1, \dots, N$, $\eta > 0$ is the distance between the points of the integration grid, $v_j = \eta(j-1)$, $j = 1, \dots, N$, and δ is the Dirac function.

As can be seen in Figure 1.3 the differences between the FFT method and the (semi-)analytical formula (1.12) are relatively small. The differences result from the fact that the former method yields ‘exact’ values only on the grid k_u . In order to preserve the speed of the FFT-based algorithm we use linear interpolation between the grid points. This approach, however, slightly overestimates the true prices since the option price is a convex function of the strike (Borak, Detlefsen, and Härdle, 2005). It can be clearly seen that near the grid points the prices obtained by both methods coincide, while between the grid points the FFT-based algorithm generates higher prices than the analytical solution, see the right panels in Figure 1.3.

1.4 Calibration

1.4.1 Qualitative Effects of Changing the Parameters

Before calibrating the model to market data we will show how changing the input parameters affects the shape of the fitted smile curve. This analysis will help in reducing the dimensionality of the problem. In all plots of this subsection the solid blue curve with x’s is the smile obtained for $v_0 = 0.01$, $\sigma = 0.2$, $\kappa = 1.5$, $\theta = 0.015$, and $\rho = 0.05$.

First, take a look at the initial variance (top left panel in Figure 1.4). Apparently, changing v_0 allows for adjustments in the height of the smile curve. On the other hand, the volatility of variance (vol of vol) has a different impact on the smile. Increasing σ increases the convexity of the fit, see the top right panel in Figure 1.4. In the limiting case, setting σ equal to zero would produce a deterministic process for the variance and, hence, a volatility which does not admit any smile (a constant curve).

The effects of changing the long-run variance θ are similar to those observed by changing the initial variance, compare the left panels in Figure 1.4. It seems promising to choose the initial variance a priori, e.g. set $\sqrt{v_0} =$ implied at-the-money (ATM) volatility, and only let the long-run variance θ vary. In particular, a different initial variance for different maturities would be inconsistent.

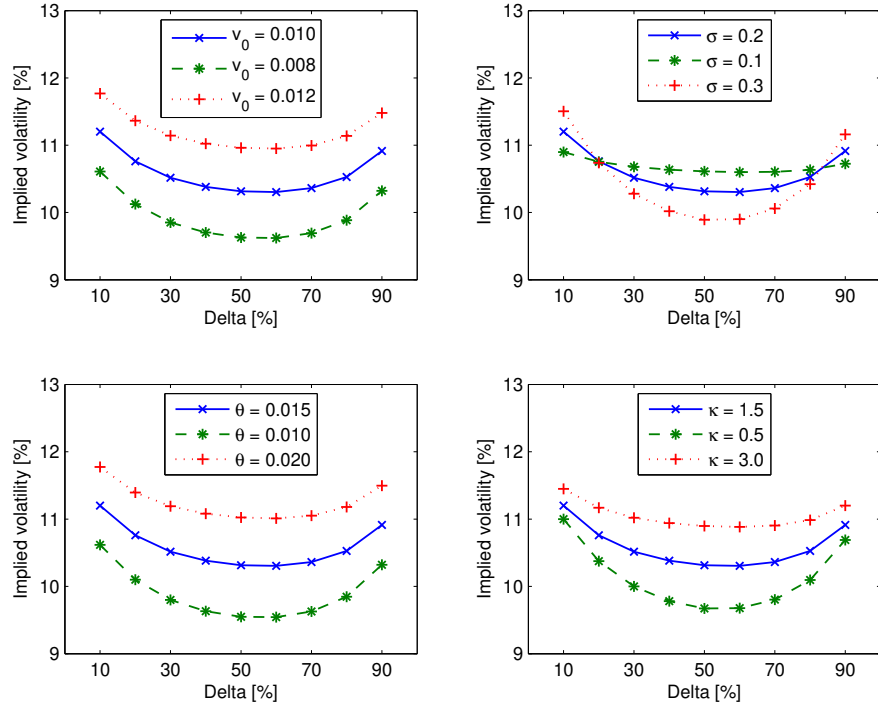



Figure 1.4: The effects of changing the model parameters on the shape of the smile: initial variance v_0 (*top left*), volatility of variance σ (*top right*), long-run variance θ (*bottom left*), and mean reversion level κ (*bottom right*).

