Stochastic volatility models and Kelvin waves

2008 J. Phys. A: Math. Theor. 41 344012

(http://iopscience.iop.org/1751-8121/41/34/344012)

IOP PUBLISHING JOURNAL OF PHYSICS A: MATHEMATICAL AND THEORETICAL

J. Phys. A: Math. Theor. **41** (2008) 344012 (23pp) doi:10.1088/1751-8113/41/34/344012

Stochastic volatility models and Kelvin waves

Alex Lipton[1,3] **and Artur Sepp**[2]

[1] Merrill Lynch, Mlfc Main, 2 King Edward Street, London EC1A 1HQ, UK
[2] Merrill Lynch, 4 World Financial Center, New York, NY 10080, USA

E-mail: Alex_Lipton@ml.com and Artur_Sepp@ml.com

Received 11 December 2007, in final form 22 May 2008
Published 11 August 2008
Online at stacks.iop.org/JPhysA/41/344012

Abstract

We use stochastic volatility models to describe the evolution of an asset price, its instantaneous volatility and its realized volatility. In particular, we concentrate on the Stein and Stein model (SSM) (1991) for the stochastic asset volatility and the Heston model (HM) (1993) for the stochastic asset variance. By construction, the volatility is not sign definite in SSM and is non-negative in HM. It is well known that both models produce closed-form expressions for the prices of *vanilla option* via the Lewis–Lipton formula. However, the numerical pricing of *exotic options* by means of the finite difference and Monte Carlo methods is much more complex for HM than for SSM. Until now, this complexity was considered to be an acceptable price to pay for ensuring that the asset volatility is non-negative. We argue that having negative stochastic volatility is a psychological rather than financial or mathematical problem, and advocate using SSM rather than HM in most applications. We extend SSM by adding volatility jumps and obtain a closed-form expression for the density of the asset price and its realized volatility. We also show that the current method of choice for solving pricing problems with stochastic volatility (via the affine ansatz for the Fourier-transformed density function) can be traced back to the Kelvin method designed in the 19th century for studying wave motion problems arising in fluid dynamics.

PACS numbers: 02.50.−r, 89.65.−s

(Some figures in this article are in colour only in the electronic version)

[3] Visiting Professor at Department of Mathematics, Imperial College London, London, SW7 2BZ, UK.

1. Introduction

1.1. Volatility modeling

Empirical studies indicate that asset price volatility is a random process and, in general, cannot be described by a single number. Yet, the option pricing models assuming a deterministic asset volatility originated by Black–Scholes (1973) and Merton (1973) (BSMM) are very popular in practice because of their relative simplicity. Vanilla option traders understand the limitations of BSMM, and know how to adjust the model and manage their risks appropriately. However, more complex financial products, such as forward-starting options, or options on the realized variance of an asset, derive their value from the asset volatility rather than its price, so for the pricing and risk management of such products traders tend to use stochastic volatility models.

We note that there exists a large class of so-called local volatility models (LVM), originated by Cox (1975) in a parametric form (the constant elasticity of variance model) and by Dupire (1994) in a non-parametric form (the local volatility surface model), which specify that the asset volatility depends only on the asset price level and time, so that all the uncertainty in the volatility dynamics is driven by the uncertainty of the asset price. Trading experience suggests that although most of these models can explain today's market data for simple options almost perfectly, they tend to have poor predictive and explanatory power, and are not satisfactory for the risk management of complex trades.

The use of a stochastic volatility model in practice consists of two major steps: first, the adjustment of the model parameters to fit vanilla options prices (model calibration); and, second, the application of the calibrated model to compute the values and risk parameters of complex trades. The first step is important because we want to express the risks of complex trades in terms of the risks of liquid vanilla options, which we shall subsequently use to hedge against these risks. As a result, it is important that our model is consistent with the values of these liquid vanilla options. The second step is typically achieved through numerical solution of the corresponding partial differential equations (PDEs), or Monte Carlo (MC) simulations of the corresponding stochastic differential equations (SDEs). It is common for academics to concentrate only on the first part by deriving closed-form formulae for vanilla option values and using them to estimate model parameters. However, for stochastic volatility models to be useful in practice, we have to formulate the pricing problem in either the PDE or MC frameworks (or both) and to ensure that the chosen stochastic volatility model allows robust implementation of the appropriate numerical algorithms.

In this paper, we study two stochastic volatility models: the first one assumes that the asset volatility is driven by the Ornstein–Uhlenbeck process (OUP), the Stein and Stein (1991) model (SSM); and the second one assumes that the asset variance is driven by the square root process (SRP), the Heston (1993) model (HM).

Of all the stochastic volatility models, HM appears to be the most popular one for pricing and hedging of price-dependent options depending on the future level of the asset volatility, such as forward-starting options and options on the realized variance[4]. Nonetheless, as we will show later, while yielding closed-form pricing formulae for vanilla options, HM is not well suited for pricing exotic options because it is hard to deal with numerically. The main reason is that in this model the asset variance is constrained to be positive (provided that the standard Feller condition is satisfied), and non-negative in general, which imposes a nontrivial

[4] The SABR model (Hagan *et al* 2002), which is ubiquitous in the fixed-income market is static in nature because it has no mean-reversion built into the volatility dynamics; it is mostly used for quoting and fitting of implied volatility surfaces for vanilla options rather than for dynamic pricing and hedging of exotic path-dependent deals.

boundary condition. In practice HM often implies a relatively high probability for observing small values of the variance, so that the corresponding boundary condition cannot be ignored and has to be dealt with numerically.

SSM has often been criticized (and discarded) on the grounds that it does not preclude the asset volatility from becoming negative. However, the non-negativity constraint should be imposed on the variance rather than the volatility itself, so that there are no conceptual or computational issues with having negative volatility (also see Aase (2004) for a specific example). Since the natural domain for the asset volatility covers the entire real axis, we can use natural boundary conditions in our calculations which are much easier to deal with. An additional drawback of the original SSM (1991) is that it assumes zero correlation between the asset value and its instantaneous volatility (the correlation was introduced later by Schobel–Zhu (1999)), which is important to model the so-called leverage effect observed in the marketplace. Equity market data imply negative correlation since most market participants hedge against large negative price jumps; the correlation is positive in commodity and energy markets since market participants hedge against big positive price jumps; finally, in currency and interest rate markets it is close to zero.

General references for the original HM and its generalizations (which are typically referred to as affine volatility models) include, for example, Carr–Madan (1999), Duffie *et al* (2000), Lewis (2000), Lipton (2001) and (2002), Sepp (2007). HM is popular among econophysicists, see for example Dragulescu–Yakovenko (2002), Silva *et al* (2004), Silva–Yakovenko (2003) and (2007), Perello *et al* (2004). References dealing with SSM (which is commonly referred to as a quadratic volatility model) include, for example, Schobel–Zhu (1999), Cheng *et al* (2004), Galluccio–Cam (2005), Cheng–Scaillet (2007).

In this paper, we will compare and contrast SSM and HM, and show that the former produces more robust market calibrations and is easier to deal with numerically. We derive the small-time asymptotics of the asset log-price density in SSM and show how to incorporate volatility jumps into the model. It turns out that volatility jumps are important to model skew effects observed for options on the realized volatility. We also study the distribution of the realized variance for SSM.

Our contribution is fourfold. First, we introduce volatility jumps in SSM and obtain the corresponding closed-form solution for the density of the asset price and its realized volatility. Second, we obtain the short-time asymptotics for the density of the asset log-price in SSM. Third, we illustrate both SSM and HM by calibrating them to fit historical and implied market data for the S&P 500 index (SPX) and the General Motors stock (GM). We argue that even for vanilla options SSM is at least as good as HM. Finally, we show that contemporary methods for solving affine pricing problems with stochastic volatility can be traced to the Kelvin method designed in the 19th century for studying wave motion problems arising in fluid dynamics.

1.2. Mathematical formulation

For option pricing and risk-management purposes we model the uncertainty associated with the random evolution of an asset price by the following SDE:

$$\mathrm{d}S(t) = \mu(t, S(t), \omega)\,\mathrm{d}t + \sigma(t, S(t), \omega)\,\mathrm{d}W_1(t), \tag{1}$$

where $W_1(t)$ is a standard Brownian motion, and $\mu(t, S, \omega)$ and $\sigma(t, S, \omega)$ represent the drift and volatility of the asset price process, respectively. We assume that they depend on time, the asset price $S(t)$ and a realization of the random variable ω.

To ensure the stability of the pricing problem we naturally require that the first and the second moments of the asset price are finite:

$$\begin{aligned} &\mathbb{E}^{\mathbb{Q}}\left[\int_0^t |\mu(t', S(t'), \omega)|\,\mathrm{d}t'\right] < \infty, \\ &\mathbb{E}^{\mathbb{Q}}\left[\int_0^t \sigma^2(t', S(t'), \omega)\,\mathrm{d}t'\right] < \infty, \qquad 0 \leqslant t < \infty, \end{aligned} \tag{2}$$

where the expectation is taken over all possible realizations of ω under the pricing measure $\mathbb{Q}$.

Because of their relative simplicity, we prefer to work with log-normal models for the asset price SDE which assume that the asset drift and volatility are proportional to $S(t)$:

$$\mathrm{d}S(t) = \mu(t)S(t)\,\mathrm{d}t + \sigma(t, \omega)S(t)\,\mathrm{d}W_1(t). \tag{3}$$

Under assumption (2) and conditional on the realization of ω, SDE (3) has the following exponential solution:

$$S(t) = S(0)\,\mathrm{e}^{\int_0^t \mu(t')\,\mathrm{d}t' - \frac{1}{2}I(t) + Y(t)}, \tag{4}$$

where

$$I(t) = \int_0^t \sigma^2(t', \omega)\,\mathrm{d}t', \qquad Y(t) = \int_0^t \sigma(t', \omega)\,\mathrm{d}W_1(t'), \tag{5}$$

$I(t)$ is the realized variance, and $Y(t)$ is the realized uncertainty factor, respectively.

It is important to note that we do not have any restrictions on the signs of $Y(t)$ and $\sigma(t)$. It is often stated or assumed that $\sigma(t)$ has to be positive. However, as we see from SDE (4) it is not necessary to impose this restriction. By omitting it, we can broaden our choice of possible volatility models as long as (2) is satisfied.

For robust modeling of the volatility dynamics, we assume that the asset volatility is a function of a stochastic variable $Z(t)$, $\sigma(t, \omega) \to \sigma(t, Z(t))$, which is driven by the following SDE:

$$\mathrm{d}Z(t) = \theta(t, Z(t))\,\mathrm{d}t + \varepsilon(t, Z(t))\,\mathrm{d}W_2(t). \tag{6}$$

Typically, we assume that the volatility process is mean-reverting and has a steady-state distribution. We also assume that Brownian motions $W_1(t)$ and $W_2(t)$ are correlated with a constant correlation parameter ρ.

