

Gamma Expansion of the Heston Stochastic Volatility Model

Paul Glasserman · Kyoung-Kuk Kim

Received: date / Accepted: date

Abstract We derive an explicit representation of the transitions of the Heston stochastic volatility model and use it for fast and accurate simulation of the model. Of particular interest is the integral of the variance process over an interval, conditional on the level of the variance at the endpoints. We give an explicit representation of this quantity in terms of infinite sums and mixtures of gamma random variables. The increments of the variance process are themselves mixtures of gamma random variables. The representation of the integrated conditional variance applies the Pitman-Yor decomposition of Bessel bridges. We combine this representation with the Broadie-Kaya exact simulation method and use it to circumvent the most time-consuming step in that method.

Keywords Stochastic Volatility model · Monte Carlo methods

Mathematics Subject Classification (2000) 60H35 · 65C05 · 91B70

JEL Classification C63 · G12 · G13

1 Introduction

The Heston [18] stochastic volatility model is among the most fundamental models in both the theory and practice of financial mathematics. It provides a natural extension beyond geometric Brownian motion as a description of asset price dynamics, modeling volatility as a positive, mean-reverting, stochastic process. It also offers some tractability: Heston [18] showed that standard European options can be priced in this model through an extension of the Black-Scholes formula; evaluation of this formula reduces to numerical inversion of a characteristic function.

Despite this degree of tractability, the Heston model has proved notoriously difficult to simulate accurately. Standard Euler and Milstein discretization schemes (as in Kloeden and Platen [24]) often produce erratic results when applied to the model, largely as a consequence of the square root functions in its dynamics, which violate the Lipschitz conditions typically used to

P. Glasserman
Graduate School of Business, Columbia University, New York, NY 10027, USA
E-mail: pg20@columbia.edu

K.-K. Kim
Corresponding author, Barclays Capital, New York, NY 10019, USA
E-mail: catenoid@gmail.com

ensure convergence. Indeed, even the convergence of discretization schemes for the univariate variance process (the Cox-Ingersoll-Ross [10] process) has received considerable attention (see, in particular, Alfonsi [3], Berkaoui, Bossy, and Diop [7], Deelstra and Delbaen [11], and Higham and Mao [19]). Discretization methods for the Heston model (and, in some cases, its extensions) include Alfonsi [4], Andersen [5], Van Haastrecht and Pelsser [31], Kahl and Jäckel [21], and several methods compared in Lord, Koekkoek, and van Dijk [26]. Some of these methods are limited to certain parameter ranges; Andersen [5] uses a combination of approximations to make his method applicable to the variety of parameters that result from calibrating the Heston model to different markets.

Broadie and Kaya [8] provided a breakthrough by showing that the transitions of the Heston model can be simulated *exactly*, with no discretization error. Key to their method is sampling the integral of the variance process conditional on its endpoints. Broadie and Kaya [8] give the characteristic function of this integral and generate samples through numerical transform inversion. As noted by Broadie and Kaya [8] and subsequently by Lord et al. [26], this step is rather time-consuming because of the form of the characteristic function and, most importantly, because of its dependence on values of the variance process. This dependence precludes computing and storing a numerical approximation to the distribution in advance — the relevant distribution changes at each step of the simulation.

We address this key issue through a gamma expansion of the integral of the variance process conditional on its endpoints; that is, we provide an exact representation of its distribution as an infinite sum of mixtures of gamma random variables. This representation is based on results of Pitman and Yor [27] on decompositions of Bessel bridges and results of Pitman and Yor [28] on representations of certain infinitely divisible distributions. We note that a gamma expansion is available for the variance process itself. Thus, a transition of the Heston model can in principle be simulated exactly using sums and mixtures of gamma random variables and a single normal random variable. For practical implementation, we truncate certain infinite series, and we examine the resulting error both theoretically and numerically; we also consider some computational shortcuts. Across a variety of parameter values, we find that the method is both fast and accurate.

The rest of this paper is organized as follows. We present our main representation result in Section 2. Section 3 examines the truncation error and provides a central limit theorem and a gamma approximation for the remainder from truncation. Section 4 details the practical application to simulation. Section 5 reports numerical results, and Section 6 provides an alternative representation for time-varying step sizes. Section 7 concludes the paper.

2 Main Result

The Heston model is described by the pair of stochastic differential equations

$$\frac{dS_t}{S_t} = \mu dt + \sqrt{V_t} \left(\rho dW_t^1 + \sqrt{1 - \rho^2} dW_t^2 \right) \quad (2.1)$$

$$dV_t = \kappa(\theta - V_t) dt + \sigma \sqrt{V_t} dW_t^1, \quad (2.2)$$

in which (W^1, W^2) is a standard two-dimensional Brownian motion. The variable S_t describes the level of an underlying asset and V_t the variance of its instantaneous returns. The parameters κ , θ , σ (and typically also μ) are positive, and ρ takes values in $[-1, 1]$. We take the initial conditions S_0 and V_0 to be strictly positive.

The main result of this section provides a representation of the increments of the Heston model. Because the process (S, V) is a time-homogeneous Markov process, a representation of

(S_t, V_t) as a function of t and (S_0, V_0) extends to a representation of (S_{t_2}, V_{t_2}) given (S_{t_1}, V_{t_1}) for arbitrary $t_1 < t_2$.

We begin by considering the variance process (2.2). This is a square-root diffusion or Cox-Ingersoll-Ross [10] process, and its transition law is a scaled noncentral chi-square distribution; more precisely, we have

$$V_t = \frac{\sigma^2(1 - e^{-\kappa t})}{4\kappa} \chi_\delta'^2 \left(\frac{4\kappa e^{-\kappa t}}{\sigma^2(1 - e^{-\kappa t})} V_0 \right), \quad t > 0, \quad \delta = \frac{4\kappa\theta}{\sigma^2}, \quad (2.3)$$

where $\chi_\delta'^2(\lambda)$ denotes a noncentral chi-square random variable with δ degrees of freedom and non-centrality parameter λ . This representation allows exact simulation of V_t given V_0 , as discussed in Scott [30].

Broadie and Kaya [8] build on this representation to simulate the Heston model. Writing

$$S_t = S_0 \exp \left(\mu t - \frac{1}{2} \int_0^t V_s ds + \rho \int_0^t \sqrt{V_s} dW_s^1 + \sqrt{1 - \rho^2} \int_0^t \sqrt{V_s} dW_s^2 \right),$$

they observe that $\log(S_t/S_0)$ is conditionally normal, given $\int_0^t V_s ds$ and $\int_0^t \sqrt{V_s} dW_s^1$,

$$\log \frac{S_t}{S_0} \sim \mathcal{N} \left(\mu t - \frac{1}{2} \int_0^t V_s ds + \rho \int_0^t \sqrt{V_s} dW_s^1, (1 - \rho^2) \int_0^t V_s ds \right) \quad (2.4)$$

as W^2 is independent of V . From (2.2), we also have

$$\int_0^t \sqrt{V_s} dW_s^1 = \frac{1}{\sigma} \left\{ V_t - V_0 - \kappa\theta t + \kappa \int_0^t V_s ds \right\}.$$

Hence, if one can sample from the joint distribution of the pair

$$\left(V_t, \int_0^t V_s ds \right),$$

then simulating S_t given (S_0, V_0) reduces to sampling from the conditional normal distribution in (2.4). As V_t can be sampled using (2.3), the Broadie-Kaya method reduces the problem of exact simulation of the Heston model to the problem of sampling from

$$\left(\int_0^t V_s ds \mid V_0, V_t \right), \quad (2.5)$$

the conditional distribution of the integrated variance over $[0, t]$, given the level of the variance at the endpoints. Broadie and Kaya [8] sample from this distribution through numerical inversion of its characteristic function (as discussed in Section 4); this method is rather time-consuming, particularly because the characteristic function changes at each step along a simulated path as the corresponding values of V at the endpoints of each interval change. Our focus is on deriving a more tractable representation of (2.5), for which we use properties of squared Bessel bridges studied by Pitman and Yor [27] and investigate associated characteristic functions.

Remark 2.1 The notation in (2.5) is somewhat informal, since it involves conditioning on an outcome of probability 0, but it can be made precise as follows. Define a process B through the time change

$$V_t = e^{-\kappa t} B \left(\frac{e^{\kappa t} - 1}{4\kappa/\sigma^2} \right); \quad (2.6)$$

then B satisfies

$$dB(t) = \delta dt + 2\sqrt{B(t)}dW(t), \quad W\left(\frac{e^{\kappa t} - 1}{4\kappa/\sigma^2}\right) = \int_0^t \frac{\sigma}{2} e^{\kappa s/2} dW_s^1,$$

which makes B a δ -dimensional squared Bessel process, with δ as defined in (2.3). The conditional law $(B_s, 0 \leq s \leq t | B_0, B_t)$ is made precise by Pitman and Yor [27], p.446, and this defines the conditional law $(V_s, 0 \leq s \leq t | V_0, V_t)$ through (2.6).

To state our main result, we need to recall the definition of the Bessel distribution, which we denote by $BES(\nu, z)$, with parameter $\nu > -1$ and $z > 0$ (see Yuan and Kalbfleisch [32] for a study of this distribution). The $BES(\nu, z)$ distribution is supported on the nonnegative integers with probability mass function

$$p_n \equiv p_n(\nu, z) = \frac{(z/2)^{2n+\nu}}{I_\nu(z)n!\Gamma(n+\nu+1)}, \quad n \geq 0,$$

where I_ν is a modified Bessel function of the first kind. We now have the following result:

Theorem 2.2 *The distribution (2.5) of the integrated variance conditional on its endpoints admits the representation*

$$\left(\int_0^t V_s ds \mid V_0 = v_0, V_t = v_t\right) \stackrel{d}{=} X_1 + X_2 + X_3 \equiv X_1 + X_2 + \sum_{j=1}^{\eta} Z_j, \quad (2.7)$$

in which $X_1, X_2, \eta, Z_1, Z_2, \dots$ are mutually independent, the Z_j are independent copies of a random variable Z , and η is a Bessel random variable with parameters $\nu = \delta/2 - 1$ and

$$z = \frac{2\kappa/\sigma^2}{\sinh(\kappa t/2)} \sqrt{v_0 v_t}.$$

Moreover, X_1, X_2 and Z have the following representations:

$$X_1 \stackrel{d}{=} \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1), \quad X_2 \stackrel{d}{=} \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \Gamma_n(\delta/2, 1), \quad Z \stackrel{d}{=} \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \Gamma_n(2, 1), \quad (2.8)$$

where $\delta = 4\kappa\theta/\sigma^2$,

$$\lambda_n = \frac{16\pi^2 n^2}{\sigma^2 t (\kappa^2 t^2 + 4\pi^2 n^2)}, \quad \gamma_n = \frac{\kappa^2 t^2 + 4\pi^2 n^2}{2\sigma^2 t^2},$$

the N_n are independent Poisson random variables with respective means $(v_0 + v_t)\lambda_n$, the $\text{Exp}_j(1)$ are independent, unit-mean exponential random variables, and the $\Gamma_n(\alpha, \beta)$ denote independent gamma random variables with shape parameter α and scale parameter β .

