Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

CHANGE OF MEASURE FOR THE SQUARE-ROOT PROCESS

Daniel Dufresne

Centre for Actuarial Studies
University of Melbourne
111 Barry Street
Carlton VIC 3053, AUSTRALIA

Felisa Vázquez-Abad

Department of Computer Science
Hunter College CUNY
695 Park ave
New York NY 10021, USA

Stephen Chin

Ingensoma Arbitrage Ltd
Hanwha Securities Bldg 22F
23-5 Youngdeungpo-gu
Seoul, KOREA, 150-717

ABSTRACT

The square-root process is used to model interest rates and volatility in financial mathematics. The pricing of derivatives involving that process often requires simulating it, since there are often no explicit formulas for prices. We study how a change of measure (CM) may improve those simulations. We compare with Andersen's quadratic-exponential scheme (QE), which so far appears to be the most efficient technique to simulate the stochastic differential equation satisfied by the square-root process. An integer-dimension squared Bessel process, easy to simulate, is used to generate the law of the square-root process using a change of measure. The new method performs very well, and the two algorithms execute at similar speeds; however, CM is slower than QE if random number generation is taken into account, because CM requires more random numbers. The Radon-Nikodym derivative sometimes has a rather intriguing behavior, which is itself of interest. We propose an explanation.

1 INTRODUCTION

Suppose W is standard Brownian motion. The stochastic differential equation (SDE)

$$dV_t \;=\; (aV_t + b)\,dt + c\sqrt{V_t}\,dW_t, \quad V_0 \;=\; v_0, \tag{1}$$

has a unique strong solution when $a \in \mathbb{R}$, $b, v_0 \geq 0$ and $c > 0$ (Revuz and Yor 1999, Chapter 9). The solution is called the square-root process. This process has a number of explicit properties, one of them to remain non-negative at all times. For appropriately chosen parameters it also has a stationary distribution. In financial mathematics, this process has been used as a model for interest rates (Cox, Ingersoll, and Ross 1985) and for squared volatility (Heston 1993). Despite the explicit results known since the 1950s (Feller 1951) about the square-root process, there is no closed form expression for the prices of many options involving that process. As a consequence, pricing those derivatives requires simulating (1). This is not a straightforward exercise, since the discretized version of (1), when generated in the usual way, will take negative values with positive probability at every step. Even if $\sqrt{V}$ (that multiplies dW) is replaced with $\sqrt{|V|}$ the resulting simulated process is not a good approximation of the true square-root process.

Andersen (2008) reviews the previous schemes proposed to simulate this process, and describes one, the "quadratic-exponential" (QE) scheme, that does better than the previous ones. The QE scheme is base on an approximation of the transition density function of the process V.

This paper proposes an alternative to the QE scheme, based on the intimate connection between the square-root process and squared Bessel processes. The idea is to simulate a squared Bessel process of integer dimension (see below), which can be done without discretization error, and then apply a change of probability measure to compute expectations involving the square-root process. Section 2 summarizes the needed theorems, Section 3 shows a successful application to option pricing, while Section 4 describes an unexpected phenomenon, that may affect results obtained when applying this method.

An important point is that we are looking for methods that work for path-dependent as well as path-independent options. For the latter, in particular European calls or puts, the Fourier inversion method, proposed by Heston (1993) and others, works fine and is much faster than any simulation scheme. In order to compare the change of measure and QE methods, this paper uses of formulas in Dufresne, Garrido, and Morales (2009); for European, path-independent derivatives those formulas are hundreds or thousands of times faster than any method involving simulating whole paths of the process S (for instance, compare the execution times in Table 2 below); this is simply because the Mellin transform of S_t is known in closed form. Hence, Fourier inversion is the method of choice for all European options with a payoff function that has a closed form Laplace transform, such as calls or puts. In those cases it is overkill to simulate the process, or to find the explicit probability distribution of S_t.

However, in the context of the Heston model there are no explicit formulas or transforms for path-dependent options, for example average or barrier options. For those options it is required to simulate the stock price for at least all the time points where the options are monitored. For instance, a one-year average option with weekly averaging requires obtaining the values of the underlying at least 52 times. Generating the "exact" transition density, as del Baño Rollin, Ferreiro-Castilla, and Utzet (2010) and Glasserman and Kim (2011) have shown, is clearly more complicated and much less efficient than QE, as the transition density needs to be generated anew at every time step. This is why methods that generate an "exact" transition density are not included in our numerical investigations. (The word "exact" is a misnomer, as the known expressions for the density of S_t include either infinite series or integrals, and need to be approximated.)