Next, we introduce two independent Brownian motions $w_1(t)$ and $w_2(t)$ and present SDEs (3) and (6) as follows:

$$\begin{aligned} \mathrm{d}S(t) &= \mu(t)S(t)\,\mathrm{d}t + \sigma(t, Z(t))S(t)\left(\sqrt{1-\rho^2}\,\mathrm{d}w_1(t) + \rho\,\mathrm{d}w_2(t)\right), \\ \mathrm{d}Z(t) &= \theta(t, Z(t))\,\mathrm{d}t + \varepsilon(t, Z(t))\,\mathrm{d}w_2(t). \end{aligned} \tag{7}$$

Conditioning on the path of $w_2(t)$ and applying the isometry and scaling property of the Brownian motion, we solve the SDE for $S(t)$ as follows:

$$S(t) = S(0)\,\mathrm{e}^{\int_0^t \mu(t')\,\mathrm{d}t' - \frac{1}{2}I^*(t) + \sqrt{(1-\rho^2)I^*(t)}w_1(1) + \rho Y^*(t)}, \tag{8}$$

with

$$I^*(t) = \int_0^t \sigma^2(t', Z(t'))\,\mathrm{d}t', \qquad Y^*(t) = \int_0^t \sigma(t', Z(t'))\,\mathrm{d}w_1(t').$$

Thus we are able to present the evolution of the asset price in terms of realizations of the random variables $w_1(1)$, $I^*(t)$ and $Y^*(t)$.

1.3. Vanilla option pricing

Now we apply the above framework for pricing options on the underlying asset whose price process is modeled by $S(t)$. In particular, we concentrate on the vanilla call (put) option which gives its holder the right but not the obligation to buy (sell) the asset at a pre-determined price at the contract maturity time. Under the risk-neutral pricing measure $\mathbb{Q}$, and assuming that the risk-free interest rate is deterministic, the call option value is computed as follows:

$$C(t, S; T, K) = \mathrm{e}^{-\int_t^T r(t')\,\mathrm{d}t'}\mathbb{E}^{\mathbb{Q}}[\max(S(T) - K, 0)|\mathcal{F}(t)], \tag{9}$$

where filtration $\mathcal{F}(t)$ contains all information available at time t and the risk-neutral asset drift is $\mu(t) = r(t) - \mathrm{d}(t)$. Here $r(t)$ is the risk-free interest rate, and $\mathrm{d}(t)$ is the instantaneous dividend yield.

First, we consider the famous BSMM which postulates that the volatility process is a deterministic function of time, so that $I(t) = \int_0^t \sigma^2(t')\,\mathrm{d}t'$, $Z(t) \equiv 0$ and $\rho \equiv 0$. Combining expressions (8) and (9) and computing the integral of the call payoff function with respect to the standard normal density, we obtain:

$$\begin{aligned} &C^{BS}(0, S, I; T, K) = \mathrm{e}^{-\int_0^T \mathrm{d}(t')\,\mathrm{d}t'} S\mathcal{N}(d_+) - \mathrm{e}^{-\int_0^T r(t')\,\mathrm{d}t'} K\mathcal{N}(d_-), \\ &d_\pm = \frac{\ln(S/K) + \left(\int_0^T \mu(t')\,\mathrm{d}t' \pm \frac{1}{2}I(T)\right)}{\sqrt{I(T)}}, \end{aligned} \tag{10}$$

where $\mathcal{N}(x)$ is the distribution density function (CDF) of the standard normal random variable.

Under the more general setup (8), we note that the effective asset variance is $(1-\rho^2)I^*(T)$ and the effective asset price is $S\,\mathrm{e}^{-\frac{1}{2}\rho^2 I^*(t)+\rho Y^*(T)}$, and obtain (Romano–Touzi (1997) and Willard (1997)):

$$C(0, S; T, K) = \mathbb{E}^{\mathbb{Z}}\left[C^{BS}(0, S\,\mathrm{e}^{-\frac{1}{2}\rho^2 I^*(T)+\rho Y^*(T)}, (1-\rho^2)I^*(T); T, K)\right], \tag{11}$$

where the expectation $\mathbb{E}^{\mathbb{Z}}$ is computed over all realizations of the volatility process $Z(t)$.

2. Kelvin waves and option pricing via the Fourier transform

In this section, we discuss the analogy between the Kelvin waves arising in hydrodynamics and parabolic problems arising in option pricing. We note in passing that Darryl Holm was one of the pioneers of using Kelvin waves for studying the stability of complicated fluid flows, see, e.g., Bayly *et al* (1996).

2.1. Euler equations and Kelvin waves

The Euler equations describing the motion of an inviscid, incompressible fluid in a three-dimensional domain are given by

$$\frac{\partial \mathbf{V}}{\partial t} + (\mathbf{V}\cdot\nabla)\mathbf{V} - \nabla P = 0, \qquad \nabla\cdot\mathbf{V} = 0, \tag{12}$$

where $\mathbf{V}(t, \mathbf{x})$ is the velocity vector of the fluid and $P(t, \mathbf{x})$ is the scalar pressure field.

We consider a linear solution to problem (12):

$$\mathbf{V}(t, \mathbf{x}) = \mathcal{L}(t)\mathbf{x}, \qquad P(t, \mathbf{x}) = \tfrac{1}{2}(\mathcal{M}(t)\mathbf{x}\cdot\mathbf{x}), \tag{13}$$

where the matrices $\mathcal{L}$, $\mathcal{M} = \mathcal{M}^*$ satisfy the following system of ordinary differential equations (ODEs):

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{L} + \mathcal{L}^2 - \mathcal{M} = 0, \qquad \mathrm{trace}(\mathcal{L}) = 0. \tag{14}$$

Now we introduce a perturbation $\mathbf{v}(t, \mathbf{x})$, $p(t, \mathbf{x})$ and simplify equations (12) by neglecting the nonlinear term in $\mathbf{v}(t, \mathbf{x})$ as follows:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathcal{L}\mathbf{x} \cdot \nabla)\mathbf{v} + \mathcal{L}\mathbf{v} - \nabla p = 0, \qquad \nabla \cdot \mathbf{v} = 0. \tag{15}$$

A remarkable observation due to Kelvin (1887) is that the above equation possesses a wave-like solution of the form

$$\mathbf{v}(t, \mathbf{x}) = \mathbf{a}(t) \exp\{\mathrm{i}\mathbf{k}(t) \cdot \mathbf{x}\}, \qquad p(t, \mathbf{x}) = \pi(t) \exp\{\mathrm{i}\mathbf{k}(t) \cdot \mathbf{x}\}, \tag{16}$$

where $\mathrm{i} = \sqrt{-1}$ and $\mathbf{a}(t)$, $\mathbf{k}(t)$ satisfy the following system of ODEs:

$$\begin{aligned} &\frac{\mathrm{d}\mathbf{k}}{\mathrm{d}t} + \mathcal{L}^*\mathbf{k} = 0, \\ &\frac{\mathrm{d}\mathbf{a}}{\mathrm{d}t} + \mathcal{L}\mathbf{a} - 2\frac{\mathcal{L}\mathbf{a} \cdot \mathbf{k}}{\mathbf{k} \cdot \mathbf{k}} = 0, \\ &\mathbf{k} \cdot \mathbf{a} = 0, \end{aligned}$$

and $\pi(t)$ can be found from the incompressibility condition.

Equally remarkable observation was made by Chandrasekhar (1961), who noted, that due to incompressibility, the combined solution $\mathbf{V} + \mathbf{v}$, $P + p$ solves the system of Euler equations exactly, since the corresponding nonlinear term $(\mathbf{v} \cdot \nabla)\mathbf{v}$ vanishes identically.

Below we show that Kelvin waves have natural analogs in mathematical finance. We first show how to apply the Kelvin ansatz (16) in the BSMM framework and then apply the same technique in the stochastic volatility context.

2.2. *Option pricing via the fourier transform*

We consider option pricing in the BSMM and analyze the European option on a certain asset with price $S(t)$, with payoff at maturity time T has the form $u(S(T))$. At time t the option value, which is denoted by $\overline{U}(t, S)$, can be computed by analogy with equation (9). Applying the Feynman–Kac formula we represent the pricing problem for $\overline{U}(t, S)$ as follows:

$$\begin{aligned} &\overline{U}_t + \tfrac{1}{2}\sigma^2(t)S^2\overline{U}_{SS} + \mu(t)S\overline{U}_S - r(t)\overline{U} = 0, \\ &\overline{U}(T, S) = u(S). \end{aligned} \tag{17}$$

Introducing the logarithm of the drift-adjusted asset price normalized by a constant K and the modified value function:

$$\begin{aligned} &S \to X = \ln\frac{S}{K} + \int_t^T \mu(t')\,\mathrm{d}t', \qquad t \to I = \int_t^T \sigma^2(t')\,\mathrm{d}t', \\ &\overline{U}(I, S) \to U(t, X) = \mathrm{e}^{-\frac{1}{2}X + \int_t^T r(t')\,\mathrm{d}t'}\overline{U}(t, S), \end{aligned} \tag{18}$$

we observe that $U(I, X)$ solves the following backward Kolmogoroff equation:

$$U_I = \tfrac{1}{2}U_{XX}, \qquad U(0, X) = \mathrm{e}^{-\frac{1}{2}X}u(K\mathrm{e}^X). \tag{19}$$

Now we look for the Green function (or the state price density) of X which is denoted by $G(I, X)$ and which satisfies equation (19) supplied with the terminal condition $G(0, X) = \delta(X - X')$, where δ is the Dirac delta function, by applying the Fourier transform in X:

$$\widehat{G}(I, k) = \int_{-\infty}^{\infty} \mathrm{e}^{-\mathrm{i}kX'}G(I, X')\,\mathrm{d}X', \tag{20}$$

where $k \in \mathbb{R}$.

Performing straightforward calculations we obtain:

$$\widehat{G}(I,k) = \mathrm{e}^{-\mathrm{i}kX-\frac{1}{2}k^2 I}. \tag{21}$$

Now we can solve the backward problem (19) by applying the inverse Fourier transform:

$$U(I,X) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \Re[\check{u}(k)\widehat{G}(I,k)]\,\mathrm{d}k = \frac{1}{\pi}\int_{0}^{\infty} \Re[\check{u}(k)\widehat{G}(I,k)]\,\mathrm{d}k, \tag{22}$$

where $\check{u}(k)$ is the inverse Fourier transform of the payoff function:

$$\check{u}(k) = \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{i}kX'}\mathrm{e}^{-\frac{1}{2}X'}u(K\mathrm{e}^{X'})\,\mathrm{d}X'. \tag{23}$$

For the call option payoff function with strike K, formula (23) becomes:

$$\check{u}(k) = \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{i}kX'}\mathrm{e}^{-\frac{1}{2}X'}\max(K\mathrm{e}^{X'}-K,0)\,\mathrm{d}X' = \frac{K}{k^2+\frac{1}{4}}. \tag{24}$$

In the general setup we proceed by analogy: (1) we write the Green function for X through the backward equation akin to (19); (2) we find the solution for the transformed Green function (21); (3) we compute the transformed option payoff (23); (4) we invert formula (22) by means of the Gaussian quadratures or the FFT method. Details can be found in Lewis (2000), Lipton (2001) and (2002), where the general Lewis–Lipton formula is derived.