Remark 2.3 Before proceeding with the proof, we note that the distribution of the increments of V can be put in a related form. Let $\gamma = 4\kappa/(\sigma^2(1 - \exp(-\kappa t)))$, then

$$(V_t | V_0 = v_0) \stackrel{d}{=} \frac{1}{\gamma} \left(\Gamma_0(\delta/2, 2) + \sum_{j=1}^{N_{\lambda/2}} \Gamma_j(1, 2) \right), \quad \lambda = \gamma e^{-\kappa t} v_0,$$

with $\Gamma_0, \Gamma_1, \dots$ denoting independent gamma random variables and $N_{\lambda/2}$ a Poisson random variable with mean $\lambda/2$. This follows from the representation of the noncentral chi-square distribution as a Poisson mixture of ordinary chi-square distributions. The Bessel random variable η

can also be expressed in terms of gamma and Poisson random variables; for example, it follows from (5.j) in Pitman and Yor [27] that

$$\eta \sim \left(N \mid \Gamma(\nu + 1, 1) + \sum_{j=1}^N \text{Exp}_j(1) = z^2/4 \right),$$

with N a unit-mean Poisson random variable. However, we will not use this representation to generate η .

Proof We prove the result in two steps. We first decompose (2.5) into the the sum of three independent random variables and identify their Laplace transforms; we then derive the series representations from the transforms.

For the first step, we apply a result of Pitman and Yor [27] through a transformation of the variance process. Fix $t > 0$ and define $\{A_s\}_{0 \leq s \leq 1}$ by setting

$$A_s = \frac{4}{\sigma^2 t} V_{st}. \quad (2.9)$$

Then A solves the stochastic differential equation

$$dA_s = (\delta + 2aA_s) ds + 2\sqrt{A_s} dW_s$$

with $a = -\kappa t/2$ and W a standard Brownian motion, $W_s = W_{st}^1/\sqrt{t}$. This is the equation for a δ -dimensional squared Ornstein-Uhlenbeck (OU) process with parameter a . (For integer δ , A has the law of the squared Euclidean norm of a δ -dimensional OU process; the name applies by extension to all real $\delta \geq 0$.) Let us denote by $A_{x,y}^{\delta,1} = \{A_{x,y}^{\delta,1}(s)\}_{0 \leq s \leq 1}$, a process with law

$$(A_s, 0 \leq s \leq 1 \mid A_0 = x, A_1 = y), \quad (2.10)$$

the law of A conditioned on its endpoints — a squared OU bridge. Pitman and Yor [27] show (see the discussion of Theorem 5.8 following equation (6.d) on p.456 of [27]) that the squared OU bridge admits the decomposition

$$A_{x,y}^{\delta,1} \stackrel{d}{=} A_{x,0}^{0,1} + A_{0,y}^{0,1} + A_{0,0}^{\delta,1} + A_{0,0}^{4\eta,1},$$

in which the four squared OU bridges on the right are independent processes, and η is an independent Bessel random variable with parameters $\nu = \delta/2 - 1$ and $z = \sqrt{xy}a/\sinh(a)$. As explained in Section 5.3 of [27], $A_{0,y}^{0,1}$ should be understood as the time-reversal of the process $A_{y,0}^{0,1}$ because 0 is an absorbing state for a 0-dimensional squared OU process.

From the decomposition above, we get

$$\int_0^1 A_{x,y}^{\delta,1}(s) ds \stackrel{d}{=} \int_0^1 A_{x,0}^{0,1}(s) ds + \int_0^1 A_{0,y}^{0,1}(s) ds + \int_0^1 A_{0,0}^{\delta,1}(s) ds + \int_0^1 A_{0,0}^{4\eta,1}(s) ds.$$

The law of the second integral on the right would be unchanged if we replaced its integrand with an independent copy of $A_{y,0}^{0,1}$, in view of the definition of $A_{0,y}^{0,1}$ through time reversal. Moreover, independent OU bridges have the following additivity property:

$$A_{x,0}^{\delta,1} + A_{y,0}^{\delta',1} \stackrel{d}{=} A_{x+y,0}^{\delta+\delta',1};$$

this follows from a similar additivity property for squared Bessel bridges (equation (1.b) in Pitman and Yor [27]) through the transformation (6.b) in [27]. Therefore, we have

$$\int_0^1 A_{x,y}^{\delta,1}(s) ds \stackrel{d}{=} \int_0^1 A_{x+y,0}^{0,1}(s) ds + \int_0^1 A_{0,0}^{\delta,1}(s) ds + \int_0^1 A_{0,0}^{4\eta,1}(s) ds. \quad (2.11)$$

The additivity property also allows us to express the last term on the right as

$$\int_0^1 A_{0,0}^{4\eta,1}(s) ds \stackrel{d}{=} \sum_{j=1}^{\eta} \int_0^1 \left(A_{0,0}^{4,1}\right)^{(j)}(s) ds,$$

using independent copies $(A_{0,0}^{4,1})^{(j)}$ of $A_{0,0}^{4,1}$.

Next, we convert (2.11) into a decomposition of (2.5). It follows from (2.9) and (2.10) that

$$\left(\int_0^t V_s ds \mid V_0 = v_0, V_t = v_t\right) \stackrel{d}{=} \frac{\sigma^2 t^2}{4} \int_0^1 A_{x,y}^{\delta,1}(s) ds$$

with $x = 4v_0/(\sigma^2 t)$ and $y = 4v_t/(\sigma^2 t)$. Thus, (2.7) holds with

$$\begin{aligned} X_1 &= \frac{\sigma^2 t^2}{4} \int_0^1 A_{x+y,0}^{0,1}(s) ds, \\ X_2 &= \frac{\sigma^2 t^2}{4} \int_0^1 A_{0,0}^{\delta,1}(s) ds, \\ Z &= \frac{\sigma^2 t^2}{4} \int_0^1 A_{0,0}^{4,1}(s) ds. \end{aligned}$$

Since $a = -\kappa t/2$, $x = 4v_0/(\sigma^2 t)$ and $y = 4v_t/(\sigma^2 t)$, we have

$$z = \frac{a}{\sinh(a)} \sqrt{xy} = \frac{2\kappa/\sigma^2}{\sinh(\kappa t/2)} \sqrt{v_0 v_t}.$$

To derive the series representations of X_1 , X_2 , and Z , we record their Laplace transforms in the following lemma, which is proved in the appendix.

Lemma 2.4 *The Laplace transforms Φ^1 , Φ^2 , Φ^Z of X_1 , X_2 and Z are, for $b \geq 0$,*

$$\Phi^1(b) = \exp\left(\frac{(v_0 + v_t)}{\sigma^2} \left(\kappa \coth \frac{\kappa t}{2} - L \coth \frac{Lt}{2}\right)\right), \quad (2.12)$$

$$\Phi^2(b) = \left(\frac{L}{\kappa} \cdot \frac{\sinh \kappa t/2}{\sinh Lt/2}\right)^{\delta/2}, \quad (2.13)$$

$$\Phi^Z(b) = \left(\frac{L}{\kappa} \cdot \frac{\sinh \kappa t/2}{\sinh Lt/2}\right)^2, \quad (2.14)$$

where $L = \sqrt{2\sigma^2 b + \kappa^2}$.

A useful tool in working with these Laplace transforms is the following infinite product from Pitman and Yor [28], p.308:

$$\prod_{n=1}^{\infty} \left(1 + \frac{x^2}{\pi^2 n^2}\right)^{-1} = \frac{x}{\sinh x}. \quad (2.15)$$

From Ahlfors [2], p.192, Theorems 5 and 6, we know that an infinite product $\prod(1+a_n)$ with $a_n \neq -1$ converges simultaneously with $\sum \log(1+a_n)$ (using the principal branch in \mathbb{C} if necessary) and that this product converges absolutely if and only if $\sum |a_n|$ does. These properties and (2.15) imply that

$$\sum_{n=1}^{\infty} \log \left(1 + \frac{x^2}{\pi^2 n^2}\right) = -\log \frac{x}{\sinh x}.$$

The series on the left as well as its termwise derivative is uniformly convergent on compact intervals, so we can deduce that

$$\sum_{n=1}^{\infty} \frac{2x^2}{x^2 + \pi^2 n^2} = x \coth x - 1.$$

It follows that, for real values $x \geq y \geq 0$, we have

$$x \coth x - y \coth y = \sum_{n=1}^{\infty} \frac{2x^2}{x^2 + \pi^2 n^2} - \sum_{n=1}^{\infty} \frac{2y^2}{y^2 + \pi^2 n^2} = \sum_{n=1}^{\infty} \frac{2\pi^2 n^2 (x^2 - y^2)}{(x^2 + \pi^2 n^2)(y^2 + \pi^2 n^2)}.$$

Plugging $x = Lt/2$, $y = \kappa t/2$ into this formula (with $L = \sqrt{2\sigma^2 b + \kappa^2}$) and rearranging terms, we get

$$\frac{1}{\sigma^2} \left(\kappa \coth \frac{\kappa t}{2} - L \coth \frac{Lt}{2} \right) = - \sum_{n=1}^{\infty} \frac{16\pi^2 n^2}{\sigma^2 t (\kappa^2 t^2 + 4\pi^2 n^2)} \cdot \frac{b}{b + \frac{\kappa^2 t^2 + 4\pi^2 n^2}{2\sigma^2 t^2}} = - \sum_{n=1}^{\infty} \frac{\lambda_n b}{b + \gamma_n}.$$

This provides an alternative expression for the Laplace transform of X_1 in (2.12).