 STF2hes04.m

Changing the mean reversion κ affects the ATM part more than the extreme wings of the smile curve. The low/high deltas (Δ) remain almost unchanged while increasing the mean reversion lifts the center, see the bottom right panel in Figure 1.4. Moreover, the influence of κ is often compensated by a stronger vol of vol σ . This suggests fixing mean reversion (at some level, say $\kappa = 1.5$) and only calibrating the remaining three parameters. If the obtained parameters do not satisfy the Feller condition (1.32), it might be worthwhile to fix κ at a different level, say $\tilde{\kappa} = 3$, recalibrate the remaining parameters and check if the new estimates fulfill the condition and lead to a more realistic variance process.




Figure 1.5: The effects of changing the correlation on the shape of the smile.

 STF2hes05.m

Finally, let us look at the influence of correlation. The uncorrelated case produces a fit that looks like a symmetric smile curve centered at-the-money, see Figure 1.5. However, it is not exactly symmetric. Changing ρ changes the degree of symmetry. In particular, positive correlation makes calls more expensive, negative correlation makes puts more expensive. Note, that the model yields a volatility skew, a typically observed volatility structure in equity markets, only when the correlation is set to a very high or low value

1.4.2 The Calibration Scheme

Calibration of SV models can be done in two conceptually different ways. One is to look at the time series of historical data. Estimation methods such as Generalized, Simulated, and Efficient Methods of Moments (respectively GMM, SMM, and EMM), as well as Bayesian MCMC have been extensively applied. See Broto and Ruiz (2004) for a review. In the Heston model we could also try to fit empirical distributions of returns to the marginal distributions specified in (1.4) via a minimization scheme. Unfortunately, all historical approaches have one common flaw: they do not allow for the estimation of the market price of volatility risk $\lambda(t, v, S)$. Observing only the underlying spot price and estimating SV models with this information will not yield correct prices for the derivatives.

This leads us to the second estimation approach: instead of using the spot data we calibrate the model to the volatility smile (i.e. prices of vanilla options). In this case we do not need to worry about estimating the market price of volatility risk as it is already embedded in the market smile. This means that we can set $\lambda = 0$ by default and just determine the remaining parameters.

As a preliminary step, we have to retrieve the strikes since the smile in FX markets is specified as a function of delta. Comparing the Black-Scholes type formulas (in the FX market setting we have to use the Garman and Kohlhagen (1983) specification) for delta and the option premium yields the relation for the strikes K_i . From a computational point of view this stage requires only an inversion of the Gaussian distribution function. Next, based on the findings of Section 1.4.1, we fix two parameters (initial variance v_0 and mean reversion κ) and fit the remaining three: volatility of variance σ , long-run variance θ , and correlation ρ for a fixed time to maturity and a given vector of market Black-Scholes implied volatilities $\{\hat{\sigma}_i\}_{i=1}^n$ for a given set of delta pillars $\{\Delta_i\}_{i=1}^n$.

After fitting the parameters we compute the option prices in the Heston model using (1.12) and retrieve the corresponding Black-Scholes model implied volatilities $\{\sigma_i\}_{i=1}^n$ via a standard bisection method (function *fzero.m* in Matlab uses a combination of bisection, secant, and inverse quadratic interpolation methods). The next step is to define an objective function, which we choose to be the Sum of Squared Errors (SSE):

$$\text{SSE}(\kappa, \theta, \sigma, \rho, v_0) = \sum_{i=1}^n \{\hat{\sigma}_i - \sigma_i(\kappa, \theta, \sigma, \rho, v_0)\}^2. \quad (1.45)$$

We compare the volatilities because they are all of similar magnitude, in contrast to the prices which can range a few orders of magnitude for it-the-money (ITM) vs. out-of-the-money (OTM) options. In addition, one could introduce weights for all the summands to favor ATM or OTM fits. Finally we minimize over this objective function using a simplex search routine (*fminsearch.m* in Matlab) to find the optimal set of parameters.

