

On the Efficacy of Fourier Series Approximations for Pricing European Options

A. S. Hurn¹, K. A. Lindsay¹, A. J. McClelland²

¹School of Economics and Finance, Queensland University of Technology, Brisbane, Australia

²Numerix Pty Ltd, Sydney, Australia

Email: s.hurn@qut.edu.au, kenneth.lindsay@qut.edu.au, clelland@numerix.com

Received 23 July 2014; revised 26 August 2014; accepted 5 September 2014

Copyright © 2014 by authors and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This paper investigates several competing procedures for computing the prices of vanilla European options, such as puts, calls and binaries, in which the underlying model has a characteristic function that is known in semi-closed form. The algorithms investigated here are the half-range Fourier cosine series, the half-range Fourier sine series and the full-range Fourier series. Their performance is assessed in simulation experiments in which an analytical solution is available and also for a simple affine model of stochastic volatility in which there is no closed-form solution. The results suggest that the half-range sine series approximation is the least effective of the three proposed algorithms. It is rather more difficult to distinguish between the performance of the half-range cosine series and the full-range Fourier series. However there are two clear differences. First, when the interval over which the density is approximated is relatively large, the full-range Fourier series is at least as good as the half-range Fourier cosine series, and outperforms the latter in pricing out-of-the-money call options, in particular with maturities of three months or less. Second, the computational time required by the half-range Fourier cosine series is uniformly longer than that required by the full-range Fourier series for an interval of fixed length. Taken together, these two conclusions make a case for pricing options using a full-range range Fourier series as opposed to a half-range Fourier cosine series if a large number of options are to be priced in as short a time as possible.

Keywords

Fourier Transform, Fourier Series, Characteristic Function, Option Price

1. Introduction

The need to price vanilla European options in a rapid manner arises in numerous activities at financial institutions. Perhaps the most well-known situation in which this need occurs in practice is model calibration, in which exotic options are priced using models with values for the (risk-neutral) parameters chosen in such a way as to ensure that the model reproduces quoted prices for liquid options. For each set of parameters considered in the search space, it is therefore necessary to evaluate the prices of all options in the calibration set using the model, before comparing these with the quoted prices. In a similar vein, there is a growing literature which uses large panels of options data for estimating model parameters ([1]-[5]) in which the computational complexity is driven by the evaluation of vast numbers of option prices. Consequently, this paper investigates several competing procedures for computing the prices of vanilla European options, such as puts, calls and binaries and assesses their comparative performance on the basis of accuracy and computational speed.

Various strategies have been proposed for calculating the price of option contracts from knowledge of the conditional characteristic function of the underlying model. It is an important fact that a surprisingly large number of models have a semi-closed expression for their conditional characteristic function. For example, the identification of the conditional characteristic function for multivariate affine models with/without jump processes leads to the solution of a family of ordinary differential equations, albeit in the complex plane. In view of the Levy-Khintchine theorem, the identification of the conditional characteristic function for Levy processes is expressed in terms of various integrals with respect to the Levy measure.

The most commonly used techniques for taking advantage of a known conditional characteristic function have at their core the application of the Fast Fourier Transform (FFT). The best documented approaches is due to Carr and Madan [6] who construct an expression for the price of a European call option in terms of an integral over the characteristic function. This integral, which has an oscillatory kernel, is computed by an application of the FFT. Borak, Detlefsen and Hardle [7] apply the FFT strategy and demonstrate its efficacy by comparison with Monte Carlo simulation for a variety of models. Lord, Fang, Bervoets and Oosterlee [8] and Kwok, Leung and Wong [9] demonstrate how Fourier's convolution theorem in combination with the FFT can be used to price certain exotic options from knowledge of the conditional characteristic function of the price of the underlying asset. A different approach pioneered by Fang and Oosterlee [10] uses the characteristic function to directly approximate the marginal transitional probability density of returns by a Fourier cosine series. Recently Zhang, Grzelak and Oosterlee [11] have demonstrated how this methodology can be extended to the pricing of early-exercise commodity options under the Ornstein-Uhlenbeck process.

Rather than describing in detail the nuances of these various strategies, it is useful to point out what overarching assumptions connect them. Recall that the FFT is simply a clever piece of linear algebra that reduces the arithmetical load in implementing the Discrete Fourier Transform (DFT), namely the pair of equations connecting the coefficients of a finite Fourier series with values of the underlying function and vice versa. Therefore the decision to use the FFT implicitly makes the assumption that the underlying function is periodic over an interval of finite length, in practice determined by the range of frequencies submitted to the characteristic function, and that the function has been approximated over the interval by a finite Fourier series. The values of Fourier coefficients calculated from the characteristic function are in error by the extent to which the Finite Fourier Transform¹ differs from the Fourier transform.