Now we compare this Laplace transform with that of the first series in (2.8). Set

$$X'_1 = \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1).$$

This random variable is well-defined because the sum of variances $\sum_{n=1}^{\infty} 2(v_0 + v_t)\lambda_n/\gamma_n^2$ is finite, and thus the infinite sum converges almost surely. For any $b \geq 0$, we have

$$\begin{aligned} \log \mathbb{E} e^{-bX'_1} &= \log \mathbb{E} \exp \left(-b \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1) \right) \\ &= \sum_{n=1}^{\infty} \log \mathbb{E} \exp \left(-\frac{b}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1) \right) \\ &= - \sum_{n=1}^{\infty} \frac{(v_0 + v_t)\lambda_n b}{b + \gamma_n}. \end{aligned}$$

Hence, by the uniqueness of the Laplace transform, $X_1 \stackrel{d}{=} X'_1$, and we can express the distribution of X_1 with the series representation in (2.8).

For X_2 , the identities in (2.13) and (2.15) yield

$$\begin{aligned} \log \mathbb{E} e^{-bX_2} &= \frac{\delta}{2} \left(\log \frac{Lt/2}{\sinh Lt/2} - \log \frac{\kappa t/2}{\sinh \kappa t/2} \right) \\ &= \frac{\delta}{2} \left\{ - \sum_{n=1}^{\infty} \log \left(1 + \frac{L^2 t^2}{4\pi^2 n^2} \right) + \sum_{n=1}^{\infty} \log \left(1 + \frac{\kappa^2 t^2}{4\pi^2 n^2} \right) \right\} \\ &= -\frac{\delta}{2} \sum_{n=1}^{\infty} \log \left(1 + \frac{b}{\gamma_n} \right) \\ &= \sum_{n=1}^{\infty} \log \mathbb{E} \exp \left(-\frac{b}{\gamma_n} \Gamma_n(\delta/2, 1) \right) \\ &= \log \mathbb{E} \exp \left(-b \sum_{n=1}^{\infty} \frac{1}{\gamma_n} \Gamma_n(\delta/2, 1) \right). \end{aligned}$$

The random variable given by the infinite series in the last expression is again well-defined because the variances of the terms in the series are summable. The expansion of Z is a special case with $\delta = 4$. \square

Remark 2.5 For the practical application of Theorem 2.2, it is important to note how the terms in the decomposition (2.7) depend on the endpoints v_0 and v_t on which we condition. In particular, we highlight the fact that X_2 and Z do not depend on these values at all, so their distributions do not change as we simulate a path of the Heston model over time steps of length t . The values of v_0 and v_t appear in the distribution of the Bessel random variable η and in the means of the Poisson random variables in X_1 .

Remark 2.6 We note that X_1 , X_2 and Z belong to certain special classes of infinitely divisible distributions. Bondesson [6], p.858, defines the class \mathcal{T}_2 of distributions that are weak limits of finite convolutions of mixtures of exponential distributions; X_1 is clearly in \mathcal{T}_2 . The generalized gamma convolutions form the subclass of \mathcal{T}_2 that are weak limits of finite convolutions of exponential distributions; X_2 and Z are in this subclass.

3 Series Truncation

As explained in the previous section, sampling from the conditional distribution (2.5) is a crucial step in the Broadie-Kaya [8] method; the decomposition in Theorem 2.2 is useful in sampling from this distribution. One way (though not the only way) to apply the decomposition is through truncation of the series that define X_1 , X_2 and Z , so in this section we examine the effect of truncation.

Define the remainder random variables

$$X_1^K = \sum_{n=K+1}^{\infty} \frac{1}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1), \quad X_2^K = \sum_{n=K+1}^{\infty} \frac{1}{\gamma_n} \Gamma_n(\delta/2, 1), \quad Z^K = \sum_{n=K+1}^{\infty} \frac{1}{\gamma_n} \Gamma_n(2, 1),$$

resulting from truncating the series at K terms each. The following result shows how the means and variances of the remainders decrease with K .

Lemma 3.1 *As K increases,*

$$\begin{aligned} \mathbb{E}X_1^K &\sim \frac{2(v_0 + v_t)t}{\pi^2 K}, & \text{Var}(X_1^K) &\sim \frac{2(v_0 + v_t)\sigma^2 t^3}{3\pi^4 K^3}, \\ \mathbb{E}X_2^K &\sim \frac{\delta\sigma^2 t^2}{4\pi^2 K}, & \text{Var}(X_2^K) &\sim \frac{\delta\sigma^4 t^4}{24\pi^4 K^3}, \\ \mathbb{E}Z^K &\sim \frac{\sigma^2 t^2}{\pi^2 K}, & \text{Var}(Z^K) &\sim \frac{\sigma^4 t^4}{6\pi^4 K^3}. \end{aligned}$$

Proof Observe that

$$\begin{aligned} \text{Var}(X_1^K) &= (v_0 + v_t) \sum_{n=K+1}^{\infty} \frac{2\lambda_n}{\gamma_n^2} \\ &= (128(v_0 + v_t)\pi^2\sigma^2 t^3) \sum_{n=K+1}^{\infty} \frac{n^2}{(\kappa^2 t^2 + 4\pi^2 n^2)^3} \\ &\sim (128(v_0 + v_t)\pi^2\sigma^2 t^3) \int_K^{\infty} \frac{y^2}{(4\pi^2 y^2)^3} dy \\ &= \frac{2(v_0 + v_t)\sigma^2 t^3}{3\pi^4 K^3}. \end{aligned}$$

Similar calculations establish the other asymptotics. \square

Rather than drop the remainders from truncation, we can approximate them with simpler random variables to try to reduce the error from truncation in a convenient way. The next result provides a central limit theorem, supporting a normal approximation for each of the remainders; but we show that a gamma approximation is also asymptotically valid, and we give a sense in which it provides a closer approximation to the remainder than does the normal. The gamma approximation has the added advantage of preserving positivity.

Proposition 3.2 *Let R_K denote any of the remainder random variables X_1^K , X_2^K or Z^K , and let Γ_K denote a gamma random variable with the same mean and variance as R_K . Then*

$$\frac{R_K - \mathbb{E}(R_K)}{\sqrt{\text{Var}(R_K)}} \implies \mathcal{N}(0, 1) \quad \text{and} \quad \frac{\Gamma_K - \mathbb{E}(R_K)}{\sqrt{\text{Var}(R_K)}} \implies \mathcal{N}(0, 1), \quad \text{as } K \uparrow \infty.$$

Moreover, if we let \mathcal{N}_K denote a normal random variable with the same mean and variance as R_K , then Γ_K is closer to R_K than \mathcal{N}_K is in the following sense: for all $b > 0$ in a neighborhood of the origin,

$$0 \leq \log \mathbb{E}e^{bR_K} - \log \mathbb{E}e^{b\Gamma_K} \leq \log \mathbb{E}e^{bR_K} - \log \mathbb{E}e^{b\mathcal{N}_K}, \quad (3.1)$$

for all sufficiently large K .

Proof We detail the case $R_K = X_1^K$; the proofs for the other two cases are very similar.

The Laplace transforms in Lemma 2.4 remain valid for negative arguments sufficiently close to zero, as well as for all positive arguments. From the representation in (2.8), it follows that, for all $b > 0$ in a neighborhood of the origin,

$$\log \mathbb{E}e^{bX_1^K} = \sum_{n=K+1}^{\infty} \frac{(v_0 + v_t)\lambda_n b}{\gamma_n - b} = (v_0 + v_t) \sum_{n=K+1}^{\infty} \sum_{m=1}^{\infty} \lambda_n \left(\frac{b}{\gamma_n}\right)^m,$$

and the double sequence is absolutely convergent for each fixed b and K . This yields

$$\log \mathbb{E} \exp \left(b \frac{X_1^K - \mathbb{E}X_1^K}{\sqrt{\text{Var}(X_1^K)}} \right) = \frac{b^2}{2} + (v_0 + v_t) \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \lambda_n \left(\frac{b}{\gamma_n \sqrt{\text{Var}(X_1^K)}} \right)^m.$$

The asymptotic normality of X_1^K follows from Lemma 15.15 of Kallenberg [22] if the double sequence in this expression converges to zero.

To establish this convergence, first we observe that, for fixed $s > 0$

$$\begin{aligned} \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \lambda_n \frac{s^m}{\gamma_n^m} &\leq \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \frac{s^m}{\gamma_n^m} \\ &= \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \frac{(2\sigma^2 t^2 s)^m}{(\kappa^2 t^2 + 4\pi^2 n^2)^m} \\ &\leq \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \int_K^{\infty} \frac{(2\sigma^2 t^2 s)^m}{(\kappa^2 t^2 + 4\pi^2 y^2)^m} dy \\ &\leq \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \int_K^{\infty} \frac{(2\sigma^2 t^2 s)^m}{(4\pi^2 y^2)^m} dy \\ &\leq \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \left(\frac{\sigma^2 t^2 s}{2\pi^2} \right)^m \frac{1}{(2m-1)K^{2m-1}}. \end{aligned}$$

Then, we get

$$\begin{aligned}
& \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \lambda_n \left(\frac{b}{\gamma_n \sqrt{\text{Var}(X_1^K)}} \right)^m \\
& \leq \frac{4}{\sigma^2 t} \sum_{m=3}^{\infty} \left(\frac{\sigma^2 t^2 b}{2\pi^2 \sqrt{\text{Var}(X_1^K)}} \right)^m \frac{1}{(2m-1)K^{2m-1}} \\
& \leq \frac{4K}{\sigma^2 t} \sum_{m=3}^{\infty} \left(\frac{\sigma^2 t^2 b}{2\pi^2 K^2 \sqrt{\text{Var}(X_1^K)}} \right)^m \\
& = \frac{4K}{\sigma^2 t} \left(\frac{\sigma^2 t^2 b}{2\pi^2 K^2 \sqrt{\text{Var}(X_1^K)}} \right)^3 \left(1 - \frac{\sigma^2 t^2 b}{2\pi^2 K^2 \sqrt{\text{Var}(X_1^K)}} \right)^{-1},
\end{aligned}$$

the last equality holding for all sufficiently large K . Lemma 3.1 now shows that the last expression goes to zero as K increases.