1.4.3 Sample Calibration Results

We are now ready to calibrate the Heston model to market data. First, we take the EUR/USD volatility surface on July 1, 2004 and fit the parameters in the Heston model according to the calibration scheme discussed earlier. The results are shown in Figures 1.6–1.7. Note, that the fit is very good for intermediate

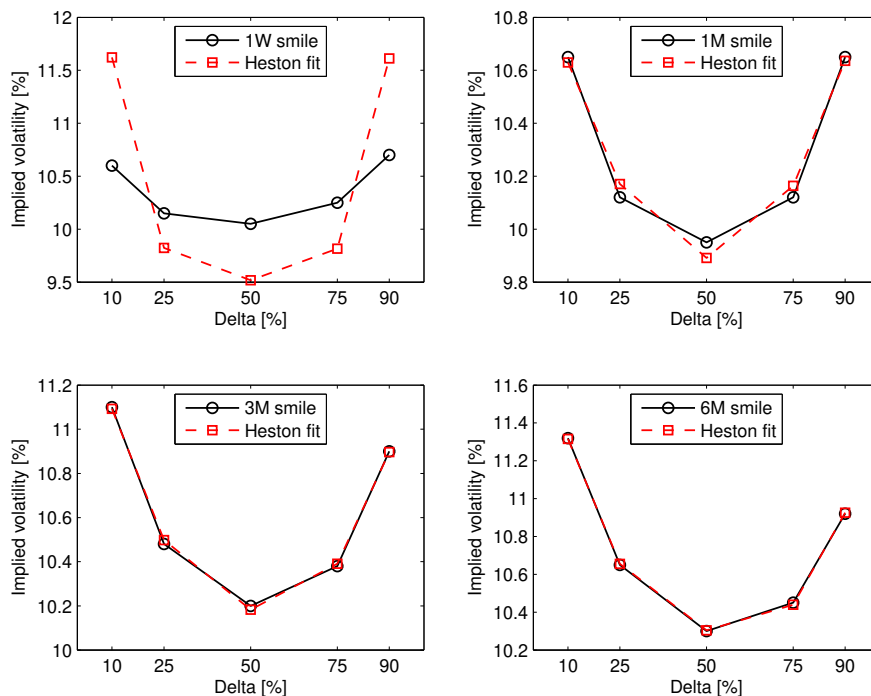



Figure 1.6: The EUR/USD market smile on July 1, 2004 and the fit obtained with the Heston model for $\tau = 1$ week (*top left*), 1 month (*top right*), 3 months (*bottom left*), and 6 months (*bottom right*).

 STF2hes06.m

and long maturities (three months and more). Unfortunately, the Heston model does not perform satisfactorily for short maturities (see Section 1.5.2 for a discussion of alternative approaches). Comparing with the fits in Weron and Wystup (2005) for the same data, the long maturity (2Y) fit is better. This is due to a more efficient optimization routine (Matlab 7.2 vs. XploRe 4.7) and utilization of the transformed formulas (1.29)-(1.31) instead of the original ones (1.14)-(1.16).