With this goal in mind (focusing on methods that could be used for path-dependent or American options), our numerical work compares the change of measure and the QE methods for speed and accuracy, for option prices that can also be computed using a third method, namely Fourier inversion. We have also checked whether the methods yielded accurate expectations for moments of the square-root process (integral moments and expectations of logarithms), as the latter can be obtained from explicit formulas. We were able to confirm the excellent accurracy of the QE method (something that was perhaps not covered enough in Andersen's original paper). All those computations show that, despite being based on a rather simplistic approximation, the QE method yields very accurate results.

2 ABSOLUTE CONTINUITY RELATIONSHIPS AMONG SQUARE-ROOT PROCESSES

For the sake of brevity we introduce radial Ornstein-Uhlenbeck (ROU) processes and then note that squared Bessel processes are particular cases; textbooks, see for example Revuz and Yor, describe Bessel processes first and then apply some transformation to obtain ROU processes. Göing-Jaeschke and Yor (2003) parametrize ROU processes as follows.

Definition 1 *For $\lambda \in \mathbb{R}$, $\delta \geq 0$ and $y \geq 0$ the unique strong solution to the SDE*

$$X_t \;=\; x + \int_0^t (\delta + 2\lambda X_s)\, dt + 2\int_0^t \sqrt{X_s}\, dW_s$$

is called the ***squared δ-dimensional radial Ornstein-Uhlenbeck process*** *with parameter λ. The probability law of the process is denoted ${}^{\lambda}Q_x^{\delta}$ or ${}^{\lambda}Q_x^{(\nu)}$, where (here and in the rest of the paper)*

$$\nu \;=\; \frac{\delta}{2} - 1$$

is the ***index*** *that corresponds to the dimension δ.*

Multiplying (1) by $4/c^2$, the square-root process becomes a ROU process, with $\lambda = 2a/c^2$ and dimension $\delta = 4b/c^2$. The transition density of a δ-dimensional ROU processes is known in closed form, and involves the Bessel function of index $\nu = \delta/2 - 1$. This density is computationally too onerous to use in simulation, so we will not need it in this paper. However, the index appears in many formulas, so we will end up using both δ and ν for the same process.

The behaviour of ROU processes near the origin is important in understanding when the method we describe can or cannot be applied. For any λ and $\delta \geq 2$ (*i.e.* for $\nu \geq 0$), with probability one the ROU process does not reach 0; for any λ and $\delta < 2$ (*i.e.* for $\nu < 0$) with probability one the ROU process reaches 0. Therefore, the laws ${}^{\lambda}Q_x^{(\nu)}$ and ${}^{\lambda'}Q_x^{(\nu')}$ cannot be mutually absolutely continuous if $\nu \geq 0$ and $\nu' < 0$. More details may be found in Göing-Jaeschke and Yor (2003). The special case $\lambda = 0$ is at the heart of the method we study in this paper.

Definition 2 *For every $\delta \geq 0$ and $y \geq 0$, the unique strong solution to the SDE*

$$X_t \;=\; x + \delta t + 2\int_0^t \sqrt{X_s}\, dW_s \tag{2}$$

is called the ***square of a δ-dimensional Bessel process*** *starting at x, and is denoted by $BESQ_x^{\delta}$, or $BESQ_x^{(\frac{\delta}{2}-1)}$. We call $\nu = \delta/2 - 1$ the index of the process. Its law on $C(\mathbb{R}_+, \mathbb{R})$ is denoted Q_x^{δ} or $Q_x^{(\nu)}$.*

Combining results from Göing-Jaeschke and Yor (2003) and Chapter 11 of Revuz and Yor (1999), we get the following theorem.

Theorem 1 Let $\mathscr{F}_t = \sigma(X_s, 0 \leq s \leq t)$. For every $\lambda \in \mathbb{R}$, $\mu, \nu \geq 0$ and $x > 0$,

$$\left.\frac{d^{\lambda}Q_x^{(\nu)}}{dQ_x^{(\mu)}}\right|_{\mathscr{F}_t} \;=\; \left(\frac{X_t}{x}\right)^{\frac{\nu-\mu}{2}} \exp\left(\frac{\lambda}{2}\left[X_t - x - (2\nu+2)t\right] - \frac{\lambda^2}{2}\int_0^t X_s\, ds + \frac{\mu^2-\nu^2}{2}\int_0^t \frac{ds}{X_s}\right). \tag{3}$$

This says that there is an explicit Radon-Nikodym derivative that allows one to calculate expectations under ${}^{\lambda}Q_x^{(\nu)}$ by simulating $Q_x^{(\mu)}$. The squared Bessel process of integer dimension is especially easy to simulate, from the following classical result.

Theorem 2 *Let $\beta \in \{1, 2, 3, \dots\}$, $y_j \in \mathbb{R}$, and suppose $W^{(1)}, \dots, W^{(\beta)}$ is a vector of independent standard Brownian motions issued from 0. Then the law of the* ***process***

$$\{(y_1 + W_t^{(1)})^2 + \cdots + (y_\beta + W_t^{(\beta)})^2, t \geq 0\}$$

is $Q_x^{\beta} = Q_x^{(\mu)}$, where $\mu = \beta/2 - 1$ and $x = y_1^2 + \cdots + y_\beta^2$.