To match the mean and the variance of X_1^K , the shape parameter α_K and the scale parameter β_K of Γ_K must satisfy

$$\alpha_K = \frac{(\mathbb{E}X_1^K)^2}{\text{Var}(X_1^K)} \sim \frac{6(v_0 + v_t)K}{\sigma^2 t}, \quad \beta_K = \frac{\text{Var}(X_1^K)}{\mathbb{E}X_1^K} \sim \frac{\sigma^2 t^2}{3\pi^2 K^2}.$$

From

$$\log \mathbb{E}e^{b\Gamma_K} = -\alpha_K \log(1 - \beta_K b) = \sum_{m=1}^{\infty} \frac{\alpha_K}{m} (\beta_K b)^m,$$

we get

$$\log \mathbb{E} \exp \left(b \frac{\Gamma_K - \mathbb{E}X_1^K}{\sqrt{\text{Var}(X_1^K)}} \right) = \frac{b^2}{2} + \sum_{m=3}^{\infty} \frac{\alpha_K}{m} \left(\frac{\beta_K b}{\sqrt{\text{Var}(X_1^K)}} \right)^m,$$

and the asymptotic normality for Γ_K follows because

$$\sum_{m=3}^{\infty} \frac{\alpha_K}{m} \left(\frac{\beta_K b}{\sqrt{\text{Var}(X_1^K)}} \right)^m \leq \alpha_K \left(\frac{\beta_K b}{\sqrt{\text{Var}(X_1^K)}} \right)^3 \left(1 - \frac{\beta_K b}{\sqrt{\text{Var}(X_1^K)}} \right)^{-1} \rightarrow 0.$$

Moreover, we have

$$\begin{aligned}
\log \mathbb{E}e^{bX_1^K} - \log \mathbb{E}e^{b\Gamma_K} &= (v_0 + v_t) \sum_{m=3}^{\infty} \left(\sum_{n=K+1}^{\infty} \frac{\lambda_n}{\gamma_n^m} \right) b^m - \sum_{m=3}^{\infty} \left(\frac{\alpha_K \beta_K^m}{m} \right) b^m \\
&= (v_0 + v_t) \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \frac{\lambda_n}{\gamma_n^m} (1 - R_{n,m}) b^m
\end{aligned}$$

with

$$R_{n,m} = \frac{\alpha_K \beta_K^m / (m(v_0 + v_t))}{\sum_{n=K+1}^{\infty} \lambda_n / \gamma_n^m}.$$

The denominator of $R_{n,m}$ satisfies

$$\begin{aligned} \sum_{n=K+1}^{\infty} \frac{\lambda_n}{\gamma_n^m} &= \frac{16\pi^2}{\sigma^2 t} \sum_{n=K+1}^{\infty} \frac{n^2 (2\sigma^2 t^2)^m}{(\kappa^2 t^2 + 4\pi^2 n^2)^{m+1}} \\ &\sim \frac{4}{\sigma^2 t} \sum_{n=K+1}^{\infty} \frac{(2\sigma^2 t^2)^m}{(4\pi^2 n^2)^m} \\ &\sim \frac{4}{\sigma^2 t} \left(\frac{\sigma^2 t^2}{2\pi^2} \right)^m \int_K^{\infty} \frac{1}{y^{2m}} dy \\ &= \frac{4}{\sigma^2 t} \left(\frac{\sigma^2 t^2}{2\pi^2} \right)^m \frac{1}{(2m-1)K^{2m-1}}, \end{aligned}$$

and the numerator of $R_{n,m}$ satisfies $\alpha_K \beta_K^m / (m(v_0 + v_t)) \sim 6 / (m\sigma^2 t K^{2m-1}) (\sigma^2 t^2 / (3\pi^2))^m$. These properties imply that

$$R_{n,m} \sim \frac{6m-3}{2m} \left(\frac{2}{3} \right)^m, \quad \text{as } K \rightarrow \infty.$$

Thus,

$$\begin{aligned} \log \mathbb{E} e^{bX_1^K} - \log \mathbb{E} e^{b\Gamma_K} &= (v_0 + v_t) \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \frac{\lambda_n}{\gamma_n^m} (1 - R_{n,m}) b^m \\ &\leq (v_0 + v_t) \sum_{m=3}^{\infty} \sum_{n=K+1}^{\infty} \frac{\lambda_n}{\gamma_n^m} b^m \\ &= \log \mathbb{E} e^{bX_1^K} - \log \mathbb{E} e^{b\mathcal{N}_K}, \end{aligned}$$

for all sufficiently large K . □

The comparison in (3.1) indicates that the cumulant generating function (the log moment generating function) of the remainder is closer to that of the gamma approximation than to that of the normal approximation. Moreover, it indicates that the tail of the remainder is actually heavier than that of the gamma approximation. These observations suggest that approximating the remainder with a gamma random variable is preferable to simply dropping the remainder or approximating it with a normal random variable, and these suggestions are supported by separate numerical comparisons (not included here). Thus, when we truncate at level K , we use an approximation of the form

$$X_1 \approx \sum_{n=1}^K \frac{1}{\gamma_n} \sum_{j=1}^{N_n} \text{Exp}_j(1) + \Gamma_K^1,$$

preserving the first two moments of X_1 , and similarly for X_2 and Z .

Remark 3.3 We noted in the previous section that X_1 , X_2 and Z belong to special classes of infinitely divisible distributions. Indeed, it is easy to see from (2.12) and (2.13) that X_1 and X_2 are Lévy processes with “time” parameters $v_0 + v_t$ and δ , respectively. From the series expansions of X_1 and X_2 , we get their Lévy densities

$$\rho_1(x) = \sum_{n=1}^{\infty} \lambda_n \gamma_n e^{-\gamma_n x}, \quad \rho_2(x) = \sum_{n=1}^{\infty} \frac{1}{2x} e^{-\gamma_n x}.$$

Table 4.1: Model parameters.

	case I	case II	case III	case IV
κ	0.5	0.3	1	6.2
θ	0.04	0.04	0.09	0.02
σ	1	0.9	1	0.6
ρ	-0.9	-0.5	-0.3	-0.7

These observations raise the possibility of simulating X_1 and X_2 using their Lévy densities and methods of the type discussed in, e.g., Bondesson [6] and Cont and Tankov [9]. However, these methods also involve truncation and do not appear advantageous in our context, given the very explicit expressions available in (2.8).

4 Simulation Procedure

We now turn to the specifics of using Theorem 2.2 for simulation of the Heston model. As explained in Section 2, the Broadie-Kaya method reduces the problem of exact simulation of the Heston model to the problem of sampling from the conditional distribution in (2.5). Broadie and Kaya [8] implement this step through numerical inversion of the characteristic function of (2.5). The characteristic function is complicated, and each inversion is somewhat time-consuming; moreover, the characteristic function depends on the endpoints v_0 and v_t , so precomputing the transform inversion and storing the distribution in a table is not an option: the relevant transform changes at each step in the simulation. This makes the overall procedure rather slow and opens the possibility for improvement using the expansions in (2.8).

We treat separately the terms X_1 , X_2 and Z . Recall that the endpoints v_0 and v_t appear only in X_1 and the Bessel random variable η ; this enhances the potential efficiency gains from Theorem 2.2. Moreover, X_1 depends only on the sum of v_0 and v_t , and η depends only on their product. Thus, although the full distribution (2.5) depends on two path-dependent parameters, no element of our decomposition depends on more than one such parameter.

For our numerical experiments, we use the four sets of parameters in Table 4.1. These are intended to reflect the range of parameter values for which the Heston model is commonly used in practice, particularly since some methods are known to work better in some ranges than others. The first three cases are taken from Andersen [5], and case IV is set close to estimated parameters in Duffie, Pan and Singleton [15]. According to these sources, case I is relevant for long-dated FX options, case II for long-dated interest rate options, case III for equity options, and case IV for S&P 500 index options. Andersen [5] describes cases I–III as challenging as well as practically relevant.

4.1 Simulation of X_1

Exact Simulation. As the Laplace transform of X_1 is available in closed form, one can avoid series truncation and sample using transform inversion. The sampling algorithm generates $U \sim \text{Unif}[0,1]$ and then finds $x \geq 0$ such that $F_{X_1}(x) = U$, with F_{X_1} the distribution of X_1 ; the resulting x is returned as a sample from F_{X_1} . Searching for x involves a root-finding procedure combined with repeated transform inversion to evaluate $F_{X_1}(x)$ at each candidate value of x . For

Table 4.2: Time to generate 100 samples of X_1 from Φ^1 and time with a dummy calculation of $I_\nu(z)$ added.

$v_0 + v_t$	no $I_\nu(z)$	one $I_\nu(z)$
2	0.57	2.08
0.2	6.67	35.08
0.02	76.59	381.77

illustrative purposes, in this section we use the MATLAB `fzero` function for root-finding; for transform inversion we use the algorithm of Abate and Whitt [1], following the approach Broadie and Kaya [8] use for their transform inversion.

Using the Abate-Whitt algorithm, we calculate the distribution F_{X_1} using

$$F_{X_1}(x) \approx \frac{hx}{\pi} + \frac{2}{\pi} \sum_{k=1}^N \frac{\sin h k x}{k} \operatorname{Re}(\Phi^1(-i h k)), \quad h = \frac{2\pi}{x + u_\epsilon}. \quad (4.1)$$

Here, $u_\epsilon = \mu_{X_1} + m\sigma_{X_1}$, with μ_{X_1} the mean of X_1 , σ_{X_1} its standard deviation, and m set equal to 5. (We increase m for higher precision.) The parameter u_ϵ controls the discretization error in the inversion integral. Truncation error is handled by increasing N until $|\Phi^1(-i h N)|/N < \pi\epsilon'/2$, with $\epsilon' = 10^{-5}$.

The mean and variance of X_1 are given by

$$\mu_{X_1} = (v_0 + v_t)\mu_{X_1}^*, \quad \sigma_{X_1}^2 = (v_0 + v_t)\sigma_{X_1}^{*2}$$

where $\mu_{X_1}^*$ and $\sigma_{X_1}^{*2}$ are the mean and variance of X_1 with $v_0 = 1$ and $v_t = 0$,

$$\begin{aligned} \mu_{X_1}^* &= \frac{1}{\kappa} \coth\left(\frac{\kappa t}{2}\right) - \frac{t}{2} \operatorname{csch}^2\left(\frac{\kappa t}{2}\right), \\ \sigma_{X_1}^{*2} &= \frac{\sigma^2}{\kappa^3} \coth\left(\frac{\kappa t}{2}\right) + \frac{\sigma^2 t}{2\kappa^2} \operatorname{csch}^2\left(\frac{\kappa t}{2}\right) - \frac{\sigma^2 t^2}{2\kappa} \coth\left(\frac{\kappa t}{2}\right) \operatorname{csch}^2\left(\frac{\kappa t}{2}\right). \end{aligned}$$

These values can be computed and stored in the initialization of the Monte Carlo simulation if we work with a fixed time step t and fixed model parameters κ and σ .

We initialize the root-finding routine `fzero` with a guess of $x = F_N^{-1}(U)$, with F_N denoting the normal distribution $\mathcal{N}(\mu_{X_1}, \sigma_{X_1}^2)$. We set the tolerance level to 10^{-5} . These steps follow the approach in Broadie and Kaya [8].