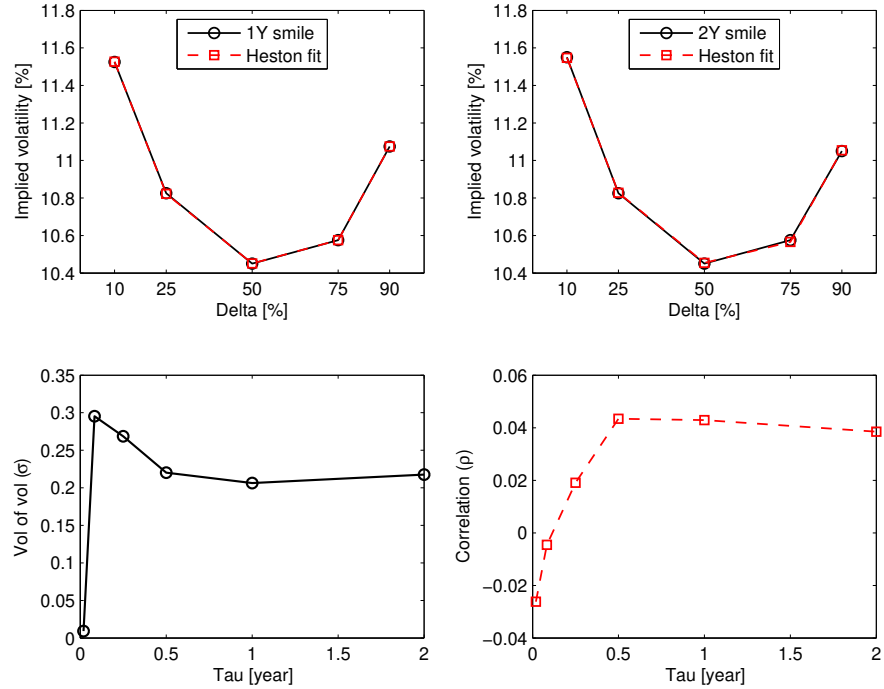



Figure 1.7: The EUR/USD market smile on July 1, 2004 and the fit obtained with the Heston model for $\tau = 1$ year (*top left*) and 2 years (*top right*). The term structure of the vol of vol and correlation visualizes the problem with fitting the smile in the short term (*bottom panels*).

 STF2hes06.m

Now we take a look at more recent data and calibrate the Heston model to the EUR/USD volatility surface on July 22, 2010. The results are shown in Figures 1.8–1.9. Again the fit is very good for intermediate and long maturities, but unsatisfactory for maturities under three months. The term structure of the vol of vol and correlation visualizes the problem with fitting the smile in the short term for both datasets. The calibrated vol of vol is very low for the 1W smiles, then jumps to a higher level. The correlation behaves similarly, for the 1W smiles it is much lower than for the remaining maturities. In 2004 it