3. SSM for stochastic volatility

Stein–Stein (1991) assume that the asset volatility is driven by an OUP:

$$\mathrm{d}\sigma(t) = \kappa(\theta-\sigma(t))\,\mathrm{d}t + \varepsilon\mathrm{d}W_2(t), \qquad \sigma(0)=\sigma, \tag{25}$$

where θ is the long-term mean volatility, κ is the mean-reversion rate to the long-term mean and ε is the volatility of the volatility.

Solving SDE (25), we find that the distribution of volatility at time $\tau > 0, \sigma(\tau) \equiv \sigma'$, conditional on $\sigma(0) \equiv \sigma$ is Gaussian:

$$G^{\sigma}(\tau,\sigma,\sigma') = \frac{1}{\sqrt{2\pi\overline{m}_2^{\sigma}(\tau)}}\,\mathrm{e}^{-\frac{(\sigma'-\overline{m}_1^{\sigma}(\tau))^2}{2\overline{m}_2^{\sigma}(\tau)}}, \tag{26}$$

with the mean and variance parameters given respectively by

$$\overline{m}_1^{\sigma}(\tau) = \mathrm{e}^{-\kappa\tau}\sigma + (1-\mathrm{e}^{-\kappa\tau})\theta, \qquad \overline{m}_2^{\sigma}(\tau) = \frac{\varepsilon^2}{2\kappa}(1-\mathrm{e}^{-2\kappa\tau}). \tag{27}$$

To solve the pricing problem we also have to study the evolution of the variance. For this purpose, we augment SDE (25) with the SDE driving the variance process by means of applying Ito's lemma to function $V = \sigma^2$:

$$\mathrm{d}V(t) = 2\kappa\left(\frac{\varepsilon^2}{2\kappa} + \theta\sigma(t) - V(t)\right)\mathrm{d}t + 2\sigma(t)\varepsilon\,\mathrm{d}W_2(t), \qquad V(0)=\sigma^2. \tag{28}$$

Using the relationship between the normal and χ-squared distributions, we can show that the distribution of the variance at time τ, $V(\tau) = \sigma^2(\tau)$, conditional on $\sigma(0)$, is χ-squared with non-centrality parameter $\vartheta(\tau) = \left(\overline{m}_1^{\sigma}(\tau)\right)^2/\overline{m}_2^{\sigma}(\tau)$ and one degree of freedom:

$$G^{V}(\tau,V,V') = \frac{1}{2}\,\mathrm{e}^{-\frac{V'+\vartheta(\tau)}{2}}\left(\frac{V'}{\vartheta(\tau)}\right)^{-\frac{1}{4}}\mathcal{I}_{-\frac{1}{2}}(\sqrt{\vartheta(\tau)V'}), \tag{29}$$

where $\mathcal{I}_d$ is the modified Bessel function of order d.

We also note that under the SSM the instantaneous correlation between the asset price and its volatility depends on the sign of the instantaneous volatility while the instantaneous correlation between the asset price and its variance does not, so that the leveraging effect between the asset price and its instantaneous variance is always preserved.

3.1. Transformed log-price density

To solve the pricing problem efficiently we find the state density function, denoted by $G(\tau, \sigma, V, X, X')$, of the asset log-price X, as defined by (18) and (19). Here, since we assume time-homogeneous model parameters, we use the time-to-maturity variable $\tau = T - t$. Then $G(\tau, \sigma, V, X, X')$ describes the probability of X reaching the value of X' at time τ and contains both σ and V as independent variables; its solves the following backward Kolmogoroff equation:

$$\begin{aligned} -G_\tau + 2\kappa\left(\frac{\varepsilon^2}{2\kappa} + \theta\sigma - V\right)G_V + \kappa(\theta - \sigma)G_\sigma + \frac{1}{2}VG_{XX} + 2\varepsilon^2 VG_{VV} \\ + \frac{1}{2}\varepsilon^2 G_{\sigma\sigma} + 2\rho\varepsilon VG_{XV} + \rho\varepsilon\sigma G_{X\sigma} + 2\varepsilon^2\sigma G_{\sigma V} = 0, \\ G(0, \sigma, V, X, X') = \delta(X - X'). \end{aligned} \tag{30}$$

Applying the exponential ansatz similar to Kelvin's we obtain:

$$\widehat{G}(\tau, \sigma, V, X, k) = \mathrm{e}^{-\mathrm{i}kX + A(\tau) + B(\tau)V + C(\tau)\sigma}, \tag{31}$$

where functions $A(\tau)$, $B(\tau)$, $C(\tau)$ solve the following system of ODEs:

$$\begin{aligned} -\dot{A} + \varepsilon^2 B + \tfrac{1}{2}\varepsilon^2 C^2 + \kappa\theta C = 0, \qquad & A(0) = 0, \\ -\dot{B} + 2\varepsilon^2 B^2 - 2(\kappa + \rho\varepsilon\mathrm{i}k)B - \tfrac{1}{2}k^2 = 0, \qquad & B(0) = 0, \\ -\dot{C} - (\kappa + \rho\varepsilon\mathrm{i}k)C + (2\varepsilon^2 C + 2\kappa\theta)B = 0, \qquad & C(0) = 0, \end{aligned} \tag{32}$$

and the dot stands for the derivative with respect to time variable τ.

Although the solution of the system of ODEs (31) can be expressed in closed-form, it is both too cumbersome and numerically expensive to deal with. Instead we solve equation (32) by applying the standard Runge–Kutta method, which can also be used when the model parameters are deterministic functions of time. For numerical inversion it is important to study the asymptotic behavior of the corresponding functions as $|k| \to \infty$. First looking at the ODE for B, we see that asymptotically, for large values of $|k|$, $B \sim \mathrm{i}Lk$, where L is a constant. Plugging this expression into the ODE and balancing terms of order k^2, we get a quadratic equation for L:

$$\varepsilon^2 L^2 - \rho\varepsilon L + \tfrac{1}{4} = 0, \tag{33}$$

so that, $L = (\rho \pm \mathrm{i}\sqrt{1 - \rho^2})/(2\varepsilon)$.

Next we find that asymptotically, for large values of $|k|$, C is a constant and $A \sim \mathrm{i}\tau\varepsilon^2 Lk$. Accordingly, the leading term in the real-valued part of $\widehat{G}$ becomes:

$$\Re[\widehat{G}(\tau, \sigma, V, X, k)] \sim \mathrm{e}^{-\frac{\tau\varepsilon^2 + V}{2\varepsilon}\sqrt{1-\rho^2}|k|}, \qquad |k| \to \infty. \tag{34}$$

Finally, for pricing options on X we apply formula (22) with $\widehat{G}$ computed by using expression (31).

3.2. Steady-state and small-time density

First we consider the small-time ($\tau \to 0$) volatility density which is normal with mean $\sigma(0) + \tau\kappa(\theta - \sigma(0))$ and variance $\varepsilon^2\tau$. Second, we consider the steady state ($\tau \to \infty$) volatility density which is normal with mean θ and variance $\varepsilon^2/(2\kappa)$. We note that the mean-reversion coefficient impacts the mean of the process in the short-term and its volatility in the long-term. Next, we find the transformed density function for $\sigma(\tau)$ and $V(\tau)$ conditioned on $\sigma(0)$:

$$\begin{aligned} Z(\Sigma, \Theta) &= \frac{1}{\sqrt{2\pi\overline{m}_2^\sigma(\tau)}} \int_{-\infty}^{\infty} \mathrm{e}^{-\Sigma\sigma(\tau)-\Theta V(\tau)-\frac{(\sigma(\tau)-\overline{m}_1^\sigma(\tau))^2}{2\overline{m}_2^\sigma}} \mathrm{d}\sigma(\tau) \\ &= \frac{1}{\sqrt{2\Theta\overline{m}_2^\sigma(\tau)+1}} \mathrm{e}^{\frac{-(\overline{m}_1^\sigma(\tau))^2\Theta-\overline{m}_1^\sigma(\tau)\Sigma+\frac{1}{2}\overline{m}_2^\sigma(\tau)\Sigma^2}{2\Theta\overline{m}_2^\sigma(\tau)+1}}, \end{aligned} \tag{35}$$

where $\Sigma, \Theta \in \mathbb{C}$ with $\Re[\Theta] > -1/\left(2\overline{m}_2^\sigma(\tau)\right)$.

Now we analyze the short-term density of the asset log-price X. This short-term density function is the key factor for fitting the model to the historical time series of, say, daily log-returns and calculating Q–Q plots. For brevity, we assume zero correlation between the spot price and its volatility and zero drift.

The integral representation of the log-price density (without applying the shift by $\mathrm{e}^{-X/2}$) conditioned on $I(\tau)$ is given by

$$G(\tau, X, X') = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Re\left[\mathrm{e}^{\mathrm{i}k(X'-X)+\frac{1}{2}\left(-k^2+\mathrm{i}k\right)I(\tau)}\right] \mathrm{d}k, \tag{36}$$

where τ is assumed to be small.

Now we approximate $I(\tau)$ by $I(\tau) \approx \frac{1}{2}(\sigma^2(0) + \sigma^2(\tau))\tau$ and assume that $\sigma(0)$ has a stationary distribution (a similar idea is used by Dragulescu and Yakovenko (2002)). We can represent integral (36) as follows:

$$\begin{aligned} G(\tau, X, X') = {} & \frac{\sqrt{2\kappa}}{4\pi^2\sqrt{\varepsilon^2\overline{m}_2^\sigma(\tau)}} \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \\ & \times \Re\left[\mathrm{e}^{\mathrm{i}k(X'-X)+\frac{1}{4}\left(-k^2+\mathrm{i}k\right)\left(\sigma^2(0)+\sigma^2(\tau)\right)\tau-\frac{\left(\sigma(\tau)-\overline{m}_1^\sigma(\tau)\right)^2}{2\overline{m}_2^\sigma(\tau)}-\frac{\kappa(\sigma(0)-\theta)^2}{\varepsilon^2}}\right] \mathrm{d}k\, \mathrm{d}\sigma(\tau)\, \mathrm{d}\sigma(0). \end{aligned} \tag{37}$$

Exchanging integration order and applying formula (35) twice, we obtain:

$$G(\tau, X, X') = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Re\left[\frac{\exp\left\{\mathrm{i}k(X'-X) - c - \frac{-\theta^2 a-\theta b+\varepsilon^2 b^2/(4\kappa)}{\sqrt{\varepsilon^2 a/\kappa+1}}\right\}}{\sqrt{\left(2\overline{m}_2^\sigma(\tau)q+1\right)\left(\varepsilon^2 a/\kappa+1\right)}}\right] \mathrm{d}k, \tag{38}$$

where

$$\begin{aligned} q &= -\frac{1}{4}(-k^2+\mathrm{i}k)\tau, & a &= q\left(\frac{\mathrm{e}^{-2\tau\kappa}}{2\overline{m}_2^\sigma(\tau)q+1}+1\right), \\ b &= q\left(\frac{2\theta\,\mathrm{e}^{-\tau\kappa}(1-\mathrm{e}^{-\tau\kappa})}{2\overline{m}_2^\sigma(\tau)q+1}\right), & c &= q\left(\frac{\theta^2(1-\mathrm{e}^{-\tau\kappa})^2}{2\overline{m}_2^\sigma(\tau)q+1}\right). \end{aligned} \tag{39}$$

4. HM for stochastic variance

We now consider HM, where we model the dynamics of the asset variance:

$$\mathrm{d}V(t) = \kappa(\theta^2 - V(t))\,\mathrm{d}t + \varepsilon\sqrt{V(t)}\,\mathrm{d}W_2(t), \qquad V(0) = V. \tag{40}$$

Here θ^2 is the long-term mean variance, κ is the reversion rate to the long-term mean and ε is the volatility of the variance.