These facts say that to calculate expectations under ${}^{\lambda}Q_x^{\delta} = {}^{\lambda}Q_x^{(\frac{\delta}{2}-1)}$, $\delta \geq 2$, it is sufficient to simulate the squared norm of a β-dimensional Browian motion, with $\beta \in \{2, 3, \dots\}$. Note that $\beta = 1$ will not work, because Theorem 1 does not apply if $\mu = \beta/2 - 1 < 0$. This implies that the change of measure method needs more random numbers than other methods for simulating a square-root process, since at each step the β-dimensional Brownian motion needs to be updated. However, the simulation of the squared Bessel process is exact, in the sense that there is no discretization error.

Table 1: Square-root process estimated parameters from various authors

Reference	a	b	c	δ	ρ	Data
Ait-Sahali and Kimmel (2007)	-5.070	0.232	0.480	4.02	-0.77	S&P 500 option
Bakshi, Cao, and Chen (1997)	-2.180	0.174	0.530	2.48	-0.70	S&P 500 option
Bollerslev and Zhou (2002)	-0.146	0.076	0.579	0.90	0.00	Deutsch Mark/US$
Dragulescu and Yakovenko (2002)	-11.350	0.250	0.618	2.62	0.00	Dow Jones index
Duffie, Pan, and Singleton (2000)	-6.210	0.118	0.610	1.27	-0.70	S&P 500 option
Eraker (2004)	-0.017	0.015	0.107	5.26	-0.37	S&P 500 option
Forbes, Martin, and Wright (2007)	-5.480	0.789	0.719	6.11	0.16	ASX option
Fouque and Lorig (2011)	-3.400	0.082	0.390	2.15	-0.64	S&P 500 index option
Pan (2002)	-5.300	0.128	0.380	3.55	-0.57	S&P 500 index option

3 APPLICATION TO A STOCHASTIC VOLATILITY MODEL

Heston (1993) proposed the following model for stock prices. The squared volatility V and the stock price S are driven by correlated standard Brownian motions $W_t^{(S)}, W_t^{(V)}$:

$$dS_t = rS_t\,dt + \sqrt{V_t}S_t\,dW_t^{(S)} \tag{4}$$

$$dV_t = (aV_t + b)\,dt + c\sqrt{V_t}\,dW_t^{(V)}. \tag{5}$$

The first SDE is the one in the Black-Scholes model, except that the "σ" is replaced with the stochastic volatility process $\sqrt{V_t}$. The squared volatility follows a square-root process. Heston derived Fourier inversion formulas for ordinary European calls and puts, but other derivatives, in particular path-dependent ones, do not have explicit expressions.

Several authors have fitted the square-root process to observed volatility, see Table 1. The change of measure method we describe only applies when $\delta \geq 2$, but two of the entries are smaller than 2. When $0 < \delta < 2$ the process reaches 0 with positive probability and is instantaneously reflected (Göing-Jaeschke and Yor 2003, p.315). We will use the parameters estimated in Eraker (2004), with $v_0 = 0.15$, to illustrate how the change of measure compares with the QE scheme. All computations were performed on a MacBookPro 11,3 2.6 GHz. Simulations were coded in C.

3.1 Option Prices by Inverse Mellin Transform

To compute the exact prices of the European call options we use to compare the CM and QE schemes, we could use the formulas derived in Heston (1993). We will instead employ a much more general formula, from Dufresne, Garrido, and Morales (2009), that applies Parseval's Theorem to the complex Mellin transform $\mathbf{E}S^{iu}$, for any non-negative variable S.

Theorem 3 *Let $S \geq 0$, $K > 0$ and*

$$h(u) = \frac{K^{-iu+1}}{iu(iu-1)}\mathbf{E}(S^{iu}).$$

(a) If $\mathbf{E}(S^{\alpha}) < \infty$ for some $\alpha < 0$, then

$$\mathbf{E}(K-S)_+ = K\mathbf{P}\{S=0\} + \frac{1}{2\pi}PV\int_{-\infty}^{\infty} h(u-i\alpha)\,du.$$

If, moreover, $\mathbf{E}(S) < \infty$, then

$$\mathbf{E}(S-K)_+ = \mathbf{E}S - K\mathbf{P}\{S>0\} + \frac{1}{2\pi}PV\int_{-\infty}^{\infty} h(u-i\alpha)\,du.$$

(b) In all cases,

$$\mathbf{E}(K-S)_+ \;=\; \frac{K}{2}[1+\mathbf{P}\{S=0\}]+\frac{1}{\pi}\int_0^\infty \mathbf{Re}[h(u)]\,du.$$