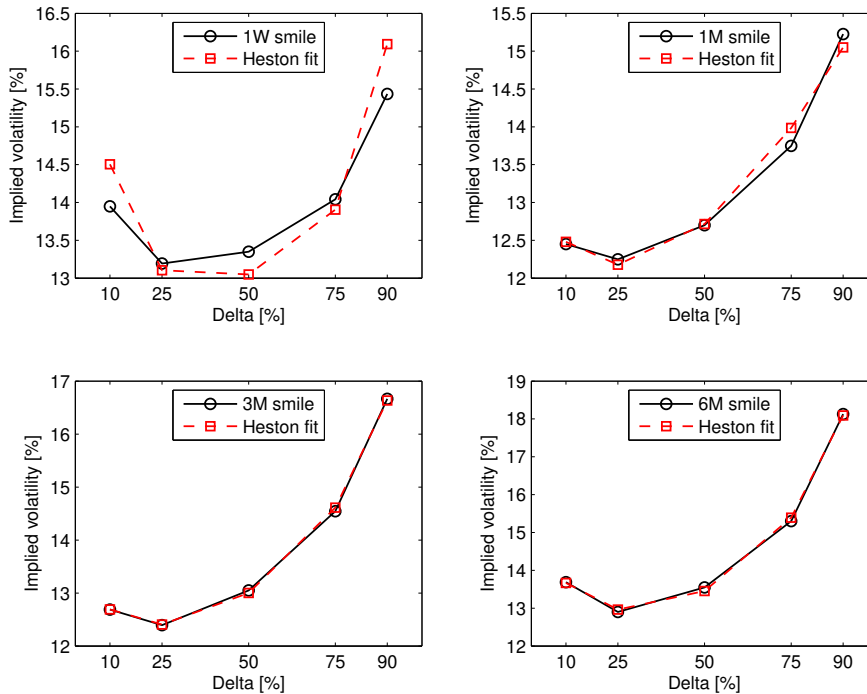



Figure 1.8: The EUR/USD market smile on July 22, 2010 and the fit obtained with the Heston model for $\tau = 1$ week (*top left*), 1 month (*top right*), 3 months (*bottom left*), and 6 months (*bottom right*).

 STF2hes07.m

additionally changes sign as the skew changes between short and longer-term tenors. Also note, that the more skewed smiles in 2010 require much higher (anti-)correlation ($-0.4 < \rho < -0.3$ for $\tau \geq 1M$) to obtain a decent fit than the more symmetric smiles in 2004 ($-0.01 < \rho < 0.05$ for $\tau \geq 1M$).

As these examples show, the Heston model can be successfully applied to modeling the volatility smile of vanilla FX options in the mid- to longer-term. There are essentially three parameters to fit, namely the long-run variance (θ), which corresponds to the ATM level of the market smile, the vol of vol (σ), which corresponds to the convexity of the smile (in the market often quoted as *but-*

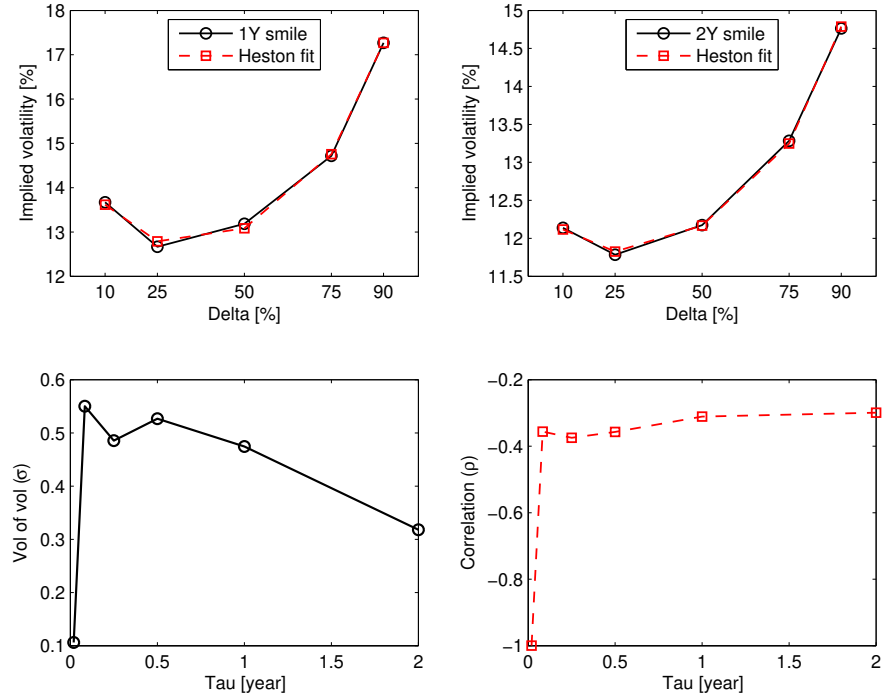



Figure 1.9: The EUR/USD market smile on July 22, 2010 and the fit obtained with the Heston model for $\tau = 1$ year (*top left*) and 2 years (*top right*). Again the term structure of the vol of vol and correlation visualizes the problem with fitting the smile in the short term (*bottom panels*).

 STF2hes07.m

terflies), and the correlation (ρ), which corresponds to the skew of the smile (quoted as *risk reversals*). It is this direct link of the model parameters to the market that makes the Heston model so attractive to practitioners.