We can show (see, for example, Lipton (2001)) that the distribution of $V(\tau) \equiv V'$ conditional on $V(0) \equiv V$ is the χ-squared with PDF:

$$G^V(\tau, V, V') = \overline{M}\,\mathrm{e}^{-\overline{M}(\mathrm{e}^{-\kappa\tau}V + V')}\left(\frac{V'}{\mathrm{e}^{-\kappa\tau}V}\right)^{\frac{\vartheta}{2}} \mathcal{I}_\vartheta(2\overline{M}\sqrt{\mathrm{e}^{-\kappa\tau}VV'}), \tag{41}$$

where $\overline{M} = 2\kappa/(\varepsilon^2(1 - \mathrm{e}^{-\kappa\tau}))$, and $\vartheta = 2\kappa\theta^2/\varepsilon^2 - 1$. According to the Feller test, if $\vartheta \geqslant 0$ the origin is a natural boundary for V while in the opposite case the origin is accessible by V.

Applying Ito's lemma, we obtain the following process for the volatility $\sigma = \sqrt{V}$ for HM:

$$\mathrm{d}\sigma(t) = \frac{\kappa}{2}\left(\frac{\theta^2 - \varepsilon^2/(4\kappa)}{\sigma(t)} - \sigma(t)\right)\mathrm{d}t + \frac{\varepsilon}{2}\,\mathrm{d}W_2(t), \tag{42}$$

from which we see that the asset volatility stays positive.

The distribution of volatility $\sigma(\tau) \equiv \sigma'$ conditional on $\sigma(0) \equiv \sigma$ for HM can be derived using PDF (41):

$$G^\sigma(\tau, \sigma, \sigma') = 2\overline{M}\,\mathrm{e}^{-\overline{M}(\mathrm{e}^{-\kappa\tau}\sigma^2 + \sigma'^2)}\left(\frac{\sigma'}{\mathrm{e}^{-\frac{1}{2}\kappa\tau}\sigma}\right)^{\vartheta} \mathcal{I}_\vartheta\left(2\overline{M}\,\mathrm{e}^{-\frac{1}{2}\kappa\tau}\sigma\sigma'\right)\sigma'. \tag{43}$$

4.1. Transformed log-price density

The backward equation for the log-price density corresponding to (19), which is denoted by $G(\tau, V, X, X')$, is given by

$$\begin{aligned} &-G_\tau + \kappa(\theta^2 - V)G_V + \tfrac{1}{2}VG_{XX} + \rho\varepsilon VG_{XV} + \tfrac{1}{2}\varepsilon^2 VG_{VV} = 0, \\ &G(0, V, X, X') = \delta(X - X'). \end{aligned} \tag{44}$$

Applying the Kelvin ansatz, we obtain

$$\widehat{G}(\tau, V, X, k) = \mathrm{e}^{-\mathrm{i}kX + A(\tau) + B(\tau)V}, \tag{45}$$

where

$$\begin{aligned} &-\dot{A} + \kappa\theta^2 B = 0, && A(0) = 0, \\ &-\dot{B} + \tfrac{1}{2}\varepsilon^2 B^2 - (\kappa + \rho\varepsilon\mathrm{i}k)B - \tfrac{1}{2}k^2 = 0, && B(0) = 0. \end{aligned} \tag{46}$$

Explicitly solving the above system yields:

$$\begin{aligned} A(\tau) &= -\frac{2\kappa\theta^2}{\varepsilon^2}\left(\frac{1}{2}\psi_+\tau + \ln(C_+\,\mathrm{e}^{-\zeta\tau} + C_-)\right), \\ B(\tau) &= -\frac{-\psi_- C_+\,\mathrm{e}^{-\zeta\tau} + \psi_+ C_-}{\varepsilon^2(C_+\,\mathrm{e}^{-\zeta\tau} + C_-)}, \end{aligned} \tag{47}$$

where

$$C_\pm = \psi_\pm/(2\zeta), \qquad \psi_\pm = \mp(\kappa + \rho\varepsilon\mathrm{i}k) + \zeta, \qquad \zeta = \sqrt{(\kappa + \rho\varepsilon\mathrm{i}k)^2 + \varepsilon^2 k^2}. \tag{48}$$

Following the same steps as for SSM, we find that asymptotically for large values of $|k|$ the leading real-valued part in formula (45) becomes:

$$\Re[\widehat{G}(\tau, V, X, k)] \sim \mathrm{e}^{-\frac{\tau\kappa\theta^2+V}{\varepsilon}\sqrt{1-\rho^2}|k|}, \qquad |k| \to \infty. \tag{49}$$

4.2. Steady-state and small-time density

First we find the steady-state density of the variance. For this purpose, we recall the following limit for the modified Bessel function:

$$\mathcal{I}_\vartheta(x) = \frac{1}{\Gamma(\vartheta+1)}\left(\frac{x}{2}\right)^\vartheta, \qquad 0 < x \ll \sqrt{\vartheta+1}, \tag{50}$$

where $\Gamma(x)$ is the gamma function.

Letting $\tau \to \infty$ and applying the above result to expression (41), we conclude that the steady-state density has the gamma distribution:

$$G^{V^\infty}(V') = \frac{(V')^{\beta-1}\,\mathrm{e}^{-\frac{V'}{\alpha}}}{\alpha^\beta\Gamma(\beta)}, \tag{51}$$

with scale parameter $\alpha = \varepsilon^2/(2\kappa)$, and shape parameter $\beta = 2\kappa\theta^2/\varepsilon^2$.

We also apply limit (50) to analyze the behavior of the density function of the variance (41) for small V':

$$G^V(\tau, V, V') = \frac{\overline{M}\left(\overline{M}V'\right)^\vartheta}{\Gamma(\vartheta+1)}\,\mathrm{e}^{-\overline{M}(\mathrm{e}^{-\kappa\tau}V+V')}. \tag{52}$$

As a result, we see that if $\vartheta < 0$, that is, if the Feller condition is not satisfied, the density diverges as V' approaches zero.

Applying the same analysis for equation (43) we find that the steady-state density of the volatility has the following distribution:

$$G^\sigma(\tau, \sigma, \sigma') = 2\frac{(\sigma')^{2\beta-1}\,\mathrm{e}^{-\frac{(\sigma')^2}{\alpha}}}{\alpha^\beta\Gamma(\beta)}. \tag{53}$$

When $\beta = 1$ this is the density of the Rayleigh distribution with volatility parameter $\theta/2$.

For small values of σ' we use formula (50) and obtain:

$$G^\sigma(\tau, \sigma, \sigma') = 2\frac{\overline{M}^{\vartheta+1}(\sigma')^{2\vartheta+1}}{\Gamma(\vartheta+1)}\,\mathrm{e}^{-\overline{M}(\mathrm{e}^{-\kappa\tau}\sigma^2+\sigma'^2)}, \tag{54}$$

so that the volatility density diverges if $4\kappa\theta^2/\varepsilon^2 - 1 < 0$, which is stronger that the Feller condition for the variance process.

Next we find the transformed density function of $V(\tau)$ conditioned on $V(0)$:

$$Z(\tau, \Theta) = \int_0^\infty \mathrm{e}^{-\Theta V'} G^V(\tau, V, V')\,\mathrm{d}V' = \mathrm{e}^{A(\tau,\Theta)+B(\tau,\Theta)V}, \tag{55}$$

where $\Theta \in \mathbb{C}$ and

$$\begin{aligned} A(\tau, \Theta) &= -\frac{2\kappa\theta^2}{\varepsilon^2}\ln\left(\frac{\varepsilon^2}{2\kappa}\Theta(1-\mathrm{e}^{-\kappa\tau})+1\right), \\ B(\tau, \Theta) &= -\frac{\Theta\,\mathrm{e}^{-\kappa\tau}}{\frac{\varepsilon^2}{2\kappa}\Theta(1-\mathrm{e}^{-\kappa\tau})+1}. \end{aligned} \tag{56}$$

Taking the limit $\tau \to \infty$ in formula (55), we obtain:

$$Z^{\infty}(\Theta) = \mathrm{e}^{A^{\infty}(\Theta)}, \qquad A^{\infty}(\Theta) = -\frac{2\kappa\theta}{\varepsilon^2} \ln\left(\frac{\varepsilon^2}{2\kappa}\Theta + 1\right). \tag{57}$$

To analyze the short-term density of the log-price X we apply formula (36) and use the approximation $I(\tau) \approx \frac{1}{2}(V(0) + V(\tau))\tau$. Then assuming that $V(0)$ has a stationary distribution, we represent expression (36) as follows:

$$G(\tau, X, X') = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \times \Re\left[\mathrm{e}^{\mathrm{i}k(X'-X) - \frac{1}{4}(-k^2+ik)(V(0)+V(\tau))\tau} G^{V}(\tau, V(0), V(\tau)) G^{V^{\infty}}(V(0))\right] \mathrm{d}k \, \mathrm{d}V(\tau) \, \mathrm{d}V(0). \tag{58}$$

Exchanging the integration order and applying formulae (55) and (57), we obtain:

$$G(\tau, X, X') = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Re\left[\mathrm{e}^{\mathrm{i}k(X'-X) + A(\tau,q) + A^{\infty}(-B(\tau,q)+q)}\right] \mathrm{d}k, \tag{59}$$

where $q = \frac{1}{4}(-k^2 + \mathrm{i}k)\tau$ and $A(\tau, q)$ and $B(\tau, q)$ are defined by (56).