If $\mathbf{E}(S)<\infty$*, then*

$$\mathbf{E}(S-K)_+ \;=\; \mathbf{E}S-\frac{K}{2}[1+\mathbf{P}\{S=0\}]+\frac{1}{\pi}\int_0^\infty \mathbf{Re}[h(u)]\,du.$$

The Mellin transform of the stock price may be found from the joint Laplace transform of $\left(V_t,\, \int_0^t X_s\,ds\right)$. The latter has been known since Feller (1951):

$$L(p,q) \;=\; \mathbf{E}\exp\left(-pV_t-q\int_0^t V_s\,ds\right) \;=\; [A(p,q,t)]^{-\frac{2b}{c^2}}\exp[B(p,q,t)],$$

where

$$\begin{aligned}
A(p,q,t) &= \frac{e^{at/2}}{P}[P\cosh(Pt/2)-a\sinh(Pt/2)+c^2p\sinh(Pt/2)]\\
B(p,q,t) &= -\frac{v_0}{c^2}\left\{P+a-\frac{(P+a-c^2p)e^{-(P-a)t/2}}{A(t)}\right\}\\
P &= P(q) \;=\; \sqrt{a^2+2c^2q}.
\end{aligned}$$

Let ρ denote the correlation between $W_t^{(S)}$ and $W_t^{(V)}$. From (5),

$$\int_0^t \sqrt{V_s}\,dW_s^{(V)} \;=\; \frac{1}{c}\left(V_t-v_0-a\int_0^t V_s\,ds-bt\right).$$

Then, from (4), S_t has the same distribution as

$$\begin{aligned}
&S_0\exp\left(rt-\frac{1}{2}\int_0^t V_s+\frac{\rho}{c}\left(V_t-v_0-a\int_0^t V_s-bt\right)+\sqrt{1-\rho^2}\left(\int_0^t V_s\right)^{\frac{1}{2}}Z\right)\\
&= S_0e^{rt-\frac{\rho}{c}(v_0+bt)}\exp\left(-\left(\frac{1}{2}+\frac{\rho a}{c}\right)\int_0^t V_s+\frac{\rho}{c}V_t+\sqrt{1-\rho^2}\left(\int_0^t V_s\right)^{\frac{1}{2}}Z\right),
\end{aligned}$$

where Z has a standard normal distribution and is independent of V. Raising this expression to power p and taking conditional expectation given the process V yields

$$\mathbf{E}S_t^p \;=\; S_0^pe^{prt-\frac{\rho p}{c}(v_0+bt)}L\left(-\frac{\rho p}{c},\frac{p}{2}+\frac{\rho ap}{c}-\frac{(1-\rho^2)p^2}{2}\right).$$

In this formula p is replaced with iu, and the result is inserted in the expression for the call price in part (b) of Theorem 3. Mathematica computed the numbers shown in the column labelled "Mellin" in Table 2. (Moments and other properties of the square-root process and its time integral may be found in Dufresne (2001).)

Table 2: Call option prices with $a = -0.017, b = 0.015, c = 0.107, r = 0, \rho = -0.37, v_0 = 0.15$

T	Mellin	CM	QE
K=80			
0.25	21.1333	21.1588 (±0.0340)	21.1530 (±0.0352)
0.50	23.0023	23.0196 (±0.0459)	23.0372 (±0.0478)
1.00	26.3282	26.3716 (±0.0641)	26.3042 (±0.0668)
K=100			
0.25	7.73778	7.75206 (±0.0239)	7.75482 (±0.0246)
0.50	10.9596	10.9728 (±0.0351)	10.9832 (±0.0364)
1.00	15.5441	15.5799 (±0.0530)	15.5298 (±0.0551)
K=120			
0.25	1.92283	1.92918 (±0.0122)	1.93001 (±0.0126)
0.50	4.46886	4.47241 (±0.0233)	4.47674 (±0.0242)
1.00	8.71782	8.74311 (±0.0415)	8.70359 (±0.0430)
Time (sec)	0.0352	142.19	88.96

3.2 Andersen's Quadratic-Exponential Scheme

The method proposed by Andersen (2008) approximates the true transition distribution from V_t to V_{t+h} (the so-called non-central chi-square) with two distributions, one for large values of the current $\hat{V}_t$, and another one for small values of $\hat{V}_t$. The first one is the distribution of

$$\alpha(\beta + Z)^2, \quad Z \sim \mathbf{N}(0,1).$$

The transition distribution used when $\hat{V}_t$ is small is an exponential distribution with parameter γ, mixed with a probability mass π at the origin. There is a boundary ψ separating small and large values of $\hat{V}_t$. The numbers $\alpha, \beta, \gamma, \pi$ and ψ are derived by moment matching in Andersen (2008). They all depend on $\hat{V}_t$. Section 4.2 of that paper describes a way to discretize the SDE for $\log S_t$, which we applied without modification, yielding column "QE" in Table 2. We used $N = 10^6$ iterations, 100 time steps per year, and an expiry t of one year.