The key application of the model is to calibrate it to vanilla options and afterwards employ it for pricing exotics (like one-touch options, see Wystup, 2003) in either a finite difference grid or a Monte Carlo simulation. Surprisingly, the results often coincide with the *traders' rule of thumb* pricing method (Wystup,

2006). This might also simply mean that a lot of traders use the same model. After all, it is a matter of belief which model reflects the reality most suitably. Recent ideas are to take prices of one-touch options along with prices of vanilla options from the market and use this common input to calibrate the Heston model.

1.5 Beyond the Heston Model

1.5.1 Time-dependent Parameters

As we have seen in Section 1.4.3, performing calibrations for different time slices of the volatility matrix produces different values of the parameters. This suggests a term structure of some parameters in the Heston model. Therefore, we need to generalize the CIR process (1.2) to the case of time-dependent parameters, i.e. we consider the process:

$$dv_t = \kappa(t)\{\theta(t) - v_t\} dt + \sigma(t)\sqrt{v_t} dW_t, \quad (1.46)$$

for some nonnegative deterministic parameter functions $\sigma(t)$, $\kappa(t)$, and $\theta(t)$. The formula for the mean turns out to be:

$$\mathbb{E}(v_t) = g(t) = v_0 e^{-K(t)} + \int_0^t \kappa(s)\theta(s)e^{K(s)-K(t)} ds, \quad (1.47)$$

with $K(t) = \int_0^t \kappa(s) ds$. The result for the second moment is:

$$\mathbb{E}(v_t^2) = v_0^2 e^{-2K(t)} + \int_0^t \{2\kappa(s)\theta(s) + \sigma^2(s)\}g(s)e^{2K(s)-2K(t)} ds, \quad (1.48)$$

and hence for the variance (after some algebra):

$$\text{Var}(v_t) = \int_0^t \sigma^2(s)g(s)e^{2K(s)-2K(t)} ds. \quad (1.49)$$

The formula for the variance allows us to compute forward volatilities of variance explicitly. Assuming known values σ_{T_1} and σ_{T_2} for times $0 < T_1 < T_2$, we want to determine the forward volatility of variance σ_{T_1, T_2} which matches the

corresponding variances, i.e.

$$\begin{aligned} \int_0^{T_2} \sigma_{T_2}^2 g(s) e^{2\kappa(s-T_2)} ds &= \\ &= \int_0^{T_1} \sigma_{T_1}^2 g(s) e^{2\kappa(s-T_2)} ds + \int_{T_1}^{T_2} \sigma_{T_1, T_2}^2 g(s) e^{2\kappa(s-T_2)} ds. \end{aligned} \quad (1.50)$$

The resulting forward volatility of variance is thus:

$$\sigma_{T_1, T_2}^2 = \frac{\sigma_{T_2}^2 H(T_2) - \sigma_{T_1}^2 H(T_1)}{H(T_2) - H(T_1)}, \quad (1.51)$$

where

$$H(t) = \int_0^t g(s) e^{2\kappa s} ds = \frac{\theta}{2\kappa} e^{2\kappa t} + \frac{1}{\kappa} (v_0 - \theta) e^{\kappa t} + \frac{1}{\kappa} \left(\frac{\theta}{2} - v_0 \right). \quad (1.52)$$

Assuming known values ρ_{T_1} and ρ_{T_2} for times $0 < T_1 < T_2$, we want to determine the forward correlation coefficient ρ_{T_1, T_2} to be active between times T_1 and T_2 such that the covariance between the Brownian motions of the variance process and the exchange rate process agrees with the given values ρ_{T_1} and ρ_{T_2} . This problem has a simple answer, namely:

$$\rho_{T_1, T_2} = \rho_{T_2}, \quad T_1 \leq t \leq T_2. \quad (1.53)$$

This can be seen by writing the Heston model in the form:

$$dS_t = S_t \left(\mu dt + \sqrt{v_t} dW_t^{(1)} \right), \quad (1.54)$$

$$dv_t = \kappa(\theta - v_t) dt + \rho\sigma\sqrt{v_t} dW_t^{(1)} + \sqrt{1 - \rho^2}\sigma\sqrt{v_t} dW_t^{(2)}, \quad (1.55)$$

for a pair of independent Brownian motions $W^{(1)}$ and $W^{(2)}$. Observe that choosing the forward correlation coefficient as stated does not conflict with the computed forward volatility.