Thus techniques using the FFT and those based on the construction of Fourier series share the same common assumptions and deficiencies. However, an important difference between an implementation using the FFT approach and one using the Fourier series approach is that the latter is parsimonious in its use of arithmetic whereas the former typically performs more arithmetic than necessary, albeit in an efficient way. For example, if the FFT is used to determine the value of a probability density function, what is recovered is the value of the function at each node of the interval, whereas all that might be needed is the value of the probability density function over a sub-interval.

The focus of this work is on the algorithm proposed by Fang and Oosterlee [10] who give a convincing demonstration of the efficacy of the Fourier cosine series. This series is more accurately called the half-range² cosine series because the actual function to be expanded is defined only on half the interval of periodicity (or range), the function being extended to the full range as an even-valued function. Half-range cosine series usually fail to

¹The Finite Fourier Transform is the integral expression defining the coefficients of a Fourier series.

²Historically, half-range Fourier series have largely arisen as analytical tools for handling various types of boundary conditions when solving partial differential equations using integral transforms.

represent derivatives whereas half-range Fourier sine series usually fail to represent function values. While the use of the half-range Fourier cosine series is a solid idea, Fang and Oosterlee [10] provide no motivation or explanation as to why this choice of approximating transitional density should be preferred over the half-range Fourier sine series or the full-range Fourier series for that matter. For example, intuition would suggest that the latter might perform better simply because it uses higher frequencies, which in turn translate to a more rapidly converging Fourier expansion. Indeed this intuition is borne out in calculation, but of course speed is not the only criterion of relevance in assessing the efficacy of a numerical procedure.

An important but subtle difference shared by the half-range Fourier cosine and full-range Fourier series approximations of density, but different from representations of density based on the half-range Fourier sine series, is that the former assign unit probability to the interval of support when in reality probability lies outside this interval, whereas the latter imposes zero probability density at the endpoints of the interval of support in contravention of reality, but on the other hand does not assign unit probability to the interval of support. Is one approach always superior to the other or is it a case of horses for courses? Intuition might suggest the latter. For example, when pricing a call option the most important component of the pricing error comes from the exclusion of contributions from asset price exterior to the finite interval of support. Because the half-range cosine and full-range Fourier series necessarily capture unit density, intuition might suggest that these approximations provide potential compensation for this component of pricing error. On the other hand intuition would suggest that the same approximations, when used to price binary options, might have a tendency to exaggerate the probability of exercise and therefore overprice this option in contrast to the half-range sine series approximation of probability density.

2. Fourier Series and Transform

Suppose that $f(y)$ satisfies the Dirichlet conditions on $[a, b]$, then there are three common ways in which $f(y)$ may be represented by a Fourier series. These are the half-range Fourier cosine series, the half-range Fourier sine series and the full-range Fourier series with respective representations

$$\begin{aligned}
 \text{(a)} \quad & \frac{a_0}{2} + \sum_{k=1}^{\infty} a_k \cos\left(\frac{k\pi(y-a)}{b-a}\right), \\
 \text{(b)} \quad & \sum_{k=1}^{\infty} b_k \sin\left(\frac{k\pi(y-a)}{b-a}\right), \\
 \text{(c)} \quad & \frac{a_0}{2} + \sum_{k=1}^{\infty} a_k \cos\left(\frac{2k\pi(y-a)}{b-a}\right) + b_k \sin\left(\frac{2k\pi(y-a)}{b-a}\right).
 \end{aligned} \tag{1}$$

The use of the term “half-range” in describing Expressions (a) and (b) simply refers to the fact that the function $f(y)$, although defined in $[a, b]$, has for the construction of the Fourier series been extended into the interval $[2a-b, a]$ as an even-valued function in the case of the half-range cosine series (so that sine contributions vanish) and as an odd-valued function in the case of the half-range sine series (so that the constant and cosine contributions vanish). Thus both half-range series are conventional Fourier series taken over the interval $[2a-b, b]$ such that the function represented by the half-range Fourier cosine series is usually not differentiable at $y = a$, whereas that represented by the half-range Fourier sine series is usually discontinuous at $y = a$.