3.3 Change of Measure

We use $\beta = 2$, that is, we simulate a two-dimensional squared Bessel process. Once again we use 10^6 iterations, 100 time steps per year, and an expiry of one year. In Table 2 the change of measure (column "CM") performs similarly to the QE scheme. Confidence intervals were most often the same, though on some occasions the confidence intervals under CM were a little wider than those under QE. Computation times are shorter for QE, because of the time it takes to generate random numbers. At every time step QE requires two random numbers, one for each Brownian motion. However, CM requires two random numbers for V and one for $\log S$. If the random numbers are separately generated then execution times of CM were 5-15% shorter than those of QE. There is no surprise in those observations, since CM only requires the simulation of the squared Bessel process and the Radon-Nikodym derivative, while QE involves a few supplementary conditions and computations. For both CM and QE, computation times have a linear relationship to the product of expiry, number of steps and number of iterations.

We have seen very similar results with the other scenarios in Table 1, CM being sometimes a little more precise, sometimes a little less precise than QE, but the differences were almost always rather small, and within the 95% confidence interval. However, after systematically testing the expected value of the Radon-Nikodym derivative for various parameter combinations we did find cases where CM gave results that were off the mark. This is analysed in the next section.

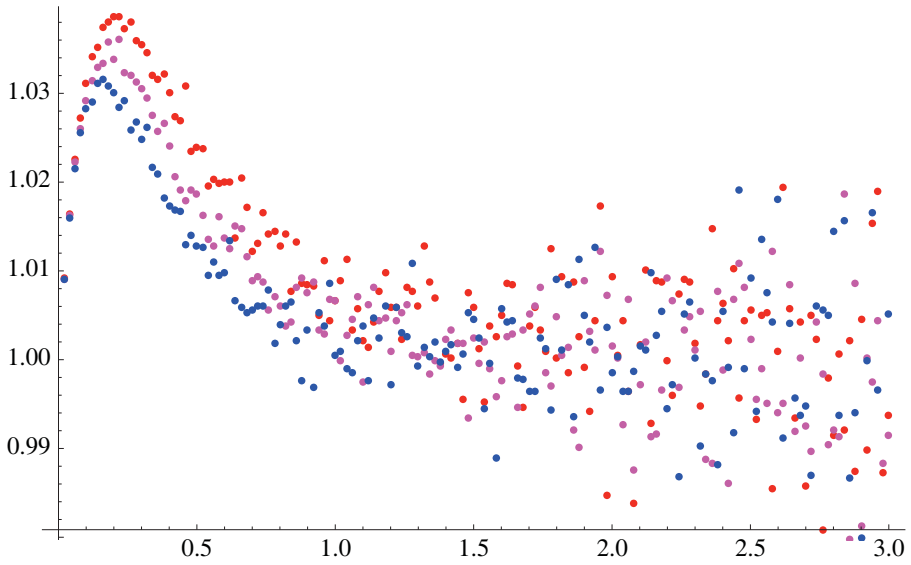

Figure 1: Average simulated values of the Radon Nikodym derivative for $\mu = 0$ as a function of ν, at time $t = 1$. The number of steps is 32 (red), 64 (magenta) and 128 (blue).

4 THE "BUMP": A CAUTIONARY TALE

One of the first things we did was to check that the computed average of the Radon-Nikodym derivative was close to 1. This appeared to happen in most cases, but in some cases even long simulations failed to produce a value close enough to 1. In Figure 1 we show the results when the simulated process is a two-dimensional squared Bessel process (*i.e.* $\mu = 0$) and the Radon-Nikodym derivative is the one given in (3), with $\lambda = 0$, $t = 1$ and $x = 0.5$; the dots correspond to the average of 10^6 replications of the Radon-Nikodym derivative, in steps of 0.02. We used independent random variables (independent replications) for each value of ν. The "bump" is therefore not a fictitious phenomenon due to the use of common random numbers. Not shown in our plot are confidence intervals: these are very small for small values of ν and increase as ν increases. Around the bump the value 1 is *outside* the estimated 95% confidence intervals. This may cause option values to be incorrect. For instance, let

$$a = -5.070, \quad b = 0.232, \quad c = 0.480, \quad v_0 = 0.046,$$

and consider a call option with expiry $t = 1$ and strike 100. Using CM with 32 steps the price was 8.956 ± 0.0659, which is pretty far from the exact price 8.204 (inverse Mellin transform). The computed average Radon-Nikodym derivative was 1.0498. These errors decrease as the number of steps is increased, for example with 256 steps the average Radon-Nikokym derivative was 1.00078 and the option price was 8.2388 ± 0.0636, which is correct. By comparison, the QE method wth 32 steps produced 8.22518 ± 0.02274.