1.5.2 Jump-diffusion Models

While trying to calibrate short term smiles, the volatility of volatility often seems to explode along with the speed of mean reversion. This is a strong indication that the process ‘wants’ to jump, which of course it is not allowed

to do. This observation, together with market crashes, has lead researchers to consider models with jumps (Gatheral, 2006; Martinez and Senge, 2002). Such models have been investigated already in the mid-seventies (Merton, 1976), long before the advent of SV. Jump-diffusion (JD) models are, in general, more challenging to handle numerically than SV models. Like the latter, they result in an incomplete market. But, whereas SV models can be made complete by the introduction of one (or a few) traded options, a JD model typically requires the existence of a continuum of options for the market to be complete.

Bates (1996) and Bakshi, Cao, and Chen (1997) suggested using a combination of jumps and stochastic volatility. This approach allows for even a better fit to market data, but at the cost of a larger number of parameters to calibrate from the same market volatility smile. Andersen and Andreasen (2000) let the stock dynamics be described by a JD process with local volatility. This method combines ease of modeling steep, short-term volatility skews (jumps) and accurate fitting to quoted option prices (deterministic volatility function). Other alternative approaches utilize Lévy processes (Cont and Tankov, 2003; Eberlein, Kallsen, and Kristen, 2003) or mixing unconditional disturbances (Tompkins and D'Ecclesia, 2006).

Bibliography

- Albrecher, H., Mayer, P., Schoutens, W. and Tistaert, J. (2006). The little Heston trap, *Wilmott Magazine*, **January**: 83–92. 231–262.
- Andersen, L. and Andreasen, J. (2000). Jump-Diffusion Processes: Volatility Smile Fitting and Numerical Methods for Option Pricing, *Review of Derivatives Research* **4**: 231–262.
- Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model, *The Journal of Computational Finance* **11(3)**, 1–42.
- Apel, T., Winkler, G., and Wystup, U. (2002). Valuation of options in Heston's stochastic volatility model using finite element methods, in J. Hakala, U. Wystup (eds.) *Foreign Exchange Risk*, Risk Books, London.
- Bakshi, G., Cao, C. and Chen, Z. (1997). Empirical Performance of Alternative Option Pricing Models, *Journal of Finance* **52**: 2003–2049.
- Bates, D. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options, *Review of Financial Studies* **9**: 69–107.
- Borak, S., Detlefsen, K., and Härdle, W. (2005). FFT-based option pricing, in P. Cizek, W. Härdle, R. Weron (eds.) *Statistical Tools for Finance and Insurance*, Springer, Berlin.
- Broto, C., and Ruiz, E. (2004). Estimation methods for stochastic volatility models: A survey, *Journal of Economic Surveys* **18(5)**: 613–649.
- Carr, P. and Madan, D. (1999). Option valuation using the fast Fourier transform, *Journal of Computational Finance* **2**: 61–73.
- Cont, R., and Tankov, P. (2003). *Financial Modelling with Jump Processes*, Chapman & Hall/CRC.