A significant computational burden in inverting the characteristic function of (2.5), as Broadie and Kaya [8] do, results from a modified Bessel function of the first kind, $I_\nu(z)$, appearing in that characteristic function. To illustrate this point, we compare the time required to sample X_1 from its transform Φ^1 when we include a superfluous evaluation of $I_\nu(z)$ (which is unnecessary for Φ^1) and compare computing times.

The results of this comparison are shown in Table 4.2, where we record the time (in seconds) to generate 100 samples of X_1 . We use values of κ and σ from case I in Table 4.1, and we use a time step of $t = 1$. The value of $v_0 + v_t$ varies along each simulated path, so we take three different levels 0.02, 0.2, 2 for $v_0 + v_t$. The results confirm that evaluating the modified Bessel function at each step adds significantly to the total computing time. The results also indicate that the computing time is quite sensitive to the value of $v_0 + v_t$.




Fig. 4.1: Convergence of gamma approximations for cases I–IV with $v_0 + v_t = 0.02$.

Gamma Expansion. Next, we test the sampling of X_1 from the series representation in (2.8). We assume that the parameters κ , θ and σ remain fixed throughout each simulation, and we also use a fixed value of t , taking $t = 1$. With these parameters fixed, $\{\lambda_n\}$ and $\{\gamma_n\}$ can be tabulated in the initialization of the Monte Carlo simulation, along with $\mu_{X_1}^*$, $\sigma_{X_1}^*$, and values of $\mathbb{E}X_1^K$ and $\text{Var}(X_1^K)$ for multiple K . Recall that the mean and variance of X_1^K are used to determine the shape and scale parameters of the Γ_K^1 random variable. In our tests, we calculate those moments for all $K = 1, \dots, 100$.

Table 4.3 reports computing times for generating X_1 from the truncated gamma expansion with $K = 1$, $K = 20$ with $v_0 + v_t = 0.02$ using MATLAB. As one would expect, the computing time increases with K . However, the approximations seem to work quite well even for small K : Figure 4.1 illustrates the distributions of the truncated gamma expansions of X_1 with $K = 1$ and $K = 20$, and the two are very close to each other. At smaller values of t (not reported here), they are even closer. We note that the computing time of moments, λ_n 's and γ_n 's is not included in reported times as they can be pre-cached and it takes only a fraction of a second to compute them (e.g., 0.032 seconds for case I).

In contrast to the transform inversion method, the gamma expansion is not very sensitive to the level of $v_0 + v_t$. For example, in case IV, the computing of 1.86 reported in Table 4.3 with $v_0 + v_t = 0.02$ becomes 2.24 at $v_0 + v_t = 0.2$ and 3.73 at $v_0 + v_t = 2$. This variation is modest, particularly in comparison to the values in Table 4.2.

Table 4.3: Time to generate 10,000 samples of X_1 with $v_0 + v_t = 0.02$ using gamma expansion at truncation level K

case	$K=1$	$K=20$
I	1.66	22.23
II	1.88	23.14
III	1.93	23.18
IV	1.86	24.27

4.2 Simulation of X_2 and X_3

Simulation of X_2 . The methods discussed for X_1 — sampling through repeated transform inversion or through the gamma expansion — are both applicable to X_2 as well. However, the fact that X_2 does not depend on v_0 or v_t opens a further option of tabulation. If we fix parameters κ , θ and σ (and therefore δ) and fix a time step t , then the distribution of X_2 does not depend on any intermediate simulated V_t values. We can therefore tabulate the distribution at the start of the simulation and then draw samples as needed by sampling from the table.

In implementing this approach, we first compute $F_{X_2}(x_i)$, the distribution function of X_2 at points

$$x_i = w\mu_{X_2} + \frac{i-1}{M}(u_\epsilon - w\mu_{X_2}), \quad i = 1, \dots, M+1,$$

with $M = 200$ and $w = 0.01$. We compute these values using transform inversion; because the distribution remains fixed, the total number of inversions is far smaller than in the iterative method described for X_1 . To sample from the tabulated distribution, we generate U uniformly over $[0, 1]$ and set $X_2 = F_{X_2}^{-1}(U)$, using linear interpolation between tabulated values. To accelerate the search through the table, we use the cutpoint method (as described, for example, in Fishman [17], Section 3.3) with 100 cutpoints.

We calculate $F_{X_2}(x_i)$ using the Abate-Whitt algorithm, as in (4.1). To achieve a high degree of precision in the tabulated values, we set $u_\epsilon = \mu_{X_2} + 12\sigma_{X_2}$. The mean and standard deviation of X_2 are given by

$$\mu_{X_2} = \delta\mu_{X_2}^*, \quad \sigma_{X_2}^2 = \delta\sigma_{X_2}^{*2}$$

with

$$\mu_{X_2}^* = \frac{\sigma^2}{4\kappa^2} \left(-2 + \kappa t \coth \left(\frac{\kappa t}{2} \right) \right) \quad (4.2)$$

$$\sigma_{X_2}^{*2} = \frac{\sigma^4}{8\kappa^4} \left(-8 + 2\kappa t \coth \left(\frac{\kappa t}{2} \right) + \kappa^2 t^2 \operatorname{csch}^2 \left(\frac{\kappa t}{2} \right) \right). \quad (4.3)$$

Evaluating the transform (2.13) with $b \in i\mathbb{R}$ for the inversion in (4.1) requires care. Some numerical packages (including MATLAB) use the complex logarithm with the principal branch $(-\pi, \pi]$. This leads to a discontinuity in $\Phi^2(-ib)$ as b moves along the real line and thus to a discontinuous approximation to the distribution of X_2 . This kind of discontinuity is also observed in Broadie and Kaya [8]. To address this issue, we keep track of how many times the expression $(L \sinh \kappa t/2)/(\kappa \sinh Lt/2)$ in (2.13) rotates around the origin as b varies. In our implementation of (4.1), we add 2π whenever $\Phi^2(-ihk)$ crosses the negative real axis, moving from the second quadrant to the third quadrant. With this implementation, the computing times required to tabulate the distribution of X_2 in cases I–IV are 1.41, 1.87, 0.36, and 0.13 seconds, respectively.

Table 4.4: Time to generate 10,000 samples from the Bessel distribution

z	$\delta = 0.1$	$\delta = 2$	$\delta = 6$
1	1.555	1.429	1.303
5	1.461	1.415	1.522
10	1.644	1.553	1.432
50	2.035	1.967	1.878
100	2.425	2.324	2.355
200	3.323	3.194	3.232
300	4.312	4.299	4.238

Simulation of X_3 . We can generate samples of Z from its series expansion or from a precomputed distribution table. The distribution of Z is a special case of that of X_2 ; it is simpler to work with because the exponent in its transform (2.14) is 2, so the complication arising from the complex logarithm can be avoided by evaluating the square as a product, without exponentiation. The mean and variance of Z are given by $\mu_Z = 4\mu_{X_2}^*$ and $\sigma_Z^2 = 4\sigma_{X_2}^{*2}$, with the parameters for X_2 as given in (4.2)–(4.3).

To simulate X_3 , we need to generate the Bessel random variable $\eta \sim BES(\nu, z)$, with parameters

$$\nu = \frac{\delta}{2} - 1, \quad z = \frac{2\kappa/\sigma^2}{\sinh(\kappa t/2)} \sqrt{v_0 v_t},$$

as given in Theorem 2.2. Devroye [13] proposed and analyzed fast acceptance-rejection algorithms for sampling Bessel random variables; Iliopoulos et al. [20] also proposed acceptance-rejection algorithms, using properties of the Bessel law studied in Yuan and Kalbfleisch [32]. However, with acceptance-rejection methods, small changes in parameter values can produce abrupt changes in the samples generated. Given the importance of sensitivity estimates in financial applications, we therefore prefer to avoid rejection methods and sample instead from the inverse of the cumulative distribution function. (Iliopoulos et al. [20] discuss this approach as well.) We generate U uniformly over $[0, 1]$ and calculate the probability mass function recursively,

$$p_{n+1} = \frac{z^2}{4(n+1)(n+1+\nu)} p_n, \quad p_0 = \frac{(z/2)^\nu}{I_\nu(z)\Gamma(\nu+1)}$$

until we reach a value η such that

$$\sum_{n=0}^{\eta-1} p_n < U \leq \sum_{n=0}^{\eta} p_n.$$

We return η as our sample from the Bessel distribution.

It turns out that the computing time with this approach increases with z but is relatively insensitive to δ ; Table 4.4 compares computing times (in seconds) across parameter values. This is convenient for our intended application because typical z -values arising in the Heston model are small. For example, with $v_0 = v_t = \theta$, the z -values for cases I–IV are 0.16, 0.2, 0.35 and 0.06, respectively. For large values of z , one could use the fast approximate sampling method of Yuan and Kalbfleisch [32]. We do not explore that option here because our focus is on exact sampling.

5 Numerical Tests

We now test our method by pricing European call options and comparing the results with the exact scheme of Broadie and Kaya [8] and the QE method of Andersen [5]. We use the parameters in Table 4.1 with $S_0 = 100$, $v_0 = \theta$, and a maturity of one year. We set the truncation level at $K = 1$ or $K = 10$ for the gamma approximations. European options in the Heston model can be priced numerically (without Monte Carlo), so they provide a convenient benchmark for comparison.

In the interest of achieving realistic speed-accuracy tradeoffs, our pricing comparisons use programs coded in the C programming language and compiled by Microsoft Visual C++ 6.0 in release mode. Execution files are run on a personal desktop computer with Intel Pentium 4 CPU 3.20 GHz and 1.0GB of RAM.

We first evaluate the bias resulting from truncation at $K = 1$ and $K = 10$ terms. The bias is small, so to estimate it precisely we use 1 billion simulation trials for each case. Table 5.1 reports the estimated biases and standard errors for at-the-money options in parameter cases I–IV together with exact option prices and confirms that these are indeed very small. Results for in-the-money and out-of-the-money options are reported in Kim [23] and are generally similar to those in Table 5.1, with a few exceptions. The exceptions result from the Abate-Whitt inversion algorithm (4.1) used to tabulate the distributions of X_2 and Z , and not from truncation of the gamma expansion. Increasing the accuracy of (4.1) by setting u_ϵ to be 14 standard deviations above the mean (rather than 12) and setting $\epsilon' = 10^{-7}$ (rather than 10^{-5}) produces consistently small biases in all cases.