The very surprising fact is that the average Radon-Nikodym derivative is further from the correct value 1 for small ν, while things improve for larger ν. As ν increases the standard deviation increases, but this was to be expected. Near $\nu = 0$ the dots resemble a "bump", and this is the name we used to describe this phenomenon. In all cases we looked at, the bump reduces in height and moves to the left as the number of time steps used increases, but it does not disappear. This behaviour is the opposite of what one would expect, the computed Radon-Nikodym derivative "should" become less accurate as one moves away from the case $\mu = 0$.

To explain what is happening, let us differentiate the Radon-Nikodym derivative with respect to ν:

$$\begin{aligned} R(\mu,\nu,t) &= \exp\left\{\frac{\nu-\mu}{2}\log\left(\frac{X_t}{x}\right)+\frac{1}{2}(\mu^2-\nu^2)J_t\right\} \\ J_t &= \int_0^t \frac{1}{X_s}\,ds \\ \frac{\partial}{\partial\nu}R(0,\nu,t) &= \left(\frac{1}{2}\log\left(\frac{X_t}{x}\right)-\nu J_t\right)R(0,\nu,t). \end{aligned}$$

Since $Q_x^{(\mu)}R(\mu,\nu,t)=1$ for all $\mu,\nu\geq 0$, we know that the partial derivative of $Q_x^{(\mu)}R(\mu,\nu,t)$ with respect to ν is 0. ($Q_x^{(\mu)}(\cdot)$ denotes expectation with respect to $Q_x^{(\mu)}$.) It is straightforward to show that, for any $\nu,x>0$,

$$0 = \frac{\partial}{\partial\nu}Q_x^{(0)}R(0,\nu,t) = Q_x^{(0)}\left[\left(\frac{1}{2}\log\left(\frac{X_t}{x}\right)-\nu J_t\right)R(0,\nu,t)\right]. \tag{6}$$

However, this does not hold for the right derivative at $\nu=0+$, because $Q_x^{(0)}\log X_t\neq 0$. Note also that

$$Q_x^{(0)}\frac{1}{X_s} = \infty \quad\Rightarrow\quad Q_x^{(0)}\int_0^t\frac{1}{X_s}\,ds = \infty.$$

In words, the expectation of the reciprocal of a two-dimensional squared Bessel process is infinite. The expectation on the right of (6) is finite, because there J_t is multiplied by $\exp(-\nu^2 J_t/2)$. If one forces $\nu=0$ on the right-hand side of (6) then one gets a result different from 0, since

$$Q_x^{(0)}\log X_t = Q_x^{(0)}\log(\sqrt{t}Z_1+\sqrt{x})^2+tZ_2^2),$$

where Z_1,Z_2 are two independent standard normals. With $t=1$ and $x=0.5$, as in Figure 1, the last expression equals 0.522. If one computes the rate of change of the observations in Figure 1 between $\nu=0$ and $\nu=.02$, one gets about 0.45, just short of 0.522. Our explanation for this is that the simulated value of J_t appears to understate its true value, implying that the estimation of

$$Q_x^{(0)}\left(J_t R(0,\nu,t)\right)$$

is also understated, at least for small ν.

Another way to look at this is to note that we have martingales $\{R(0,\nu,t),t\geq 0\}$ that are indexed by $\nu\geq 0$; hence, $\{R(0,\nu',t)-R(0,\nu,t),t\geq 0\}$ is also a martingale for every pair $\nu,\nu'\geq 0$. It can also be verified that

$$\left\{\frac{\partial}{\partial\nu}R(0,\nu,t),t\geq 0\right\}$$

is a martingale for any $\nu>0$. However,

$$\left\{\lim_{\nu\to 0+}\frac{\partial}{\partial\nu}R(0,\nu,t),t\geq 0\right\}$$

is not a martingale. The latter is simply

$$\lim_{\nu\to 0+}\frac{\partial}{\partial\nu}R(0,\nu,t) = \frac{1}{2}\log\left(\frac{X_t}{x}\right).$$

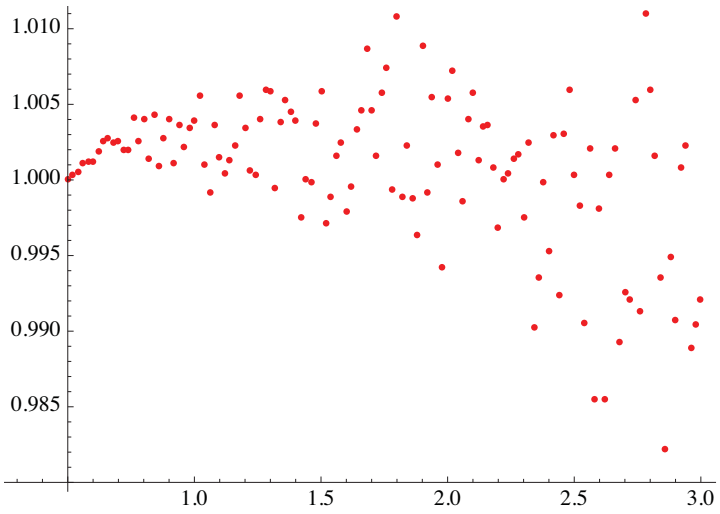

Figure 2: Simulation of the Radon Nikodym derivative for $\mu = 0.5$ as a function of ν. The number of steps used is $n = 32$.