5. Realized asset variance and volatility jumps

Now we briefly consider financial contracts on the realized asset variance which derive their values from the annualized variance of the asset spot price $S(t)$ realized over the time period from t_0 to t_N:

$$I_N(t_0, t_N) = \frac{AF}{N} \sum_{n=1}^{N} \left(\ln \frac{S(t_n)}{S(t_{n-1})}\right)^2. \tag{60}$$

Here $S(t_n)$ is the asset price observed at times t_0 (contract inception) (maturity), N is the number of observations from contract inception up to its expiration, $\ln(S(t_n)/S(t_{n-1}))$ is the return realized over the time period between t_{n-1} and t_n, and AF is the annualization factor.

In the continuous-time setting, assuming that the asset price follows a diffusion process with stochastic variance and without price jumps, $I_N(t_0, T)$ can be very accurately approximated by its continuous time limit denoted by $\hat{I}(t_0, T)$

$$\hat{I}(T) = \frac{1}{T - t_0} \int_{t_0}^{T} V(t') \, \mathrm{d}t' = \frac{1}{T - t_0} I(t_0, T), \tag{61}$$

where $V(t)$ is the instantaneous variance of assets returns, $\hat{I}(t_0, T)$ is the cumulative annualized variance realized over the period (t_0, T), $I(t_0, T)$ is the cumulative de-annualized variance. Here, we take into account the fact that the number of fixings is typically proportional to the annualization factor: $N \sim (T - t_0)AF$.

We refer to the paper by Demeterfi *et al* (1999) for the financial intuition behind contracts on the realized asset volatility. In short, these products are widely used for both hedging and speculation purposes. Particularly popular are swaps on the realized asset variance, swaps with cap and floor levels, and options on these swaps. To value and risk-manage these contracts we need to study the distribution of the realized asset variance $I(t)$. For this purpose, we augment the SDE driving the stochastic variance by the SDE driving $I(t)$ based on expression (61). More details on the augmentation technique can be found in Lipton (2001).

Finally, we mention that when modeling the realized asset variance, it is extremely important to account for the jumps in the realized asset volatility. While in principle this goal

can be achieved by including either asset jumps or volatility jumps, the volatility jumps have higher impact on the forward-start options and contracts on the future realized asset variance (VIX-type contracts). Also, it turns out that it is harder to estimate asset price jumps rather than volatility jumps, so that we concentrate on modeling volatility jumps. Although a variety of distributions can be applied to model the volatility jumps (especially in SSM), we prefer to work with the most simple distribution where the volatility jump distribution is the delta function centered at a specified jump amplitude.

5.1. SSM for the volatility dynamics

The augmented system of SDEs for σ, I is given by

$$\begin{aligned} &\mathrm{d}\sigma(t) = \kappa(\theta - \sigma(t))\,\mathrm{d}t + \varepsilon\,\mathrm{d}W(t) + J\,\mathrm{d}N(t), \qquad \sigma(0) = \sigma, \\ &\mathrm{d}I(t) = \sigma^2(t)\,\mathrm{d}t, \qquad I(0) = I, \end{aligned} \tag{62}$$

where $N(t)$ is a Poisson process with intensity γ and J is a deterministic jump with the amplitude η. In other words, upon arrival of a jump of the Poisson process $N(t)$, the asset volatility jumps by η.

Now we consider the backward Kolmogoroff equation for the state density of I:

$$\begin{aligned} -G_\tau + \kappa(\theta - \sigma)G_\sigma + 2\kappa\left(\frac{\varepsilon^2}{2\kappa} + \theta\sigma - V\right)G_V + \frac{1}{2}\varepsilon^2 G_{\sigma\sigma} + 2\varepsilon^2 V G_{VV} + 2\varepsilon^2\sigma G_{\sigma V} \\ + V G_I + \gamma(G(\sigma + J, V + 2\sigma J + J^2) - G) = 0, \\ G(0, \sigma, V, I, I') = \delta(I - I'). \end{aligned} \tag{63}$$

We introduce the Fourier transform of $G(I')$ with respect to I', denoted by $\widehat{G}(k)$, and apply the corresponding Kelvin ansatz:

$$\widehat{G}(\tau, \sigma, V, I, k) = \mathrm{e}^{-\mathrm{i}kI + A(\tau) + B(\tau)V + C(\tau)\sigma}, \tag{64}$$

and consider the term arising from the jump-part:

$$\begin{aligned} \gamma(\widehat{G}(\sigma + J, V + 2\sigma J + J^2) - \widehat{G}) &= \gamma\left(\mathrm{e}^{B(\tau)(2\sigma J + J^2) + C(\tau)J} - 1\right)\widehat{G} \\ &\approx \gamma\left(\mathrm{e}^{B(\tau)J^2 + C(\tau)J}(1 + 2B(\tau)J\sigma + 2B^2(\tau)J^2 V) - 1\right)\widehat{G}. \end{aligned} \tag{65}$$

We see that the jumps introduce nonlinear effects, which, strictly speaking, makes the problem non-analytic. However, expanding the exponent in a series with respect to σ, and keeping only terms of order $O(\sigma^2)$, we will be able to use our linear exponential ansatz. After some algebra, we obtain the following system of ODEs:

$$\begin{aligned} &-\dot{A} + \varepsilon^2 B + \tfrac{1}{2}\varepsilon^2 C^2 + \kappa\theta C + \gamma\left(\mathrm{e}^{\eta^2 B + \eta C} - 1\right) = 0, \qquad A(0) = 0, \\ &-\dot{B} + 2\varepsilon^2 B^2 - 2\kappa B - \mathrm{i}k + 2\gamma\eta^2 B^2\,\mathrm{e}^{\eta^2 B + \eta C} = 0, \qquad B(0) = 0, \\ &-\dot{C} - \kappa C + (2\varepsilon^2 C + 2\kappa\theta)B + 2\gamma\eta B\,\mathrm{e}^{\eta^2 B + \eta C} = 0, \qquad C(0) = 0. \end{aligned} \tag{66}$$

The above expansion implies that for the jump-part we linearize the impact of the cross term σJ while preserving the impact of the jump in σ and V. Numerical experiments reveal that this approximation is in excellent agreement with numerical solution based on discretization of PDE (63) using the Craig–Sneyd method which is described in the following section. The same procedure is applied when we solve for the density of the log-price under SSM with volatility jumps, in which case we add jump terms to the system of ODEs (32) similarly to what we do for the system (66).

By analogy with equation (33), we obtain that for large values of $|k|$ the leading real-valued part in formula (64) becomes:

$$\Re[\widehat{G}(\tau,\sigma,V,X,k)] \sim \mathrm{e}^{-(\frac{V}{\sqrt{2}\varepsilon}+\frac{\varepsilon\tau}{\sqrt{2}})\sqrt{|k|}}, \qquad |k|\to\infty. \tag{67}$$

Similarly, we find that the Fourier transform of the marginal density of $\sigma(t)$ is given by

$$\widehat{G}(\tau,\sigma,\Theta) = \mathrm{e}^{A(\tau)-\Theta\mathrm{e}^{-\kappa\tau}\sigma}, \tag{68}$$

with $\Theta\in\mathbb{C}$ and

$$A(\tau) = \frac{\varepsilon^2\Theta^2}{4\kappa}(1-\mathrm{e}^{-2\kappa\tau}) - \theta\Theta(1-\mathrm{e}^{-\kappa\tau}) + \gamma\mathcal{I}(\tau), \tag{69}$$

where

$$\mathcal{I}(\tau) = \int_0^\tau \mathrm{e}^{-\eta\Theta\mathrm{e}^{-\kappa t'}}\,\mathrm{d}t' - \tau = \frac{1}{\kappa}(E_1(\eta\Theta) - E_1(\eta\Theta\,\mathrm{e}^{-\kappa\tau})) - \tau, \tag{70}$$

and $E_1(z)$ is the exponential integral:

$$E_1(z) = \int_1^\infty \frac{\mathrm{e}^{-zt'}}{t'}\,\mathrm{d}t', \qquad \Re[z]\geqslant 0.$$

We see that the presence of jumps makes the problem more complex when the process is mean-reverting since in this case we have also to account for jump times. For $\kappa\to\infty$, we obtain:

$$\widehat{G}(\tau,\sigma,\Theta) = \mathrm{e}^{A(\tau)-\Theta\sigma}, \qquad A(\tau) = (\varepsilon^2\Theta^2 - \theta\Theta + \gamma(\mathrm{e}^{-\eta\Theta}-1))\tau, \tag{71}$$

so that the corresponding process is a Gaussian process augmented with the jump component. We also note that conditioned on the number of jumps occurred being l, the process is Gaussian with mean $\theta+\eta l$ and volatility ε, so that we can obtain its density by summing conditional densities over the Poisson distribution with the intensity parameter $\tau\gamma$.

Finally, we solve for the first moment of the realized asset variance, $\overline{I}(\tau) = \mathbb{E}[I(\tau)|I(0)]$. Function $\overline{I}(\tau)$ solves equation (63) supplied with the terminal condition $\overline{I}(\tau) = I$. Assuming the linear ansatz for the solution:

$$\overline{I}(\tau) = I(0) + A(\tau) + B(\tau)V + C(\tau)\sigma, \tag{72}$$

we obtain the following system of ODEs:

$$\begin{aligned} -\dot{A} + \kappa\theta C + m_2 B &= 0, & A(0) &= 0,\\ -\dot{B} - 2\kappa B + 1 &= 0, & B(0) &= 0,\\ -\dot{C} - \kappa C + 2m_1 B &= 0, & C(0) &= 0, \end{aligned} \tag{73}$$

where $m_1 = \kappa\theta+\gamma\eta$, and $m_2 = \varepsilon^2+\gamma\eta^2$. A simple integration yields:

$$\begin{aligned} A(\tau) &= \left(\frac{m_1^2}{\kappa^2}+\frac{m_2}{2\kappa}\right)\tau - \frac{m_1^2}{2\kappa^3}(1-\mathrm{e}^{-\kappa\tau})(3-\mathrm{e}^{-\kappa\tau}) - \frac{m_2}{4\kappa^2}(1-\mathrm{e}^{-2\kappa\tau}),\\ B(\tau) &= \frac{1}{2\kappa}(1-\mathrm{e}^{-2\kappa\tau}), \qquad C(\tau) = \frac{m_1}{\kappa^2}(1-\mathrm{e}^{-\kappa\tau})^2. \end{aligned} \tag{74}$$

The expected realized asset variance, $\overline{I}(\tau)$, serves as an indicator for the overall level of the expected asset variance and it is often referred to by market practitioners as the fair variance. Investors can realize their views on the expected fair variance by trading in variance swaps. For valuing nonlinear contracts on the realized asset variance we apply the general inversion formula (22) along with the transformed Green function (64).