In the case of the half-range cosine and sine series in Expressions (1a) and (1b) respectively, the coefficients a_k ($k \geq 0$) and b_k ($k \geq 1$) are calculated from the function $f(y)$ via the formulae

$$\begin{aligned}
 a_k &= \phi \int_a^b f(y) \cos\left(\frac{k\pi(y-a)}{b-a}\right) dy \\
 b_k &= \phi \int_a^b f(y) \sin\left(\frac{k\pi(y-a)}{b-a}\right) dy,
 \end{aligned} \tag{2}$$

in which

$$\phi = \frac{2}{b-a}.$$

In terms of the exponential function the real coefficients a_k and b_k are calculated respectively as the real and imaginary parts of the single complex Equation

$$a_k + ib_k = \phi \int_a^b f(y) \exp\left(\frac{k\pi i(y-a)}{b-a}\right) dy. \tag{3}$$

In the case of the full-range Fourier series in Expression (1c) the real coefficients a_k ($k \geq 0$) and b_k ($k \geq 1$) are calculated respectively as the real and imaginary parts of the complex Equation

$$a_k + ib_k = \phi \int_a^b f(y) \exp\left(\frac{2k\pi i(y-a)}{b-a}\right) dy. \tag{4}$$

Suppose now that $f(y)$ is a transitional probability density function with known characteristic function defined formally by the Equation

$$\chi(\omega) = \mathbb{E}[e^{i\omega y}] = \int_{\mathbb{R}} f(y) e^{i\omega y} dy, \tag{5}$$

where $\omega \in \mathbb{R}$ and $\chi(0) = 1$. A necessary condition for $f(y)$ to be a probability density function is that $f(y) \rightarrow 0$ as $|y| \rightarrow \infty$, and therefore there is guaranteed to be an interval $[a, b]$ such that for all $y \in (-\infty, a] \cup [b, \infty)$ it can be asserted that $f(y) < \varepsilon$ for any arbitrary small positive ε . The implication of this observation is that the Fourier coefficients in Equations (3) and (4) can be approximated from knowledge of the characteristic function via the respective formulae $a_k \approx A_k$ and $b_k \approx B_k$, where

$$\begin{aligned} A_k &= \phi \Re \left[\chi \left(\frac{k\pi}{b-a} \right) \exp \left(\frac{-k\pi i a}{b-a} \right) \right], \\ B_k &= \phi \Im \left[\chi \left(\frac{k\pi}{b-a} \right) \exp \left(\frac{-k\pi i a}{b-a} \right) \right], \end{aligned} \tag{6}$$

while the coefficients of the full-range Fourier series can be approximated from knowledge of the characteristic function via the formula $a_k \approx A_k$ and $b_k \approx B_k$, where

$$A_k + iB_k = \phi \chi \left(\frac{2k\pi}{b-a} \right) \exp \left(\frac{-2k\pi i a}{b-a} \right). \tag{7}$$

The accuracy of approximations (6) and (7) is investigated in Section 6, where it is demonstrated that the error can be made arbitrarily small by choosing a suitably large interval.

3. Approximating Probability Density Functions

The quality of this practical idea is now explored for three trial probability density functions with known closed-form expressions for their characteristic functions. The first choice is the Gaussian density which may be regarded as representative of distributions with super-exponentially decaying tail density. The second and third choices are the Gamma density and the Cauchy density which are treated as representative examples of distributions with exponentially decaying and algebraically decaying tail densities respectively.

3.1. Gaussian Density

It is well known that the Gaussian density with mean value μ and variance σ^2 has characteristic function $\chi(\omega) = \exp(i\mu\omega - \sigma^2\omega^2/2)$. In order to demonstrate the quality with which the true probability density $f(x)$ can be reconstructed from a truncated Fourier series of the form of Equation (1c), suppose that $\mu = \sigma = 1$ and take the interval of support to be $[a, b] = [-3, 5]$, i.e. four standard deviations on either side of the mean value. The approximating function in this case using $N/2$ frequencies is

$$\hat{f}(x) = \frac{1}{8} \left[c_0 + 2 \sum_{n=1}^{N/2} (c_n \cos k_n x + s_n \sin k_n x) \right], \tag{8}$$

where $k_n = n\pi/4$. **Figure 1** illustrates the quality of this approximation using 40 frequencies (solid line, $N = 80$) and using 4 frequencies (dashed line, $N = 8$).

With as few as 4 frequencies it is clear that the approximating density still provides a good representation of the true density; with 40 frequencies the approximating density function is indistinguishable from the true density function, at least in terms of the resolution in **Figure 1**. It will be seen that this excellent performance is explained by the fact that the cumulative distribution function of the Gaussian probability density function converges to zero super-exponentially as $x \rightarrow -\infty$ and to unity as $x \rightarrow \infty$.

3.2. Gamma Density

The Gamma density with shape parameter α and scale parameter β has probability density function and characteristic function give by the respective formulae

$$f(x) = \frac{1}{\Gamma(\alpha)\beta} \left(\frac{x}{\beta}\right)^{\alpha-1} e^{-x/\beta},$$

$$\chi(\omega) = \frac{1}{(1-i\beta\omega)^\alpha}.$$
(9)

The approximating density is identical to Expression (8) with $b-a=8$. The task is now to reconstruct the Gamma probability density function from $\chi(\omega)$. Because $Y = X/\beta$ is also Gamma distributed, in this case with shape parameter α and unit scale parameter, then the appropriate comparison for the Gamma distribution is between the approximating Fourier series and a Gamma distribution with unit scale parameter. **Figure 2** illustrates the quality of this approximation for $\Gamma\left(\frac{3}{2}, 1\right)$ for 50 frequencies (solid line, $N = 100$) and 20 frequencies (dashed line, $N = 40$).