Figure 5.1 illustrates the speed-accuracy tradeoff for the Broadie-Kaya exact method and the gamma expansion with $K = 1$ and $K = 10$. Each figure plots the root mean square error (RMSE) against computing time on a log-log scale for each method. The RMSE measure reflects both bias and variance. For $K = 1$, the graphs indicate that as the number of simulation trials increases the bias eventually dominates the RMSE and slows the decrease in the RMSE. But at $K = 10$, the gamma approximation exhibits the same convergence rate as the exact method while reducing the computation time by the factor of 10^2 to 10^3 .

Recall that to sample X_2 and Z , we can tabulate their distributions; this is a one-time computation for each set of parameter values. The computing times required for tabulation in cases I–IV are 0.69, 0.81, 0.44 and 0.36 seconds, respectively. As the number of simulation trials increases, the computational burden of tabulation becomes relatively negligible.

Next, we compare results with Andersen’s [5] QE method. This is a time-discretization algorithm based on approximating the transition of the Heston model over a small time step Δt . Thus, the QE method may take multiple time steps to simulate a one-year option, whereas the Broadie-Kaya algorithm and the gamma expansion simulate to maturity in a single step. Andersen [5] compared the QE method with various alternative discretizations and found that the QE method works best. For the numerical tests, we set the QE parameters at $\gamma_1 = \gamma_2 = 0.5$ and $\psi_C = 1.5$, following [5]. We do not implement the martingale correction scheme, particularly as we are dealing with at-the-money options. See Andersen [5] for the details of the QE method and other variants.

The theoretical convergence rate of the QE method is unknown, but we can estimate the order of the bias $O(\Delta t^\beta)$ numerically. Based on 1 billion simulation trials at time steps ranging from $\Delta t = 1$ to $\Delta t = 1/32$, we estimate β values of 1.28, 1.88, 1.63 and 1.27 for cases I–IV. See Table 5.2 for relevant numbers. (These estimates do indeed compare favorably with a standard Euler approximation, which would typically have $\beta = 1$ for a smooth payoff and a diffusion with smooth coefficients, and which usually performs worse with the Heston model.) With the optimal allocation rule of Duffie and Glynn [14] (which allocates computing effort optimally between

decreasing Δt and increasing the number of replications), the overall convergence rate of the RMSE is $\beta/(1 + 2\beta)$; the β estimates for cases I–IV thus yield estimated optimal convergence rates of 0.36, 0.39, 0.38 and 0.36, respectively. These values should be compared with an optimal rate of $1/2$ for unbiased simulation.

Figure 5.2 illustrates the speed-accuracy tradeoff for the QE method with $\Delta t = 1/8$ and $\Delta t = 1/32$ and the gamma expansion truncated at $K = 10$. In cases II and III, the QE method with $\Delta t = 1/8$ performs slightly better up to 160,000 simulation trials. With a larger step size, the QE method is faster, and with a smaller number of replications the bias from a larger step size is obscured by the variance, as both affect the RMSE. However, the gamma expansion shows better overall performance, particularly at high levels of precision. It exhibits a faster convergence rate, achieving a slope very close to the optimum of $-1/2$. Table 5.3 reports detailed simulation results for the QE method for comparison with Table 5.1. Similar results are reported in Kim [23] for in-the-money and out-of-the-money options.

The results in Figure 5.2 compare the QE and gamma expansion methods at a single date — a maturity of one year. But the QE method with $\Delta t = 1/32$ generates 31 intermediate values of the process in sampling the terminal value, whereas the gamma expansion method samples the terminal value directly. This is a potential advantage of the QE method in pricing path-dependent options for which the intermediate values are needed. In the numerical examples above, the gamma expansion method outperforms the QE method with $\Delta t = 1/32$ by a factor of 2–3. This suggests that generating all 32 values would take approximately 12 times as long using the gamma expansion. This is only a rough comparison because it does not incorporate potential adjustments to the gamma expansion parameters to tailor the method to a short time step.

In pricing Bermudan options, it is useful to be able to simulate backwards in time using bridge sampling or the approach developed by Dutt and Welke [16] for the method of Longstaff and Schwartz [25]. Yuan and Kalbfleisch [32], p.444, provide a sampling algorithm for a squared Bessel bridge which can be used to simulate values of V_t backwards in time; at each step, we can then use (2.5) and (2.4) to simulate values of S_t in reverse time as well. Thus, the gamma expansion method is potentially useful in pricing Bermudan options as well as European options.

6 Extension to Non-Equidistant Time Grids

In our discussion of Theorem 2.2, and particularly in the implementation in Section 4, we have highlighted the distinction between terms that depend on the endpoints (V_0, V_t) and terms that do not, but we have implicitly taken the length of the time interval t to be fixed. In practice, one might want to simulate the asset price process at dates $0 < t_1 < \dots < t_m$ without restricting these to be equally spaced.

The decomposition in Theorem 2.2 continues to apply over each transition from t_i to t_{i+1} , with the parameter t in the theorem set to $t_{i+1} - t_i$, even if these time increments vary along the simulated path. However, some of the shortcuts in Section 4 require modification if the time steps vary. In particular, Section 4 exploits the fact that the distributions of X_2 and Z do not depend on the endpoints (V_0, V_t) and thus do not vary from one simulation step to another. Their distributions do, however, depend on the value of t through the coefficients γ_n defined in Theorem 2.2; t also appears in z and the λ_n . If the value of t varies, one cannot simply use tabulated distributions to generate X_2 and Z as required for Theorem 2.2.

In this section, we present an extension of Theorem 2.2 that removes all dependence on t through a measure transformation. In order to consider a change of probability measure, we will suppose that the original model (2.1)–(2.2) is specified under a probability \mathbb{Q} ; in other words,

Table 5.1: Estimated biases of gamma approximation with standard errors in parentheses

	case I	case II	case III	case IV
Option Price	6.7304	7.0972	11.3743	7.0200
Gamma(K=1)	0.00303 (0.00015)	0.00969 (0.00033)	-0.00409 (0.00062)	0.00417 (0.00025)
Gamma(K=10)	-0.00027 (0.00015)	-0.00164 (0.00033)	-0.00121 (0.00062)	0.00093 (0.00025)

Table 5.2: Estimated biases of the QE method with standard errors in parentheses

steps	case I	case II	case III	case IV
1	-0.27287 (0.00018)	0.71166 (0.00033)	0.66345 (0.00061)	0.36716 (0.00026)
2	-0.10063 (0.00016)	0.22785 (0.00033)	0.17704 (0.00062)	0.54419 (0.00026)
4	-0.02954 (0.00016)	0.04896 (0.00033)	0.05043 (0.00062)	0.22331 (0.00025)
8	-0.01540 (0.00016)	0.01108 (0.00033)	0.02422 (0.00062)	0.06485 (0.00025)
16	-0.01034 (0.00016)	0.00463 (0.00033)	0.00709 (0.00062)	0.01704 (0.00025)
32	-0.00322 (0.00015)	0.00104 (0.00033)	0.00237 (0.00062)	0.00451 (0.00025)

Table 5.3: Simulation Results for cases I-IV

No. of trials	Gamma(K=1)		Gamma(K=10)		QE	Exact	RMSE		
	time	RMSE	time	RMSE					
10,000	0.031	0.0498	0.062	0.0489	8	0.031	0.0513	26.859	0.0486
40,000	0.156	0.0247	0.219	0.0246	8	0.125	0.0291	117.421	0.0244
160,000	0.609	0.0127	0.875	0.0123	8	0.484	0.0197	450.781	0.0122
640,000	2.484	0.0068	3.562	0.0061	32	7.906	0.0069	1775.984	0.0061
2,560,000	9.781	0.0043	14.109	0.0031	32	31.734	0.0044		
10,240,000	38.875	0.0034	56.265	0.0016	32	126.093	0.0036		
No. of trials	Gamma(K=1)		Gamma(K=10)		QE	Exact	RMSE		
	time	RMSE	time	RMSE					
10,000	0.031	0.1052	0.047	0.1050	8	0.031	0.1039	25.609	0.1047
40,000	0.156	0.0533	0.218	0.0507	8	0.125	0.0532	110.734	0.0515
160,000	0.609	0.0280	0.859	0.0257	8	0.515	0.0285	426.734	0.0262
640,000	2.484	0.0163	3.438	0.0132	8	2.046	0.0172	1735.578	0.0130
2,560,000	9.765	0.0117	13.594	0.0068	16	15.796	0.0080		
10,240,000	38.875	0.0102	54.156	0.0037	32	124.328	0.0034		
No. of trials	Gamma(K=1)		Gamma(K=10)		QE	Exact	RMSE		
	time	RMSE	time	RMSE					
10,000	0.047	0.1977	0.063	0.1883	8	0.031	0.1950	8.906	0.1937
40,000	0.172	0.0960	0.250	0.0973	8	0.141	0.1006	36.984	0.0980
160,000	0.640	0.0487	0.953	0.0489	8	0.562	0.0549	156.218	0.0489
640,000	2.641	0.0249	3.906	0.0248	16	4.421	0.0255	621.109	0.0245
2,560,000	10.375	0.0130	15.406	0.0124	16	17.546	0.0142	2397.000	0.0123
10,240,000	41.328	0.0074	61.453	0.0063	32	139.484	0.0066		
No. of trials	Gamma(K=1)		Gamma(K=10)		QE	Exact	RMSE		
	time	RMSE	time	RMSE					
10,000	0.046	0.0788	0.046	0.0784	8	0.031	0.1021	3.140	0.0795
40,000	0.140	0.0401	0.203	0.0394	16	0.296	0.0431	12.265	0.0397
160,000	0.593	0.0202	0.859	0.0198	16	1.172	0.0261	48.765	0.0198
640,000	2.422	0.0107	3.484	0.0099	32	9.312	0.0109	195.296	0.0099
2,560,000	9.500	0.0065	13.718	0.0050	32	37.031	0.0067	779.859	0.0049
10,240,000	37.828	0.0048	54.718	0.0026	32	148.172	0.0051	3106.250	0.0025