5.2. HM for the variance dynamics

For HM the augmented system of SDEs becomes:

$$\begin{aligned} &\mathrm{d}V(t) = \kappa(\theta^2 - V(t))\,\mathrm{d}t + \varepsilon\sqrt{V(t)}\,\mathrm{d}W_2(t) + J\,\mathrm{d}N(t), \qquad V(0) = V, \\ &\mathrm{d}I(t) = V(t)\,\mathrm{d}t, \qquad I(0) = I, \end{aligned} \tag{75}$$

where the jump size in the variance has fixed amplitude $J = \eta^2$. The backward equation corresponding to (75) is given by

$$\begin{aligned} &-G_\tau + \kappa(\theta^2 - V)G_V + \tfrac{1}{2}\varepsilon^2 V G_{VV} + V G_I + \gamma(G(V+J) - G) = 0, \\ &G(0, V, I, I') = \delta(I - I'). \end{aligned} \tag{76}$$

Applying the Kelvin ansatz, we obtain

$$\widehat{G}(\tau, V, I, k) = \mathrm{e}^{-\mathrm{i}kI + A(\tau) + B(\tau)V}, \tag{77}$$

where

$$\begin{aligned} &-\dot{A} + \kappa\theta^2 B + \gamma(\mathrm{e}^{\eta^2 B} - 1) = 0, \qquad A(0) = 0, \\ &-\dot{B} + \tfrac{1}{2}\varepsilon^2 B^2 - \kappa B - \mathrm{i}k = 0, \qquad B(0) = 0. \end{aligned} \tag{78}$$

Explicitly solving the above system, we obtain:

$$\begin{aligned} A(\tau) &= -\frac{\kappa\theta^2}{\varepsilon^2}\left[\psi_+\tau + 2\ln\left(\frac{\psi_- + \psi_+\,\mathrm{e}^{-\zeta\tau}}{2\zeta}\right)\right] + \gamma\left(\int_0^\tau \mathrm{e}^{\eta^2 B(t')}\,\mathrm{d}t' - \tau\right), \\ B(\tau) &= -2\Psi\frac{1 - \mathrm{e}^{-\zeta\tau}}{\psi_- + \psi_+\,\mathrm{e}^{-\zeta\tau}}, \\ \psi_\pm &= \mp\kappa + \zeta, \qquad \zeta = \sqrt{\kappa^2 + 2\varepsilon^2\mathrm{i}k}. \end{aligned} \tag{79}$$

Asymptotically, the leading real-valued part of formula (79) for large values of $|k|$ is given by

$$\Re[\widehat{G}(\tau, V, X, k)] \sim \mathrm{e}^{-\frac{V + \kappa\theta^2\tau}{\sqrt{2}\varepsilon}\sqrt{|k|}}, \qquad |k| \to \infty. \tag{80}$$

Finally, we solve for the first moment of I, $\overline{I}(\tau) = \mathbb{E}[I(\tau)|I(0)]$, for HM. Assuming the linear solution and performing the same steps as for OUP we get:

$$\overline{I}(\tau) = I(0) + \left(\frac{\theta^2}{\kappa} + \frac{\gamma\eta^2}{\kappa^2}\right)(\tau\kappa + \mathrm{e}^{-\kappa\tau} - 1) + \frac{1}{\kappa}(1 - \mathrm{e}^{-\kappa\tau})V. \tag{81}$$

By analogy, we apply formula (22) along with equation (77) for valuing nonlinear contracts on the realized asset variance for HM.

6. Relationship between SSM and HM

Now we are in a position to compare the distribution of the variance implied by SSM and HM as given by formulae (29) and (41), respectively. Although both are specified by the non-central χ-squared distribution, these distributions differ in the general case. However, they do coincide in the special case when for SSM we restrict the mean volatility to be zero $\theta_s = 0$, and choose the mean-reversion κ_s and volatility of volatility ε_s to be free parameters, and then specify the mean variance θ_h^2, the mean-reversion κ_h and ε_h for HM to be $\theta_h^2 = \varepsilon_s^2/(2\kappa_s)$, $\kappa_h = 2\kappa_s$ and $\varepsilon_h = 2\varepsilon_s$, respectively.

Given that the correlation between the asset and its volatility are the same under both SSM and HM, the above specification of model parameters also produces identical distributions for the asset price distribution. However, it is important to note that even under the above

specification the two models imply different distributions for the asset volatility: in SSM it is unrestricted in sign and is given by expression (26) while in HM the positive root of the variance is given by expression (43).

As a result, we arrive at an interesting conclusion: while, under a specific set of model parameters, two models produce identical distributions for the asset price and its variance, they produce different distributions for the asset volatility. This effect is due to the fact that in HM the volatility is defined as the positive square root of the variance. However, the square root is a two-valued function, so that the above specification is ambiguous. On the other hand, if we model the asset volatility as a primary variable, we define the asset variance with the quadratic function which is single-valued and hence unambiguous.

This conclusion is important for designing the relevant numerical methods. If we use the variance as one of the variables, we have to restrict it to be positive by imposing appropriate boundary conditions (which are difficult to deal with). However, if we use the asset volatility as a variable (and the asset variance is modeled via the augmentation procedure), we only have to impose the far-field boundary conditions for the volatility, which are natural conditions and hence easy to deal with.

In the following section we will discuss the issues which we face when implementing SSM and HM via numerical methods.

7. Numerical solution methods

As we have already mentioned, obtaining a closed-form solution for vanilla option values is only the first step in using stochastic volatility models in practice. The next important and challenging step consists in applying robust numerical methods for valuing complex deals with multiple exercise features and payoffs during the contract life. At this step we will typically apply numerical PDE solutions and (or) MC simulations. We briefly consider the challenges we face in applying these methods.

7.1. PDE methods

We introduce the value function $U(t, X, Y, I)$, where X is the asset price or its logarithm, Y is the driving factor for the volatility or variance process, and I is the accrual function of the X and Y, for the derivative contract with the payoff function $u_1(X, Y, I)$ and the instantaneous reward function $u_2(t, X, Y, I)$. The value function $U(t, X, Y, I)$ satisfies the backward Kolmogoroff equation:

$$\begin{aligned} &U_t + \mu^x U_X + \mu^y U_Y + \xi^x U_{XX} + \xi^y U_{YY} + \xi^{xy} U_{XY} + \mu^I U_I + \gamma(U(Y+\eta) - U) \\ &\qquad = u_2(t, X, Y, I), \qquad , \\ &U(T, X, Y, I) = u_1(X, Y, I) \end{aligned} \tag{82}$$

where, μ and ξ are mean and variance functions for the respective variables and ξ^{xy} is their covariance.

To solve this PDE numerically we use the Craig–Sneyd scheme (1988) (also see McKee *et al* (1996) for an extension). We first introduce the discrete time grid $\{t^n\}$ and spacial grids in X, Y and I directions, assuming, for brevity, uniform step sizes Δt, ΔX, ΔY and ΔI. Leaving the treatment of boundary conditions aside for the time being, we approximate the diffusion–advection in X and Y directions by either Crank–Nicolson or backward Euler schemes and denote the corresponding transition matrices by $\mathcal{L}_X$ and $\mathcal{L}_Y$, respectively. We treat the jump-part in the Y-direction explicitly and denote it by $\mathcal{J}$. We then approximate the

correlation via an explicit matrix $\mathcal{C}_{XY}$. Finally, we treat the advection step in the I-direction explicitly, as well, and use an appropriately modified Euler scheme with the corresponding matrix denoted by $\mathcal{L}_I$.

Schematically, the solution U^n at step t^n given the solution U^{n+1} at step t^{n+1} and the source function v^{n+1} is computed as follows:

$$\begin{aligned}
&(\mathcal{I} + \mathcal{L}_X)U^* = (\mathcal{I} - \mathcal{L}_X - 2\mathcal{L}_V + \mathcal{C}_{XY})U^{n+1} + v^{n+1} \\
&(\mathcal{I} + \mathcal{L}_Y)U^{**} = U^* + (\mathcal{L}_Y + \mathcal{J})U^{n+1} \\
&U^n = (\mathcal{I} + \mathcal{L}_I)U^{**},
\end{aligned} \tag{83}$$

where $\mathcal{I}$ is an identity matrix. The above scheme results in a system of triangular equations in X and Y directions.

In a similar fashion, we formulate the discretization scheme for the forward equation. In both schemes it is extremely important to describe the appropriate boundary conditions properly. While for backward equation these are mostly contract specific, for the forward equation they specify the marginal distribution of X (Y) for small and large values of Y (X). We now consider specifying boundary conditions for the forward equation.

First we consider HM where the variance Y is defined only on the positive axes, so that we have to specify the boundary for $Y = 0$. As formula (52) shows, the density is divergent if $\vartheta < 0$, which is almost always implied by the data. In principal, we can use formula (52) to specify the boundary condition for a small value of the lower boundary of Y. However, now we face problems specifying the boundary condition in the X-direction, since for small values of Y the density of the asset price degenerates to the delta function centered at the forward asset price, and the boundary condition becomes discontinuous for the joint density of X and Y. The same problem afflicts the backward equation.

When dealing with SSM and OUP for the volatility Y, Y is defined on the entire axis, so that we can apply the natural boundary condition, or, put differently, the convexity boundary condition $U_{YY} = 0$ and $U_{XY} = 0$ for both lower and upper boundaries of Y. The same applies for X, if X is the log-price. Now the behavior of X is the same for large positive and negative values of Y and corresponds to the case of the large variance. In this case, the natural boundary condition works well.

7.2. MC methods

Application of MC methods to the solution of the pricing problem is based on discretizing SDEs (7) for both the asset price and its variance, performing a large number of path simulations and evaluating the payoff and reward functions along each path.

Introducing the discrete time grid $\{t_n\}$ with time step Δt and applying the Euler discretization scheme to simulate the log-price X and its variance V under HM we obtain:

$$\begin{aligned}
X(t_{n+1}) &= X(t_n) + \left(\mu(t_n) - \tfrac{1}{2}V(t_n)\right)\Delta t + \sqrt{V(t_n)\Delta t}\left(\sqrt{1-\rho^2}\varsigma_{n,1} + \rho\varsigma_{n,2}\right), \\
V(t_{n+1}) &= V(t_n) + \kappa(\theta^2 - V(t_n))\Delta t + \sqrt{V(t_n)\Delta t}\,\varsigma_{n,2},
\end{aligned} \tag{84}$$

where $\varsigma_{n,1}$ and $\varsigma_{n,2}$ are independent standard normal random variables.

We immediately see that scheme (84) fails when $V(t_{n+1})$ becomes negative. There is a large (and growing) number of papers that address this issue for HM. We refer to Andersen (2008) and Lord *et al* (2006) for an overview of possible solution methods (none of which is, in our opinion, a satisfactory one). The simplest solution is to apply $V(t_{n+1}) = \max(V(t_{n+1}), 0)$ after evaluating the second SDE; however this introduces a heavy bias in the convergence of the MC scheme.