With the exception of the region very close to the origin, the true density and approximated density are not significantly different for even 20 frequencies, and with 50 frequencies the difference between the true density and the approximating density is not discernible with the exception of the origin which does not present a difficulty since the density is known to be zero there.

The quality of the approximation is again due to the fact that the cumulative distribution function of the Gamma density converges exponentially to zero as $x \rightarrow \infty$. Typical values for the coefficient of mean reversion (say, $\kappa = 3.0$), the mean volatility (say $\gamma = 0.02$) and the volatility of volatility (say $\sigma = 0.2$) in Heston's model of stochastic volatility lead to a stationary distribution of volatility described by a Gamma density with shape parameter $\alpha = 2\kappa\gamma/\sigma^2 = 3.0$. A second example with $\alpha = 3$ and $\beta = 1$ is illustrated in **Figure 3**.

The important observation from both of these experiments is that distributions with exponentially decaying tail density can also be well described by a relatively small range of frequencies.




Figure 1. Comparison of the true Gaussian density and its approximation based on 40 frequencies (solid line, $N = 80$) and 4 frequencies (dashed line, $N = 8$).




Figure 2. Comparison of the Gamma density $\Gamma(3/2,1)$ and its approximations based on 50 frequencies (solid line, $N = 100$) and 20 frequencies (dashed line, $N = 40$).




Figure 3. Comparison of the Gamma density $\Gamma(3,1)$ (solid line) and its approximations based on 10 frequencies (dashed line, $N = 20$).

3.3. Cauchy Density

The Cauchy density with median μ and scale parameter α has probability density function and characteristic function give by the respective formulae

$$f(x) = \frac{\alpha}{\pi(x-\mu)^2 + \alpha^2}, \tag{10}$$

$$\chi(\omega) = e^{i\mu\omega - \alpha|\omega|}.$$

The approximating density is again Expression (8) with $b-a=20$. The task is now to reconstruct the Cauchy probability density function from $\chi(\omega)$, and for this purpose let $\mu=2$, $\alpha=1$ and suppose that $f(x)$ is treated as a function of compact support over the interval $[a,b]=[-8,12]$. The cumulative distribution function of the Cauchy density is

$$F(x) = \frac{1}{2} + \frac{1}{\pi} \tan^{-1}\left(\frac{x-\mu}{\alpha}\right),$$

which in this case converges algebraically to zero as $x \rightarrow -\infty$ and algebraically to unity as $x \rightarrow \infty$. **Figure 4** illustrates the impact of this inferior rate of decay for the approximation based on 10 frequencies.




Figure 4. Comparison of the Cauchy density with parameters $\mu = 2$ and $\alpha = 1$ (solid line) and the approximating Fourier series based on 10 frequencies (dashed line, $N = 20$).

While some erratic behaviour is evident in the tails of the approximating density, nevertheless the quality of the approximation is remarkably good considering the small number of frequencies in use.

In general, the approximate representations of the Gaussian, Gamma and Cauchy probability density functions all share the common fact that $F(a) + 1 - F(b)$ provides an upper bound for the error in the representation of density, where $F(x)$ is the conventional cumulative distribution function of the underlying probability density function $f(x)$. In reality the error is significantly less than this upper bound due to cancellations resulting from the fact that the integrand in each of the Gaussian, Gamma and Cauchy densities is the product of a slowly varying probability density function decaying to zero and an independent rapidly oscillating function (in this case trigonometric functions represented in exponential form). However, the presence of this cancellation cannot change the character of the error which is entirely determined by the properties of $F(x)$ as $x \rightarrow -\infty$ and as $x \rightarrow \infty$.

3.4. Comparison of True and Approximating Densities

While the illustrations of **Figures 1-4** suggest that the Fourier series approximation of density is effective, it would be useful to quantify just how well density is approximated by the various Fourier methods. Instinctively it would seem that one possible way to achieve this objective is to use the Kullback-Leibler (KL) divergence criterion

$$D = \int_{\mathbb{R}} p(x) \log \frac{p(x)}{q(x)} dx \tag{11}$$

to measure the “distance” or departure of the probability density function $q(x)$ from that of $p(x)$, which in the formulation of Equation (11) is taken to be the true probability density function. The measure appeals to the fact that $D \geq 0$ with equality if and only if the density functions $p(x)$ and $q(x)$ are identical. However, the KL criterion is simply an application of Jensen’s inequality for the concave function $\log x$, and the properties of the criterion rely critically on the fact that $p(x)$ and $q(x)$ are probability densities, *i.e.* they are non-negative functions which integrate to unity.