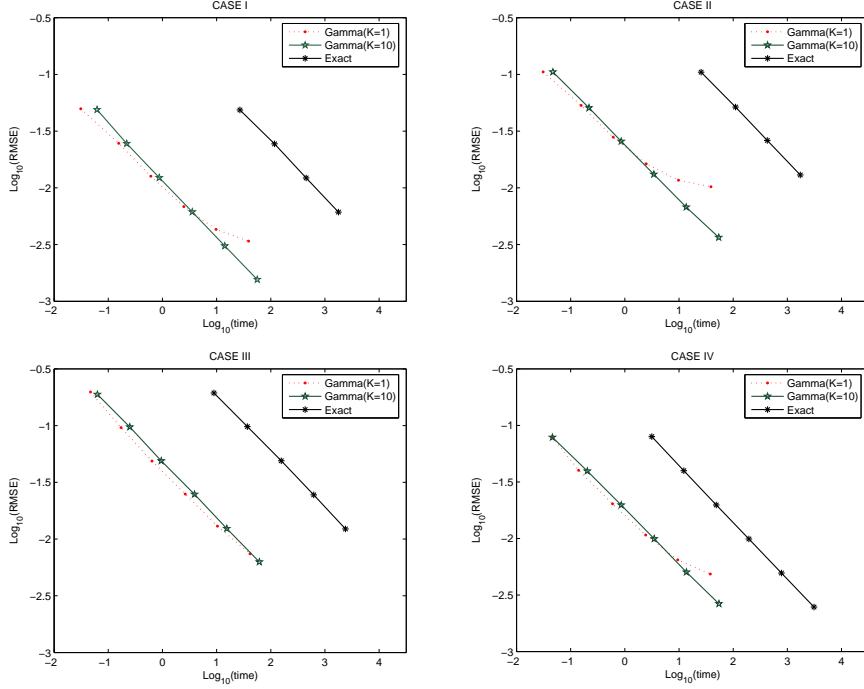


Fig. 5.1: Speed and accuracy comparisons for truncated gamma expansions and the Broadie-Kaya exact method in pricing a European call option.

(W_t^1, W_t^2) is a standard two-dimensional Brownian motion under \mathbb{Q} . We will introduce a new probability \mathbb{P} .

It will be convenient to rescale time and define $\tilde{A}_t = V_{4t/\sigma^2}$, which then satisfies

$$d\tilde{A}_t = \left(\delta - \frac{4\kappa}{\sigma^2} \tilde{A}_t \right) dt + 2\sqrt{\tilde{A}_t} d\tilde{W}_t^1,$$

where $(\tilde{W}_t^1, \tilde{W}_t^2) = (\sigma W_{4t/\sigma^2}^1/2, \sigma W_{4t/\sigma^2}^2/2)$ is a standard two-dimensional Brownian motion under \mathbb{Q} . Accordingly, we set $\tilde{S}_t = S_{4t/\sigma^2}$, which satisfies

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = \frac{4\mu}{\sigma^2} dt + \frac{2}{\sigma} \sqrt{\tilde{A}_t} \left(\rho d\tilde{W}_t^1 + \sqrt{1 - \rho^2} d\tilde{W}_t^2 \right).$$

Because S and \tilde{S} are related through a deterministic rescaling of time, any payoff that is a function of the level of S at a fixed set of dates may be written as a payoff $p(\tilde{S}_{t_1}, \dots, \tilde{S}_{t_m})$ on the path of \tilde{S} . Our objective, then, is to evaluate the \mathbb{Q} -expectation of such a payoff.

Define a new probability measure \mathbb{P} by setting

$$\frac{d\mathbb{P}}{d\mathbb{Q}} = \exp \left(q \int_0^t \sqrt{\tilde{A}_s} d\tilde{W}_s^1 - \frac{q^2}{2} \int_0^t \tilde{A}_s ds \right), \quad (6.1)$$

with $q = 2\kappa/\sigma^2$. To evaluate the \mathbb{Q} -expectation of the payoff, we can use the identity

$$\mathbb{E}^{\mathbb{Q}} \left[p(\tilde{S}_{t_1}, \dots, \tilde{S}_{t_m}) \right] = \mathbb{E}^{\mathbb{P}} \left[p(\tilde{S}_{t_1}, \dots, \tilde{S}_{t_m}) \frac{d\mathbb{Q}}{d\mathbb{P}} \right], \quad (6.2)$$

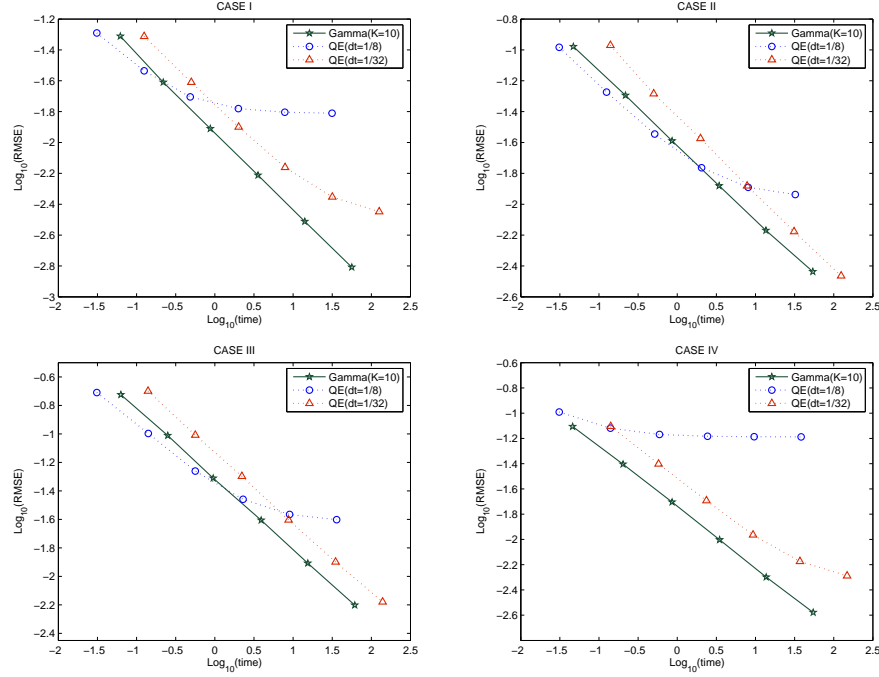


Fig. 5.2: Speed and accuracy comparisons for the truncated gamma expansion ($K = 10$) and the QE method with step sizes $\Delta t = 1/8$ and $\Delta t = 1/32$.

with $d\mathbb{Q}/d\mathbb{P}$ given by the reciprocal of (6.1). The problem, then, is to simulate $(\tilde{S}_{t_1}, \dots, \tilde{S}_{t_m})$ and $d\mathbb{Q}/d\mathbb{P}$ under \mathbb{P} .

By the Girsanov Theorem, $dW_t^{\mathbb{P}} := d\tilde{W}_t^1 - q\sqrt{\tilde{A}_t}dt$ becomes a standard Brownian motion under \mathbb{P} . Writing the dynamics of \tilde{A} as

$$d\tilde{A}_t = \delta dt + 2\sqrt{\tilde{A}_t}dW_t^{\mathbb{P}},$$

we see that \tilde{A} is a δ -dimensional squared Bessel process under \mathbb{P} . We also have

$$\frac{d\tilde{S}_t}{\tilde{S}_t} = \left(\frac{4\mu}{\sigma^2} + \frac{2\rho q}{\sigma} \tilde{A}_t \right) dt + \frac{2}{\sigma} \sqrt{\tilde{A}_t} \left(\rho dW_t^{\mathbb{P}} + \sqrt{1 - \rho^2} d\tilde{W}_t^2 \right),$$

and \tilde{W}^2 remains a standard Brownian motion independent of $W^{\mathbb{P}}$ under \mathbb{P} .

The Broadie-Kaya [8] scheme (2.3)–(2.5) for the Heston model extends to the law of (\tilde{S}, \tilde{A}) under \mathbb{P} . The transition law of \tilde{A} is again noncentral chi-square (set $\kappa = 0$ and $\sigma = 1$ in (2.3) with δ held fixed), and exact sampling of \tilde{S} reduces to the problem of sampling from the conditional law

$$\left(\int_s^t \tilde{A}_u du \mid \tilde{A}_s, \tilde{A}_t \right), \quad (6.3)$$

for arbitrary $s < t$. Moreover, the two integrals appearing in the exponent of $d\mathbb{Q}/d\mathbb{P}$ are byproducts of the Broadie-Kaya scheme applied to (\tilde{S}, \tilde{A}) , much as the integrals in (2.4) are. Thus, we can sample the Radon-Nikodym derivative *exactly*, to get an unbiased estimator of (6.2), if we can simulate (6.3).

If we define $A'_u = \tilde{A}_{s+\alpha u}/\alpha$ for some $\alpha > 0$, then A' is still a δ -dimensional squared Bessel process and

$$\left(\int_s^t \tilde{A}_u du \mid \tilde{A}_s = a_1, \tilde{A}_t = a_2 \right) = \left(\alpha^2 \int_0^{(t-s)/\alpha} A'_u du \mid A'_0 = \frac{a_1}{\alpha}, A'_{(t-s)/\alpha} = \frac{a_2}{\alpha} \right).$$

By setting $\alpha = t - s$, we reduce the problem of simulating (6.3) to the problem of simulating

$$\left(\int_0^1 A'_u du \mid A'_0 = a_1/(t-s), A'_1 = a_2/(t-s) \right).$$

In this expression, the length of the time increment $t - s$ appears only in the values of the endpoints on which we condition. As an application of the decomposition of squared Bessel bridges in Pitman and Yor [27], we obtain an expansion similar to that in Theorem 2.2.

Theorem 6.1 *The distribution of the integral of A' conditional on endpoints A'_0, A'_1 admits the representation*

$$\left(\int_0^1 A'_u du \mid A'_0 = a'_0, A'_1 = a'_1 \right) \stackrel{d}{=} Y_1 + Y_2 + Y_3 \equiv Y_1 + Y_2 + \sum_{j=1}^{\eta'} Z'_j$$

in which $Y_1, Y_2, \eta', Z'_1, Z'_2, \dots$ are mutually independent, the Z'_j are independent copies of a random variable Z' , and η' is a Bessel random variable with parameters $\nu = \delta/2 - 1$ and $z = \sqrt{a'_0 a'_1}$. Moreover, Y_1, Y_2 and Z' have the following representations:

$$Y_1 = \sum_{n=1}^{\infty} \frac{2}{\pi^2 n^2} \sum_{j=1}^{N_n} \text{Exp}_j(1), \quad Y_2 = \sum_{n=1}^{\infty} \frac{2}{\pi^2 n^2} \Gamma_n(\delta/2, 1), \quad Z' = \sum_{n=1}^{\infty} \frac{2}{\pi^2 n^2} \Gamma_n(2, 1)$$

where the N_n are independent Poisson random variables with mean $a'_0 + a'_1$.