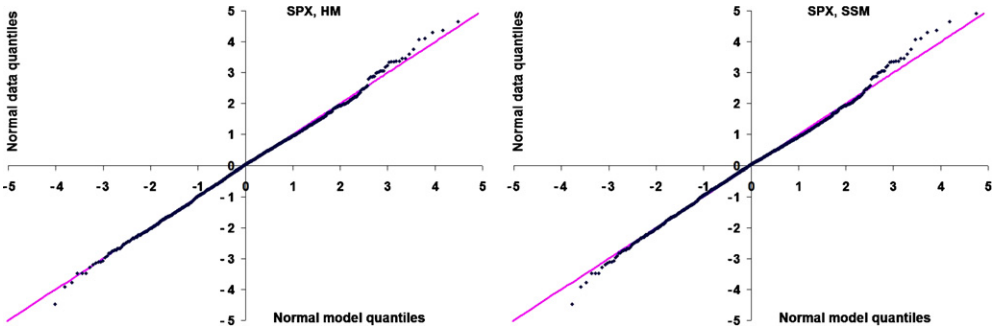

Figure 1. QQ plot for the residuals of HM (left) and SSM (right) fitted to the SPX time series. We see that both stochastic volatility models can adequately describe the historical distribution of SPX returns, and it is hard to decide which of the models is better on these grounds.

Applying the same scheme to simulate log-price X and its volatility σ in SSM yields:

$$\begin{aligned} X(t_{n+1}) &= X(t_n) + \left(\mu(t_n) - \tfrac{1}{2}\sigma^2(t_n)\right)\Delta t + \sigma(t_n)\sqrt{\Delta t}\left(\sqrt{1-\rho^2}\varsigma_{n,1} + \rho\varsigma_{n,2}\right), \\ \sigma(t_{n+1}) &= \sigma(t_n) + \kappa(\theta - \sigma(t_n))\Delta t + \sqrt{\Delta t}\,\varsigma_{n,2}. \end{aligned} \tag{85}$$

Now we see that there is no extreme case when scheme (85) can fail. This makes the MC implementation of SSM simple and robust especially when compared to HM.

8. Illustration

We now provide some illustrations for the usage of SSM and HM by applying them to historical and option implied data. For this purpose, we will use the S&P 500 stock index (SPX) and General Motor's stock (GM), since the value of a large index and the value of a single company equity tend to have different implied and historical distributions.

8.1. Parameter estimation using the time series of returns

For parameter estimation using the time series we use daily log-returns observed from 5 January 1998, to 26 October 2007 with the total number of observations of 2468. To estimate parameters of HM we use the small-time formula (59) assuming zero correlation between the spot and the variance and zero asset price drift. For SSM we use the small-time formula (38) and the same assumptions. We assume that the time series is stationary and ergodic, so that the time distribution is equivalent (in probability) to the spacial distribution, and estimate model parameters by minimizing the differences between the empirical and theoretical Q–Q plots. The corresponding Q–Q plots are shown in figures 1 and 2 for the SPX and GM time series, respectively. We see that for both time series, both SSM and HM adequately describe the shape of the distribution in the middle, while they fail to describe a small number of large positive and negative returns. To describe these rare big negative and positive returns, we would need to introduce asset price jumps.

In table 1 we report parameter estimates. We observe that while both SSM and HM imply similar estimates for mean and volatilities of the SPX volatility, the mean reversion parameters are different. In our experience, the mean reversion and volatility parameters in HM are not stable and there can be different combinations of the two which yield almost identical distributions.

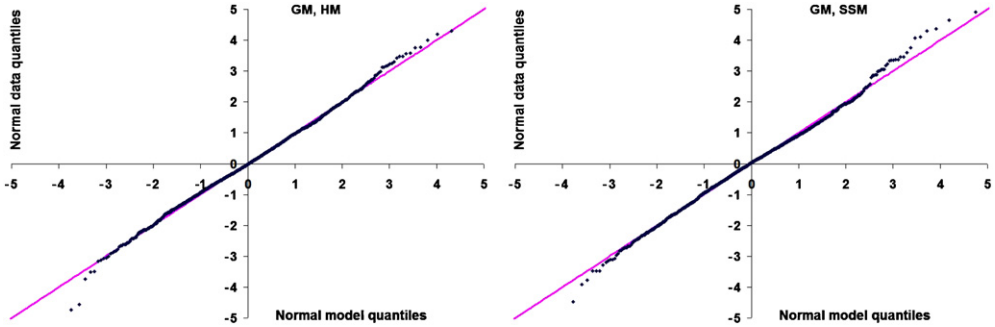

Figure 2. QQ plot for the residuals of HM (left) and SSM (right) fitted to the GM time series. Although both stochastic volatility models describe the central part of the distribution of returns, they cannot fit outliers in its left tail.

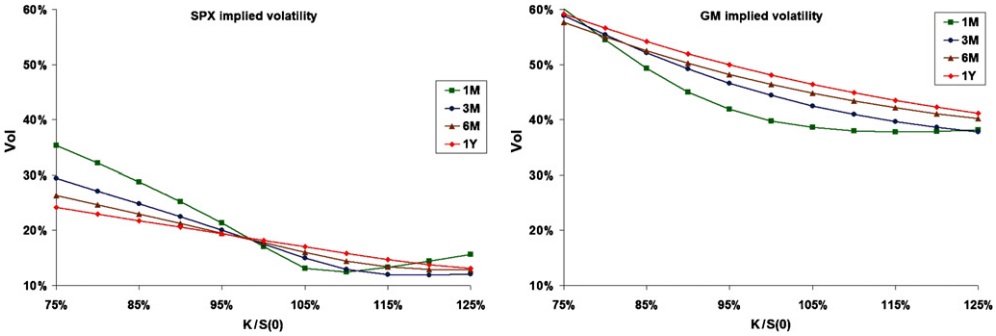

Figure 3. Implied volatilities for SPX (left) and GM (right) options as functions of the moneyness $K/S(0)$. As might be expected, GM has a higher level of the implied volatility and a steeper skew compared to the SPX implied volatility. This is due to the averaging effect present in the implied volatility of a stock index.

Table 1. HM and SSM parameters with $\mu = 0$ and $\rho = 0$ estimated by fitting both models to the time series for SPX and GM from 5 January 1998 to 26 October 2007. HM and SSM have similar estimates for the mean volatility and mean-reversion rate, however, the volatility of volatility is higher for HM than SSM.

	SPX, HM	SPX, SSM	GM, HM	GM, SSM
θ	0.1759	0.1573	0.3425	0.3184
κ	9.0101	4.0968	2.3349	1.1392
ϵ	0.6214	0.1693	0.5125	0.1563

8.2. Estimation using options data

To estimate model parameters using the volatilities implied from quoted options, we use data for 4 September 2007. In figure 3, we show the implied volatility, which is backed-out from the market price of an option using BSM formula (10) with constant volatility parameter, as functions of the moneyness $K/S(0)$ for four maturities. A typical pattern for the index implied volatility is that the left wing of the skew decreases with time while the level of ATM volatility, roughly corresponding to $K/S(0) = 100\%$, increases. GM implied volatilities

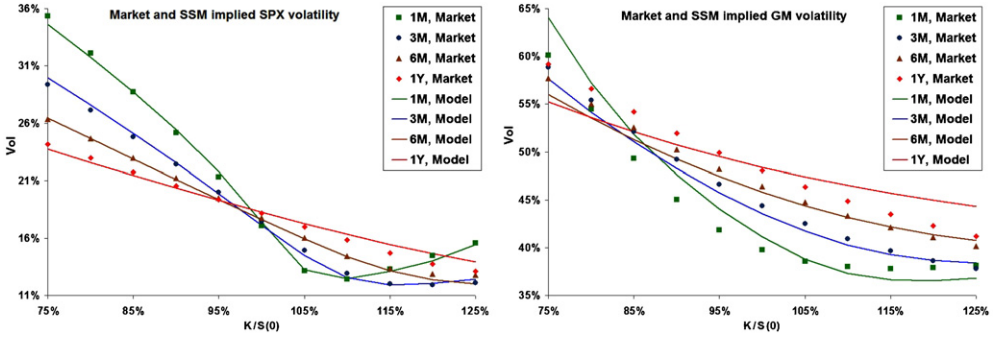

Figure 4. SSM implied volatilities versus market implied volatilities for SPX (left) and GM options (right). The SSM model reproduces the SPX implied volatility remarkably well; when augmented with the jump-to-default process, it fits GM implied volatilities with acceptable accuracy.

Table 2. HM and SSM parameters estimated from option implied volatilities for SPX and GM. For SPX options, HM and SSM models roughly satisfy their compatibility condition discussed in section 6. For GM options, SSM implies stronger mean-reversion and higher volatility of volatility compared to HM; both models imply similar probabilities of GM defaulting in one year (about 4%).

	SPX, HM	SPX, SSM	GM, HM	GM, SSM
$\sigma(0)$	0.1835	0.1883	0.4055	0.3844
θ	0.1936	−0.0177	0.6978	0.3795
κ	9.1658	4.2969	0.5284	4.0386
ϵ	1.2082	0.5743	0.9382	0.7937
ρ	−0.7579	−0.7696	−0.5362	−0.4675
α	0.0000	0.0000	0.0449	0.0435

are remarkably higher than the SPX ones, and the left wing of the skew is persistent across all maturities. The latter observation advocates augmenting the asset price process with the jump-to-default process, which is modeled as a first jump of the Poison process with constant intensity α (see, for example, Sepp (2007) for more details). We use SSM and HM augmented with the jump-to-default process with intensity α to fit GM implied volatilities by minimizing the squared differences between the market and model implied volatilities.

The parameter estimates are reported in table 2. We see that for SPX the compatibility conditions, which we have already discussed in section 6, of HM and SSM are roughly satisfied, so that both models imply almost identical distribution of the SPX spot price as well as its implied and realized variance. For GM data, the mean-reversion parameter for HM turns out to be small, once again indicating the instability of this parameter for HM. Both models imply almost identical default intensity parameter α and initial volatilities.

In figure 4 we show SSM fit to the market implied volatility. We see that for SPX options, the model adequately describes market implied volatilities across all strikes and maturities. For GM options, while fitting the short-term maturities, the model is off for longer-term maturities. This result is typical when fitting a stochastic volatility model for single name implied volatilities. To improve the fit, we can use the term structure of one or more model parameters.