The central idea of each Fourier approximation is to replace the true probability density function by a function of compact support, that is, $q(x) \equiv 0$ outside its interval of support. Consequently using measure (11) to compute the departure of the approximating density $q(x)$ from the true density $p(x)$ will typically require division by zero thereby rendering inappropriate the application of the KL criterion when used with $p(x)$ as the true probability density. However, inverting the problem, one can use the KL criterion to measure the distance between the approximating density $q(x)$ and the true density $p(x)$, that is, one may compute the value of D in Expression (11) taking $p(x)$ to be the approximating density and $q(x)$ to be the true density. In this case the integrand of expression (11) is well defined for all values of x , and therefore in order to proceed in this direction it remains to ensure that the approximating density integrates to unity. Approximations of density

based on either the half range Fourier cosine series or the full range Fourier series are guaranteed to have this property, whereas an approximation based on the half range Fourier sine series need not share this property and therefore is inappropriate for analysis using the KL criterion. **Tables 1-3** show values for the KL criterion for the Gaussian distribution with mean value one and unit standard deviation, the Gamma distribution with shape parameter 3 and scale parameter one, and the Cauchy distribution with median two and scale parameter one.

As might be anticipated, the Gaussian distribution is most efficiently approximated by Fourier methods followed by the Gamma distribution and finally the Cauchy distribution. However, it is clear that all of these distributions are well approximated by the half range Fourier cosine and full range Fourier series for sufficiently large intervals of support and an adequate number of frequencies. The results in these tables also reinforce the

Table 1. Values of the KL measure (D_c for the half range Fourier cosine series and D_f for the full Fourier series) are given for the Gaussian density with unit mean and unit standard deviation. Values measure the deviation of the true density deviates from the approximating density for intervals of length 6, 8, 10, 12 and 14 standard deviations centred about the mean of the Gaussian distribution using 5, 10, 25, 50, 100 and 200 frequencies.

Interval	KL	Number of frequencies used in Fourier sums					
Length	Measure	5	10	25	50	100	200
6	D_c	0.0038	0.0031	0.0031	0.0031	0.0031	0.0031
	D_f	0.0031	0.0031	0.0031	0.0031	0.0031	0.0031
8	D_c	0.0498	0.0001	0.0001	0.0001	0.0001	0.0001
	D_f	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
10	D_c	0.4284	0.0005	0.0000	0.0000	0.0000	0.0000
	D_f	0.0005	0.0000	0.0000	0.0000	0.0000	0.0000
12	D_c	1.2491	0.0181	0.0000	0.0000	0.0000	0.0000
	D_f	0.0181	0.0000	0.0000	0.0000	0.0000	0.0000
14	D_c	2.4615	0.1301	0.0000	0.0000	0.0000	0.0000
	D_f	0.1301	0.0000	0.0000	0.0000	0.0000	0.0000

Table 2. Values of the KL measure (D_c for the half range Fourier cosine series and D_f for the full Fourier series) are given for the Gamma density with shape parameter 3 and unit scale factor. Values measure the deviation of the true density deviates from the approximating density for intervals $[0, L]$ of length 6, 8, 10, 12 and 14 scale factors using 5, 10, 25, 50, 100 and 200 frequencies.

Interval	KL	Number of frequencies used in Fourier sums					
Length	Measure	5	10	25	50	100	200
6	D_c	0.0793	0.0720	0.0715	0.0715	0.0715	0.0715
	D_f	0.0890	0.0874	0.0871	0.0871	0.0871	0.0871
8	D_c	0.0396	0.0174	0.0158	0.0158	0.0158	0.0158
	D_f	0.0196	0.0173	0.0168	0.0168	0.0167	0.0167
10	D_c	0.0669	0.0076	0.0033	0.0032	0.0032	0.0032
	D_f	0.0105	0.0039	0.0031	0.0031	0.0031	0.0031
12	D_c	0.1532	0.0142	0.0008	0.0006	0.0006	0.0006
	D_f	0.0290	0.0038	0.0007	0.0006	0.0005	0.0005
14	D_c	0.2944	0.0391	0.0006	0.0001	0.0001	0.0001
	D_f	0.0745	0.0111	0.0006	0.0001	0.0001	0.0001

Table 3. Values of the KL measure (D_c for the half range Fourier cosine series and D_f for the full Fourier series) are given for the Cauchy density with median two and unit scale factor. Values measure the deviation of the true density deviates from the approximating density for intervals of length 14, 20, 26, 32 and 38 scale factors centred about the median of the Cauchy distribution using 5, 10, 25, 50, 100 and 200 frequencies.