- Derman, E. and Kani, I. (1994). Riding on a Smile, *RISK* **7(2)**: 32–39.
- Dragulescu, A. A. and Yakovenko, V. M. (2002). Probability distribution of returns in the Heston model with stochastic volatility, *Quantitative Finance* **2**: 443–453.
- Dupire, B. (1994). Pricing with a Smile, *RISK* **7(1)**: 18–20.
- Eberlein, E., Kallsen, J., and Kristen, J. (2003). Risk Management Based on Stochastic Volatility, *Journal of Risk* **5(2)**: 19–44.
- Fengler, M. (2005). *Semiparametric Modelling of Implied Volatility*, Springer.
- Fouque, J.-P., Papanicolaou, G., and Sircar, K.R. (2000). *Derivatives in Financial Markets with Stochastic Volatility*, Cambridge University Press.
- Garman, M. B. and Kohlhagen, S. W. (1983). Foreign currency option values, *Journal of International Money & Finance* **2**: 231–237.
- Gatheral, J. (2006). *The Volatility Surface: A Practitioner’s Guide*, Wiley.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*, Springer-Verlag, New-York.
- Hakala, J. and Wystup, U. (2002). Heston’s Stochastic Volatility Model Applied to Foreign Exchange Options, in J. Hakala, U. Wystup (eds.) *Foreign Exchange Risk*, Risk Books, London.
- Heston, S. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* **6**: 327–343.
- Hull, J. and White, A. (1987). The Pricing of Options with Stochastic Volatilities, *Journal of Finance* **42**: 281–300.
- Jäckel, P. and Kahl, C. (2005). Not-So-Complex Logarithms in the Heston Model, *Wilmott Magazine* **19**: 94–103.
- Jeanblanc, M., Yor, M. and Chesney, M. (2009). *Mathematical Methods for Financial Markets*, Springer.
- Kluge, T. (2002). *Pricing derivatives in stochastic volatility models using the finite difference method*, Diploma thesis, Chemnitz Technical University.
- Lee, R., (2004). Option pricing by transform methods: extensions, unification and error control, *Journal of Computational Finance* **7(3)**: 51–86.

- Martinez, M. and Senge, T. (2002). A Jump-Diffusion Model Applied to Foreign Exchange Markets, in J. Hakala, U. Wystup (eds.) *Foreign Exchange Risk*, Risk Books, London.
- Merton, R. (1973). The Theory of Rational Option Pricing, *Bell Journal of Economics and Management Science* **4**: 141–183.
- Merton, R. (1976). Option Pricing when Underlying Stock Returns are Discontinuous, *Journal of Financial Economics* **3**: 125–144.
- Reiss, O. and Wystup, U. (2001). Computing Option Price Sensitivities Using Homogeneity, *Journal of Derivatives* **9(2)**: 41–53.
- Rogers, L.C.G. and Williams, D. (2000). *Diffusions, Markov Processes, and Martingales – Volume Two: Itô Calculus*, McGraw-Hill.
- Rubinstein, M. (1994). Implied Binomial Trees, *Journal of Finance* **49**: 771–818.
- Rudin, W., (1991). *Functional Analysis*, McGraw-Hill.
- Schoutens, W., Simons, E. and Tistaert, J. (2004). A perfect calibration! Now what? *Wilmott March*: 66–78.
- Schmelzle, M., (2010). Option pricing formulae using Fourier transform: Theory and application, *Working paper*.
- Stein, E. and Stein, J. (1991). Stock Price Distributions with Stochastic Volatility: An Analytic Approach, *Review of Financial Studies* **4(4)**: 727–752.
- Tompkins, R. G. and D’Ecclesia, R. L. (2006). Unconditional Return Disturbances: A Non-Parametric Simulation Approach, *Journal of Banking and Finance* **30(1)**: 287–314.
- Weron, R. (2004). Computationally intensive Value at Risk calculations, in J.E. Gentle, W. Härdle, Y. Mori (eds.) *Handbook of Computational Statistics*, Springer.
- Weron, R., and Wystup, U. (2005). Heston’s model and the smile, in P. Cizek, W. Härdle, R. Weron (eds.) *Statistical Tools for Finance and Insurance*, Springer.
- Wystup, U. (2003). The market price of one-touch options in foreign exchange markets, *Derivatives Week*, **12(13)**.
- Wystup, U. (2006). *FX Options and Structured Products*, Wiley.