Now we will analyze some of the model implications. For this purpose we use SSM and HM with parameter estimates obtained from SPX data. In figure 5 we show the long-term

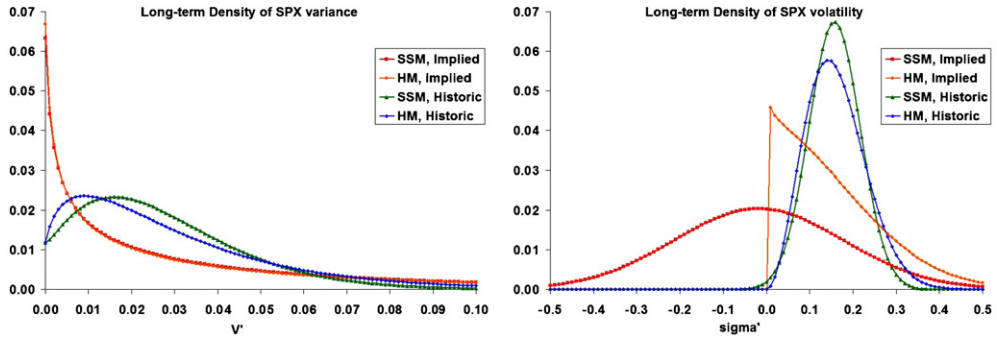

Figure 5. Left side: the model implied density of the long-term variance for HM and SSM. Right side: the model implied density for the long-term volatility. HM and SSM under both implied and historical measures show non-zero probability of observing zero instantaneous variance, which makes proper numerical implementation of HM very difficult.

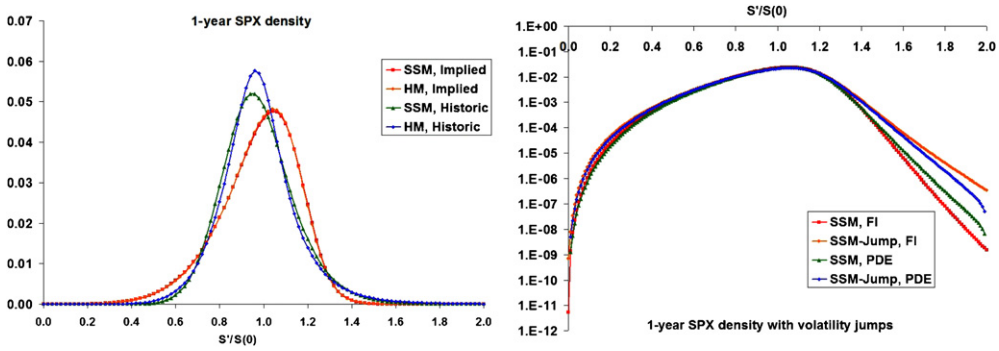

Figure 6. Left side: SPX implied density. Right side: SPX implied density with volatility jump parameters $\gamma = 1, \eta = 0.2$ on the log scale. These graphs show that the left tail of the SPX distribution is heavier under the implied measure, and that our closed-form solution based on the proposed discretization of the jump term (65) provides a good approximation to the actual density of the asset price.

density of the variance and volatility implied by these parameters and computed using formulas (26) and (29) for SSM and (41) and (43) for HM, respectively. We see that while both SSM and HM produce identical density for the variance under the pricing measure, the volatility density is clearly different. We also note that under the pricing measure there is a significant probability of observing small values of the variance and this makes HM difficult to handle numerically.

In figure 6 on the left side we show the 1 year SPX density implied by both historical and pricing measures under SSM and HM computed by formulae (45) and (31), respectively. We see again that under the pricing measure SSM and HM are equivalent. We also see that the pricing measure implies heavy left tail for the SPX price. On the right side we show the SPX density implied by SSM under the risk-neutral measure with the volatility jump parameters $\gamma = 1, \eta = 0.2$ on the log scale. First we compare the numerical PDE solution with the inversion of transformed density (31) to ensure that both methods do agree. (The right tail does seem a little off; however it is too far from the initial level and would not affect real calculations.) Then we add jump parameters and contrast the numerical PDE solution

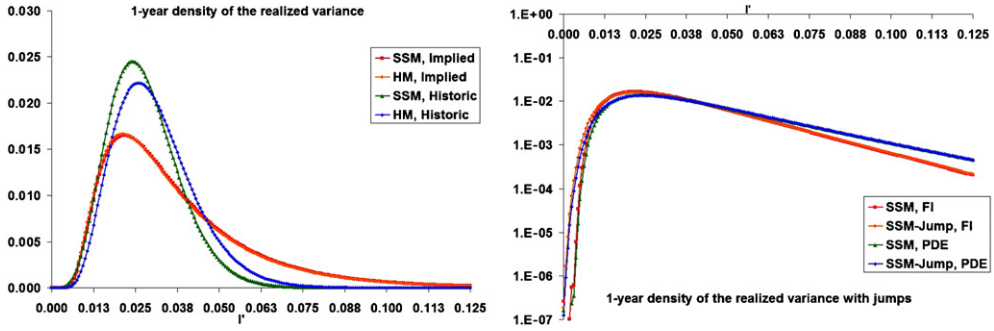

Figure 7. Left side: the implied density of the 1 year SPX realized variance. Right side: the density of the realized variance with volatility jump parameters $\gamma = 1, \eta = 0.2$ on the log scale. The left graph shows that the right tail of the distribution of the SPX realized variance is heavier under the implied measure. The right graph shows that our closed-form solution provides a good approximation to the actual density of the realized variance.

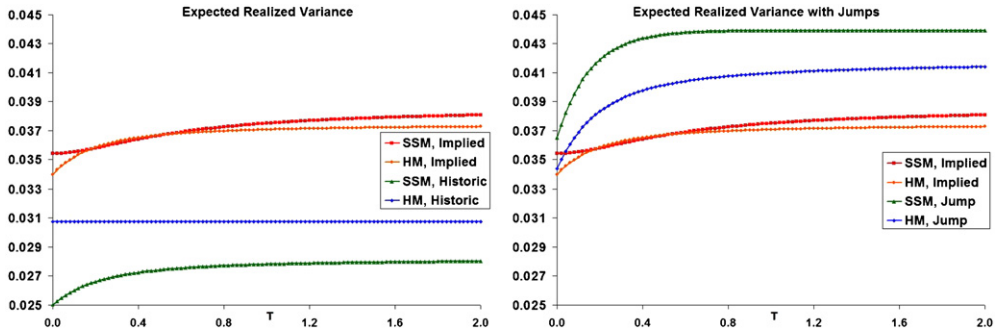

Figure 8. Left side: Expected realized variance of SPX as a function of maturity time T. Right side: expected realized variance under the pricing measure with volatility jump parameters $\gamma = 1, \eta = 0.2$. Both SSM and HM imply comparable levels of the expected realized variance which is higher than the historical one. By including volatility jumps we acquire more flexibility for modeling the term structure of the expected realized variance. This flexibility is important by pricing and hedging options on the realized variance.

(PDE) with our inversion formula (FI) where functions $A(\tau)$, $B(\tau)$, $B(\tau)$ are now calculated similarly to formula (66). We make two observations: first, our closed-form solution based on the expansion (SSM-Jump, FI) is in agreement with the PDE method (SSM-Jump, PDE) and, second, jumps in volatility have little effect on the left tail of the distribution whilst increase the right tail. The reason is that, all being equal, jumps in volatility decrease the terminal correlation between the asset and its volatility which shifts the right tail of the distribution. In general, volatility jumps cannot be implied from the vanilla options data, they should be estimated from options on the realized asset variance.

In figure 7 on the left side we show the density of the 1 year realized variance for SSM and HM computed using formulae (64) and (77), respectively. It follows that the density of the realized variance under the pricing measure has heavier right tails. On the right side, we show the density implied by SSM under the pricing measure with the volatility jump parameters $\gamma = 1$ and $\eta = 0.2$ on the log scale. Again, we first compare the inversion formula (64) (SSM-Jump, FI) with the numerical PDE solution (SSM-Jump, PDE) and then we do the same analysis for SSM with volatility jumps to ensure that both methods do agree.

Finally, in figure 8 on the left side we show the term structure of the expected realized variance for SSM and HM, computed by formulae (72) and (81) respectively, under both historical and pricing measures. Under the historical measure the expected realized asset variance is almost flat, while under the risk-neutral measure it has a positive slope due to the mean-reversion. On the right side of figure 8, we show the impact of the volatility jump with parameters $\gamma = 1$ and $\eta = 0.2$ for both SSM and HM. We see that SSM implies a higher impact of volatility jumps on the realized asset variance because of the nonlinear relationship between volatility and variance.

Acknowledgments

We are grateful to our colleagues at Merrill Lynch and Imperial College for many useful discussions. One of us (AL) wishes to thank Darryl Holm for many years of friendship.

References

Aase K 2004 *Wilmott Mag.* **4** 64
Andersen L 2008 *J. Comp. Finance* **11** 1
Bayly B, Holm D and Lifschitz A 1996 *Phil. Trans. R. Soc. London* A **354** 895
Black F and Scholes M 1973 *J. Political Economy* **81** 637
Carr P and Madan D 1999 *J. Comp. Finance* **2** 61
Chandrasekhar S 1961 *Hydrodynamic and Hydromagnetic Stability* (Oxford: Clarendon)
Cheng A and Scaillet O 2004 *Math. Finance* **17** 575
Cheng L, Filipovic D and Poor V 2004 *Math. Finance* **14** 515
Craig I and Sneyd A 1988 *Comput. Math. Appl.* **16** 341
Cox J 1975 Stanford University, unpublished
Cox J 1996 *J. Portfolio Manage.* **22** 15 (reprint)
Demeterfi K, Derman E, Kamal M and Zou J 1999 Goldman Sachs Quantitative Strategies Research Notes
Dragulescu A and Yakovenko V M 2002 *Quant. Finance* **2** 443
Duffie D, Pan J and Singleton K 2000 *Econometrica* **68** 1343
Dupire B 1994 *Risk* **7** 18
Galluccio S and Cam Y 2005 *Preprint* EconWPA, Finance, 0510028
Hagan P, Kumar D, Lesniewski A and Woodward D 2002 *Wilmott Mag.* **1** 84
Heston S L 1993 *Rev. Financ. Stud.* **6** 327
Kelvin L 1887 *Phil. Mag.* **24** 188
Lewis A 2000 *Option Valuation under Stochastic Volatility* (Newport Beach, California: Finance Press)
Lipton A 2001 *Mathematical Methods for Foreign Exchange: A Financial Engineer's Approach* (Singapore: World Scientific)
Lipton A 2002 *Risk* **15** 81
Lord R, Koekkoek R and Dijk D 2006 Tinbergen Institute, *Preprint* T1 2006-046/4
McKee S, Wall D and Wilson S 1996 *J. Comput. Phys.* **126** 64
Merton R 1973 *Bell J. Economics Manage. Sci.* **4** 141
Perello J, Masoliver J and Anento N 2004 *Physica* A **344** 134
Romano M and Touzi N 1997 *Math. Finance* **7** 399
Schobel R and Zhu J 1999 *Eur. Finance Rev.* **3** 23
Sepp A 2007 Affine models in mathematical finance: an analytical approach *PhD Thesis* University of Tartu
Silva A C, Prange R E and Yakovenko V M 2004 *Physica* A **344** 227
Silva A C and Yakovenko V M 2003 *Physica* A **324** 303
Silva A C and Yakovenko V M 2007 *Physica* A **382** 278
Stein E and Stein J 1991 *Rev. Financ. Stud.* **4** 727
Willard G 1997 *J. Derivatives* **5** 45