Interval		Number of frequencies used in Fourier sums					
Length	KL Measure	5	10	25	50	100	200
14	D_c	0.2425	0.1148	0.1076	0.1075	0.1075	0.1075
	D_f	0.1148	0.1076	0.1075	0.1075	0.1075	0.1075
20	D_c	0.4117	0.1296	0.0760	0.0753	0.0753	0.0753
	D_f	0.1296	0.0775	0.0753	0.0753	0.0753	0.0753
26	D_c	0.6356	0.2188	0.0630	0.0580	0.0580	0.0580
	D_f	0.2188	0.0713	0.0580	0.0580	0.0580	0.0580
32	D_c	0.8602	0.3391	0.0670	0.0472	0.0471	0.0471
	D_f	0.3391	0.0932	0.0472	0.0471	0.0471	0.0471
38	D_c	1.0715	0.4885	0.0937	0.0403	0.0397	0.0397
	D_f	0.4885	0.1423	0.0403	0.0397	0.0397	0.0397

idea that the number of frequencies used in the approximation is of secondary importance to the size of the interval of support once sufficient frequencies are in use. This observation accords with intuition in the respect that larger intervals of support capture more of the true density and reduce the distortion in the approximating density, which as has already been commented, will always integrate to one. It would seem that approximations based on 50 frequencies are adequate in all of these examples. The results suggest that using more frequencies provides no meaningful improvement in accuracy. Moreover, for practical purposes there is little to choose between approximation based on the half range Fourier cosine and the full range Fourier series, although the latter has a slight edge for sufficiently large intervals and an adequate number of frequencies.

4. Pricing European Options

The successfully pricing of European option contracts for affine models of stochastic volatility requires knowledge of the marginal density of the asset price under the risk-neutral measure. The difficulty, however, is that no closed-form expression for this density is available for even the simplest of the multivariate affine models used in finance, although it is well known that such models have characteristic function, $\chi(t, \omega)$, of generic form

$$\exp \left[\beta_0(\tau, \omega) + \beta_1(\tau, \omega)Y + \sum_{k=2}^M \beta_k(\tau, \omega)X_k \right], \tag{12}$$

where T is the maturity of the option, $\tau = T - t$ is the backward variable and ω is the characteristic variable associated with the non-dimensional variable Y , here defined to be the logarithm of the ratio of asset price to strike price. The state variables $\mathbf{X} = (X_2, \dots, X_M)$ in Expression (12) denote the values of the latent states of the system at time t . Typically the functions β_0, \dots, β_M are the solution of a system of $(M + 1)$ ordinary differential Equations with initial conditions $\beta_1(0, \omega) = i\omega$ and $\beta_0(0, \omega) = \beta_2(0, \omega) = \dots = \beta_M(0, \omega) = 0$. In overview, there is a well trodden procedure that starts with the specification of the multivariate affine model and ends with the construction of $\chi(t, \omega)$.

4.1. European Call Option

In the case of a European call option with strike price K and maturity T on an asset with spot price S_0 , the price of the option is

$$e^{-rT} \int_K^\infty (S - K) \tilde{f}_Q(S_0, X_2 \dots X_M, T | S; \theta) dS,$$

where $\tilde{f}_Q(S_0, X_2 \dots X_M, T | S; \theta)$ is the marginal density of the asset price at the maturity of the option when

S_0 is the spot price of the asset and (X_2, \dots, X_M) are the spot values of the latent states. When expressed in terms of $y = \log(S/K)$, the price of the option, say C , becomes

$$e^{-rT} K \int_0^\infty (e^y - 1) f_Q(\xi, X_2 \dots X_M, T | y; \theta) dy, \tag{13}$$

where $\xi = \log(S_0/K)$. The value of this integral is now approximated under the assumption that $f_Q(\xi, X_2 \dots X_M, T | y; \theta)$ is well approximated by a function of compact support over the interval $[a, b]$. Of course, the specific form taken by this approximation will depend of the choice of Expressions (a), (b) or (c) in Equations (1), but in each case the approximation of $f_Q(\xi, X_2 \dots X_M, T | y; \theta)$ by a function of compact support in $[a, b]$ necessarily changes the interval of integration in Expression (13) from $[0, \infty)$ to $[0, b]$. Thereafter it follows that the costs, C , of a call option on the basis of approximations (a) and (b) are respectively

$$\begin{aligned} \text{(a)} \quad C &= e^{-rT} K \left[\frac{a_0}{2} (e^b - b - 1) + \sum_{k=1}^\infty a_k \frac{\omega_k e^b (-1)^k - \omega_k \cos \omega_k a - \sin \omega_k a}{\omega_k (1 + \omega_k^2)} \right], \\ \text{(b)} \quad C &= e^{-rT} K \sum_{k=1}^\infty \frac{b_k}{\omega_k} \left[(e^b - 1) (-1)^{k-1} + \frac{\omega_k \sin \omega_k a - \cos \omega_k a + e^b (-1)^k}{1 + \omega_k^2} \right], \end{aligned} \tag{14}$$