In this representation, dependence on the time increment $t - s$ is absorbed into the endpoints a'_0 and a'_1 . As in Theorem 2.2, dependence on these endpoints is limited to Y_1 and η' . In particular, the distributions of Y_2 and Z' could be tabulated and re-used over time steps of different lengths.

Remark 6.2 To take this observation a step further, note that the distribution of Z' is universal, in the sense that it does not depend on any parameters of the model or simulation; this distribution could be tabulated for once and for all and then re-used with any model parameters and time steps. Indeed, the model parameters appear in the decomposition of Theorem 6.1 only through the value of $\delta = 4\kappa\theta/\sigma^2$, which appears in Y_2 and η . This feature of Theorem 6.1 is potentially valuable in model calibration, where one is interested in iterative calculation of prices while varying parameters in order to find parameters that best fit market prices.

Remark 6.3 One can extend the gamma approximation to variants of the Heston model such as the SVJ and SVCJ models, following the approach used by Broadie and Kaya [8] for exact simulation of these models. These extensions are straightforward and explained well in [8], so we do not consider them further here.

7 Concluding Remarks

We have developed a representation of the transition law of the Heston stochastic volatility model using sums and mixtures of gamma random variables for the volatility process and mixtures of normal random variables for the log asset price. This expansion builds on connections between the volatility process and Bessel processes and results of Pitman and Yor [27] on the decomposition of Bessel bridges. When combined with the method of Broadie and Kaya [8], the gamma expansion allows exact sampling of the transitions of the Heston model. We have also analyzed the error resulting from truncating infinite series in this expansion and proposed some computational shortcuts. Numerical examples illustrate the speed and accuracy of the approach.

The methods developed here are applicable in other contexts as well, particularly in light of the widespread use of the square-root diffusion (2.2) as an element of other models. Similar ideas should be applicable to the constant elasticity of variance model, given its connections to Bessel processes; see Delbaen and Shirakawa [12]. We note two other examples. In the Cox-Ingersoll-Ross [10] model, the square-root diffusion models a continuously compounded instantaneous interest rate, and its integral appears in the exponent of the discount factor. Theorem 2.2 can be used for exact simulation of the discount factor along with the interest rate over long time horizons; indeed, this was the motivating application in Scott's [30] work.

The square-root diffusion is also a standard model of a stochastic default intensity. In this type of model, default occurs the first time the integrated intensity reaches the level of an independently sampled exponential random variable. Simulating the default time thus requires simulating the square-root diffusion together with its time integral, and this is facilitated by Theorem 2.2.

Appendix

Proof (Lemma 2.4) The calculations use a change of measure argument, so we begin with a process \tilde{W} that is a Brownian motion with respect to a measure $\tilde{\mathbb{Q}}$. A process B^0 satisfying $dB_t^0 = 2\sqrt{B_t^0}d\tilde{W}_t$ is then a 0-dimensional squared Bessel process with respect to $\tilde{\mathbb{Q}}$. From Revuz and Yor [29], Corollary 3.3, p.465, we have the Laplace transform

$$\tilde{\mathbb{E}} \left[\exp \left(-\frac{b^2}{2} \int_0^1 B_s^0 ds \right) \middle| B_0^0 = x, B_1^0 = 0 \right] = \exp \left(\frac{x}{2} (1 - b \coth b) \right)$$

for $b \in \mathbb{R}$ and $x \geq 0$, where $\tilde{\mathbb{E}}$ means expectation under $\tilde{\mathbb{Q}}$.

We now apply the change of measure formula (6.d) in Pitman and Yor [27], under which the law of the squared OU bridge $A_{x+y,0}^{0,1}$ becomes the law of a squared Bessel bridge with $B_0^0 = x + y$ and $B_1^0 = 0$. For $b \geq 0$ we have (with $a = -\kappa t/2$),

$$\begin{aligned} \Phi^1(b) &= \mathbb{E} \left[\exp \left(-b \frac{\sigma^2 t^2}{4} \int_0^1 A_{x+y,0}^{0,1}(s) ds \right) \right] \\ &= \tilde{\mathbb{E}} \left[\exp \left(- \left(\frac{b\sigma^2 t^2}{4} + \frac{a^2}{2} \right) \int_0^1 A_{x+y,0}^{0,1}(s) ds \right) \right] \div \tilde{\mathbb{E}} \left[\exp \left(-\frac{a^2}{2} \int_0^1 A_{x+y,0}^{0,1}(s) ds \right) \right] \\ &= \exp \left(\frac{(v_0 + v_t)}{\sigma^2} \left(\kappa \coth \frac{\kappa t}{2} - L \coth \frac{Lt}{2} \right) \right) \end{aligned}$$

with $L = \sqrt{2\sigma^2 b + \kappa^2}$. This is also how the Laplace transform of (2.5) is derived in Broadie and Kaya [8].

For X_2 and Z , we start from the Laplace transform

$$\tilde{\mathbb{E}} \left[\exp \left(-\frac{b^2}{2} \int_0^1 B_s^f ds \right) \middle| B_0^f = B_1^f = 0 \right] = \left(\frac{b}{\sinh b} \right)^{f/2}$$

of the f -dimensional squared Bessel process B^f , $dB_t^f = f dt + 2\sqrt{B_t^f} d\tilde{W}_t$ under $\tilde{\mathbb{Q}}$. The same change of measure argument used above changes the law of $A_{0,0}^{f,1}$, $f > 0$, to that of B^f conditional on $B_0^f = B_1^f = 0$, and a similar calculation then completes the proof. \square

Acknowledgements The authors thank Mark Broadie and Özgür Kaya for helpful discussions and comments and for providing source code for their method and Leif Andersen and a referee for helpful comments. This research is supported in part by NSF grant DMS0410234.

References

1. Abate, J., Whitt, W.: The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems* **10**, 5–88 (1992)
2. Ahlfors, L.V.: *Complex Analysis*, 3rd and international edition. McGraw-Hill, New York (1979)
3. Alfonsi, A.: On the discretization schemes for the CIR (and Bessel squared) processes. *Monte Carlo Methods and Applications* **11**, 355–384 (2005)
4. Alfonsi, A.: High order discretization schemes for the CIR process: application to affine term structure and Heston models. *Mathematics of Computation*, forthcoming. http://cermics.enpc.fr/~alfonsi/papiers_fr.html (2008)
5. Andersen, L.: Efficient simulation of the Heston stochastic volatility model. Working paper, Banc of America Securities, New York. <http://papers.ssrn.com> (2007)
6. Bondesson, L.: On simulation from infinitely divisible distributions. *Advances in Applied Probability* **14**, 855–869 (1982)
7. Berkaoui, A., Bossy, M., Diop, A.: Euler scheme for SDEs with non-Lipschitz diffusion coefficient: strong convergence. *ESAIM: Probability and Statistics* **12**, 1–11 (2008)
8. Broadie, M., Kaya, Ö.: Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* **54**, 217–231 (2006)
9. Cont, R., Tankov, P.: *Financial Modelling with Jump Processes*. Chapman & Hall/CRC (2004)
10. Cox, J.C., Ingersoll, J.E., Ross, S.A.: A theory of the term structure of interest rates. *Econometrica* **53**, 1943–1978 (1985)
11. Deelstra, G., Delbaen, F.: Convergence of discretized stochastic (interest rate) processes with stochastic drift term. *Applied Stochastic Models and Data Analysis* **14**, 77–84 (1998)
12. Delbaen, F., Shirakawa, H.: No arbitrage condition for positive diffusion price processes. *Asia-Pacific Financial Markets* **9**, 159–168 (2002)
13. Devroye, L.: Simulating Bessel random variables. *Statistics & Probability Letters* **57**, 249–257 (2002)
14. Duffie, D., Glynn, P.: Efficient Monte Carlo simulation of security prices. *Annals of Applied Probability* **5**, 897–905 (1995)
15. Duffie, D., Pan, J., Singleton, K.: Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* **68**, 1343–1376 (2000)
16. Dutt, S.K., Welke, G.M.: Just-in-time Monte Carlo for path-dependent american options. *Journal of Derivatives*, summer 29–47 (2008)
17. Fishman, G.S.: *Monte Carlo: Concepts, Algorithms, and Applications*. Springer-Verlag, New York (1996)
18. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* **13**, 585–625 (1993)
19. Higham, D.J., Mao, X.: Convergence of Monte Carlo simulations involving the mean-reverting square root process. *Journal of Computational Finance* **8**, 35–62 (2005)
20. Iliopoulos, G., Karlis, D.: Simulation from the Bessel distribution with applications. *Journal of Statistical Computation and Simulation* **73**, 491–506 (2003)
21. Kahl, C., Jäckel, P.: Fast strong approximation Monte Carlo schemes for stochastic volatility models. *Quantitative Finance* **6**, 513–536 (2006)
22. Kallenberg, O.: *Foundations of Modern Probability*, 2nd edition. Springer-Verlag, Berlin (2002)
23. Kim, K.-K.: *Affine Processes in Finance: Numerical Approximation, Simulation and Model Properties*. Ph.D. thesis, Columbia University (2008)
24. Kloeden, P.E., Platen, E.: *Numerical Solution of Stochastic Differential Equations*, 3rd edition. Springer-Verlag, Berlin (1999)
25. Longstaff, F.A., Schwartz, E.S.: Valuing american options by simulation: a simple lease-squares approach. *Review of Financial Studies*, **14**, 113–147 (2001)

26. Lord, R., Koekoek, R., van Dijk, D.: A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, forthcoming. <http://www.rogerlord.com> (2008)
27. Pitman, J., Yor, M.: A decomposition of Bessel bridges. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **59**, 425–457 (1982)
28. Pitman, J., Yor, M.: Infinitely divisible laws associated with hyperbolic functions. *Canadian Journal of Mathematics* **55**, 292–330 (2003)
29. Revuz, D., Yor, M.: *Continuous Martingales and Brownian Motion*, 3rd edition. Springer-Verlag, New York (1999)
30. Scott, L.O.: Simulating a multi-factor term structure model over relatively long discrete time periods, in *Proceedings of the IAFE First Annual Computational Finance Conference*. Graduate School of Business, Stanford University (1996)
31. Van Haastrecht, A., Pelsser, A.: Efficient, almost exact simulation of the Heston stochastic volatility model. Working paper, University of Amsterdam. <http://papers.ssrn.com> (2008)
32. Yuan, L., Kalbfleisch, J.D.: On the Bessel distribution and related problems. *Annals of the Institute of Statistical Mathematics* **52**, 438–447 (2000)