This process is a *local* martingale, because

$$d\log X_t \;=\; \frac{2}{\sqrt{X_t}}\,dW_t.$$

This is one of the rare instances where one meets a local martingale that is not a martingale in an applied problem.

The integral J_t is computed as

$$\hat{J}_t(n) = \frac{1}{n}\sum_{j=1}^{\lfloor nt \rfloor} \frac{1}{X_{\frac{j}{n}}}.$$

Moment matching cannot be used to improve this approximation of the integral, since its expectation under $Q_x^{(0)}$ is infinite. If we use a larger value of n, *i.e.* a smaller time step size, then precision improves. For instance, we ran the simulations with 512 steps per year, and found a bump with a maximum value of 1.026, reached around $\nu = 0.15$. Nevertheless, the estimated derivative at $\nu = 0$, with a step size of .001, was 0.517 (recall the theoretical value 0.522 above).

A possible way out we explored is to use a three-dimensional squared Bessel process, rather than a two-dimensional one. Figure 2 shows the results using $\mu = 0.5$, which corresponds to $\delta = 3$, with $n = 32$ steps. The other parameters of the model are the same as in Figure 1. A bump is also present, with a positive slope near $\nu = 0$, but it is less pronounced than for $\mu = 0$. Once again the bump shrinks and moves closer to $\nu = 0$ when the number of steps is increased. Figure 3 shows what happens if a four-dimensional BESQ is simulated. There is apparently no bump when $\delta = 4$.

We have priced derivatives with $\mu = 0.5$ and $\mu = 1.0$. An obvious consequence is the larger number of random numbers required. We have also observed numerical instabilities when $\nu < \mu$, that we are currently investigating.

5 CONCLUSION

We have formulated a change of measure method for pricing derivatives in models that involve the square-root process. A limitation of the method is that it does not apply if the dimension of the square-root process

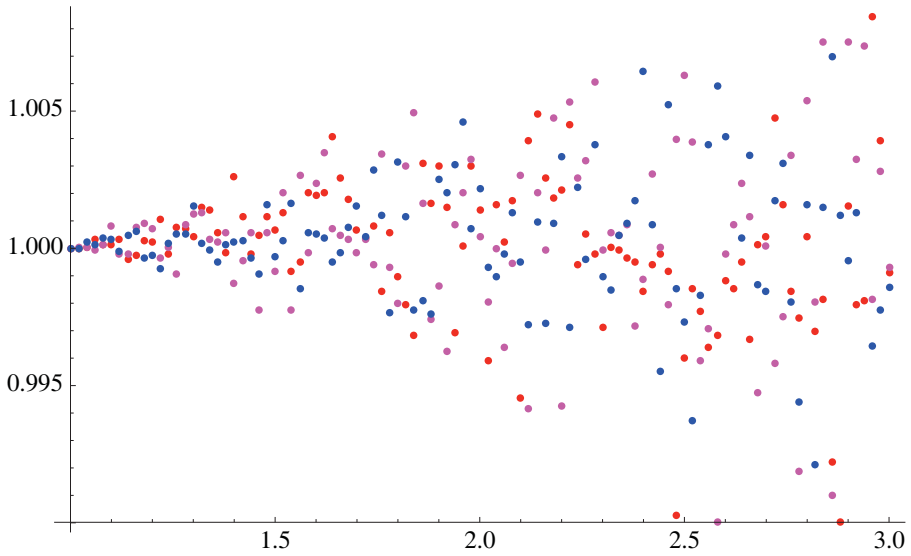

Figure 3: Simulation of the Radon Nikodym derivative with $\mu = 1$ as a function of ν. The number of steps used are 32 (red), 64 (magenta) and 128 (blue).

is less than 2. The method is simple to code and produces very good results in a majority of cases. However, there is a difficulty in that results may be inaccurate for some combinations of parameters, if the index (ν) of the square-root process is close to 0, which is the same as the dimension (δ) of the process being close to 2 (how close depends on $t$and v_0). There may be what we called a "bump" in computed Radon-Nikokym derivatives. It is always possible to get rid of the bumb by increasing the number of time steps, but of course this means longer execution times. Our understanding of this phenomenon is still incomplete. The method needs refining, and we are working on a correction to the integral of 1 over X ("J_t" in this paper), which would produce accurate results in all cases. In summary, the method has the advantage of simulating the squared Bessel process without any discretization error, but the Radon-Nikodym derivative involves integrals of the process, and those do give rise to discretization error.