where $\omega_k = k\pi/(b - a)$. Based on approximation (c) the price of a call option is

$$\begin{aligned} \text{(c)} \quad C &= e^{-rT} K \left[\frac{a_0}{2} (e^b - b - 1) + \sum_{k=1}^\infty a_k \frac{\lambda_k e^b - \lambda_k \cos \lambda_k a - \sin \lambda_k a}{\lambda_k (1 + \lambda_k^2)} \right. \\ &\quad \left. + \frac{b_k}{\lambda_k} \left(1 - e^b + \frac{\lambda_k \sin \lambda_k a - \cos \lambda_k a + e^b}{1 + \lambda_k^2} \right) \right], \end{aligned} \tag{15}$$

where $\lambda_k = 2k\pi/(b - a)$. The primary computational load in the computation of Expressions (a), (b) and (c) resides in the calculation of trigonometric functions, and consequently individual terms of Expression (c) require more arithmetical effort than the equivalent terms of either Expression (a) or (b). However the difference in computational load is minor when the more rapid convergence of summation (c) is compared with that of summations (a) and (b) simply because the frequencies used in approximation (c) are exactly double those used in approximations (a) and (b).

4.2. Binary Option

The price of a binary (call) option with strike price K and maturity T on an asset with spot price S_0 is

$$e^{-rT} \int_K^\infty \tilde{f}_Q(S_0, X_2 \dots X_M, T | S; \theta) dS,$$

where $\tilde{f}_Q(S_0, X_2 \dots X_M, T | S; \theta)$ is the marginal density of the asset price at the maturity of the option when the asset has spot price S_0 and (X_2, \dots, X_M) are the spot values of the latent states. When expressed in terms of $y = \log(S/K)$, the price of the binary becomes

$$D = e^{-rT} \int_0^\infty f_Q(\xi, X_2 \dots X_M, T | y; \theta) dy, \tag{16}$$

where $\xi = \log(S_0/K)$. The value of this integral is now approximated under the assumption that $f_Q(\xi, X_2 \dots X_M, T | y; \theta)$ is well represented by the procedures proposed in Section 2. Specifically, the price of a binary option computed from the half-range Fourier cosine and sine series are respectively

$$\begin{aligned} \text{(a)} \quad D &= e^{-rT} \left[\frac{a_0}{2} b + \sum_{k=1}^\infty a_k \frac{\sin \omega_k a}{\omega_k} \right], \\ \text{(b)} \quad D &= e^{-rT} \sum_{k=1}^\infty \frac{b_k (\cos \omega_k a - (-1)^k)}{\omega_k}, \end{aligned} \tag{17}$$

where $\omega_k = k\pi/(b-a)$. The price of the binary option based on the full-range Fourier series (c) is

$$(c) \quad D = e^{-rT} \left[a_0 b + \sum_{k=1}^{\infty} a_k \frac{\sin \lambda_k a}{\lambda_k} + \sum_{k=1}^{\infty} \frac{b_k}{\lambda_k} (\cos \lambda_k a - 1) \right], \tag{18}$$

where $\lambda_k = 2k\pi/(b-a)$. As is the case with the pricing of a call option, the frequencies used in approximation (c) are double those used in approximations (a) and (b) suggesting that calculations based on Expression (c) can be expected to converge more rapidly than those based on Expressions (a) and (b).

5. The Heston Model of Stochastic Volatility

Heston’s [12] risk-neutral model of stochastic volatility has expression

$$\begin{aligned} dY &= (r - q - V/2)dt + \sqrt{V} \left(\sqrt{1 - \rho^2} dW_1 + \rho dW_2 \right), \\ dV &= \kappa_Q (\gamma_Q - V)dt + \sigma \sqrt{V} dW_2, \end{aligned} \tag{19}$$

where $Y = \log S/K$, V is the diffusion of asset price, r is the risk-free rate of interest, q is the dividend rate, κ_Q is the risk-neutral rate of mean reversion of volatility to the risk-neutral long run value of γ_Q , σ scales the volatility of diffusion, ρ is the local correlation between returns and volatility and dW_1, dW_2 are increments in the independent Brownian motions W_1 and W_2 .