In our numerical work we have used some of the formulas in Dufresne, Garrido, and Morales (2009), that allow for the very fast computation of European put and call prices. For those options there is no need to resort to the transition density function, as done in del Baño Rollin, Ferreiro-Castilla, and Utzet (2010) and Glasserman and Kim (2011), since Fourier inversion works and is much more efficient numerically.

We have also confirmed the value of the quadratic-exponential scheme, by showing that it produces results in agreement with "exact" option prices obtained by inverse Mellin transform.

ACKNOWLEDGMENTS

This research was partially supported by the CUNY Institute for Computer Simulation, Stochastic Modeling and Optimization, and also by Linkage Grant LP0882686 from the Australian Research Council and BHP Billiton.

REFERENCES

Ait-Sahali, Y., and R. Kimmel. 2007. "Maximum likelihood estimation of stochastic volatility models". *Journal of Financial Econometrics* 83:413–452.

Andersen, L. 2008. "Simple and efficient simulation of the Heston stochastic volatility model". *Journal of Computational Finance* 11 (3): 1–42.

Bakshi, G., C. Cao, and Z. Chen. 1997. "Empirical performance of alternative option pricing models". *Journal of Finance* 52 (5): 2003–2049.

Bollerslev, T., and H. Zhou. 2002. "Estimating stochastic volatility diffusion using conditional moments of integrated volatility". *Journal of Econometrics* 109 (1): 33–65.

Cox, J., E. Ingersoll, and S. Ross. 1985. "A theory of the term structure of interest rates". *Econometrica* 53 (2): 385–407.

del Baño Rollin, S., Ferreiro-Castilla, A., and F. Utzet, 2010. "On the density of log-spot in the Heston volatility model". *Stochastic Processes and their Applications* 120 (2): 2037-2063.

Dragulescu, A. A., and V. M. Yakovenko. 2002. "Probability distribution of returns in the Heston model with stochastic volatility". *Quantitative Finance* 2:443–453.

Duffie, D., J. Pan, and K. Singleton. 2000. "Transform analysis and asset pricing for affine jump-diffusion". *Econometrica* 68 (6): 1343–1376.

Dufresne, D. 2001. "The integrated square-root process". Working Paper No.90, Centre for Actuarial Studies, University of Melbourne.

Dufresne, D., J. Garrido, and M. Morales. 2009. "Fourier inversion formulas in option pricing and insurance". *Methodology and Computing in Applied Probability* 11:359–383.

Eraker, B. 2004. "Do stock prices and volatility jump? Reconciling evidence from spot and option prices". *The Journal of Finance* 59 (3): 1367–1403.

Feller, W. 1951. "Two singular diffusion problems". *Annals of Mathematics* 54 (1): 173–182.

Forbes, C. S., G. M. Martin, and J. Wright. 2007. "Inference for a class of stochastic volatility models using option and sport prices: application of bivariate Kalman filter". *Econometric Reviews* 26 (2-4): 387–418.

Fouque, J.P. and M.J. Lorig 2011. "A fast mean-reverting correction to Heston stochastic volatility model". *SIAM J. Finan. Math.* 2(1): 221-254.

Glasserman, P. and K.Kim. 2011. "Gamma expansion of the Heston stochastic volatility model". *Finance and Stochastics* 15 (2): 267-296.

Göing-Jaeschke, A., and M. Yor. 2003. "A survey and some generalizations of Bessel processes". *Bernoulli* 9 (2): 313–349.

Heston, S. L. 1993. "A closed-form solution for options with stochastic volatility with applications to bond and currency options". *The Review of Financial Studies* 6 (2): 327–343.

Pan, J. 2002. "The jump-risk premia implicit in options: Evidence from an integrated time- series study.". *Journal of Financial Econometrics* 63 (1): 3–50.

Revuz, D., and M. Yor. 1999. *Continuous Martingales and Brownian Motion*. Third ed. Springer.

DANIEL DUFRESNE is Professor at the Centre for Actuarial Studies, Department of Economics at the University of Melbourne. His email address is dufresne@unimelb.edu.au.

FELISA VÁZQUEZ-ABAD is Professor in the Computer Science Department at Hunter College, City University of New York. She received her PhD in Applied Mathematics from Brown University. Her email address is felisav@hunter.cuny.edu.

STEPHEN CHIN is a quant trader at Ingensoma Arbitrage. He holds a PhD degree in financial mathematics from the University of Melbourne. His email address is stephen.chin@ingensoma.com.