Suppose that $f(Y, V, t | y, v, T)$ is the probability density function of Equations (19) expressed in terms of the backward state (Y, V) and the forward state (y, v) , then the backward Kolmogorov Equation satisfied by the transitional probability density function of Equations (19) is

$$\frac{\partial f}{\partial t} + (r - q - V/2) \frac{\partial f}{\partial Y} + \kappa_Q (\gamma_Q - V) \frac{\partial f}{\partial V} + \frac{V}{2} \left(\frac{\partial^2 f}{\partial Y^2} + 2\rho\sigma \frac{\partial^2 f}{\partial Y \partial V} + \sigma^2 \frac{\partial^2 f}{\partial V^2} \right) = 0. \tag{20}$$

Let $\chi(Y, V, t, \omega_Y, \omega_V)$ be the characteristic function of $f(Y, V, t | y, v, T)$ with respect to the forward variables, that is,

$$\chi(Y, V, t, \omega_Y, \omega_V) = \int_{\mathbb{Q}_2} f(Y, V, t | y, v, T) e^{i(\omega_Y y + \omega_V v)} dy dv. \tag{21}$$

By taking the Fourier transform of Equation (20) with respect to the backward variables, the characteristic function $\chi(Y, V, t, \omega_Y, \omega_V)$ is seen to satisfy the partial differential Equation

$$\frac{\partial \chi}{\partial t} + (r - q - V/2) \frac{\partial \chi}{\partial Y} + \kappa_Q (\gamma_Q - V) \frac{\partial \chi}{\partial V} + \frac{V}{2} \left(\frac{\partial^2 \chi}{\partial Y^2} + 2\rho\sigma \frac{\partial^2 \chi}{\partial Y \partial V} + \sigma^2 \frac{\partial^2 \chi}{\partial V^2} \right) = 0. \tag{22}$$

with terminal condition

$$\chi(Y, V, T, \omega_Y, \omega_V) = \exp \left[i(\omega_Y Y + \omega_V V) \right].$$

Thereafter, it is straightforward to show that the ansatz

$$\chi(Y, V, t, \omega_Y, \omega_V) = e^{\beta_0(\tau) + \beta_1(\tau)Y + \beta_2(\tau)V} \tag{23}$$

with $\tau = T - t$ is a solution of Equation (22) provided the coefficient functions $\beta_0(\tau)$, $\beta_1(\tau)$ and $\beta_2(\tau)$ satisfy the ordinary differential equations

$$\begin{aligned} \frac{d\beta_0}{d\tau} &= (r - q) \beta_1 + \kappa_Q \gamma_Q \beta_2, \\ \frac{d\beta_1}{d\tau} &= 0, \\ \frac{d\beta_2}{d\tau} &= -\frac{\beta_1}{2} - \kappa_Q \beta_2 + \frac{1}{2} (\beta_1^2 + 2\rho\sigma \beta_1 \beta_2 + \sigma^2 \beta_2^2). \end{aligned} \tag{24}$$

The characteristic function of the marginal density of the terminal value of $y = \log(S_T/K)$ requires the solution of Equations (24) with initial conditions $\beta_0(0)=0$, $\beta_1(0)=i\omega_y$ and $\beta_2(0)=0$ for the particular values of ω_y needed to construct the half-range and full-range Fourier series approximations of transitional probability density.

On a practical note, the fact that the characteristic function of $\log(S/K)$ embeds the difference between the risk-free rate of interest and the dividend rate suggest at first sight that the coefficients β_0, \dots, β_M must be computed whenever the risk-free rate or dividend rate changes, potentially each day, and not just when the parameters of the model are changed. The key to avoiding this difficulty is to note that the difference $(r-q)$ enters the calculation of β_0 alone, and that the behaviours of β_1, \dots, β_M are independent of the behaviour of β_0 . Consequently, the resolution of this practical dilemma is to divide the calculation of β_0 into two stages with the first stage treating only those calculations that involve the difference $(r-q)$ and the second stage treating everything that is independent of the value of $(r-q)$. Both stages are brought together each day in the calculation of the expected payoff, but only the first stage of calculation needs be done on a day-to-day basis. In particular, this calculation is very easy as it amounts to adding $(r-q)T\omega_n$ to the second stage calculation of β_0 in the computation of $\chi(\omega_n)$.

6. Error Analysis

Let $f_Q(y)$ denote the marginal density of $Y = \log(S/K)$, where K denotes the strike price of an asset. The purpose of this section is to demonstrate that the error in estimating the price of a call option using only that part of the density $f_Q(y)$ in the finite interval $[a, b]$ can be made arbitrarily small by choosing a suitably large interval. The function $f_Q(y)$, when truncated to the interval $[a, b]$, is assumed to satisfy the Dirichlet conditions and therefore is guaranteed to have a convergent Fourier series on $[a, b]$. The approximation procedure introduces three different errors which are now described.

6.1. Truncation Error

Truncation error occurs when the semi-infinite interval in Expression (13) or (16) is replaced by an integral over a finite interval, say $[0, b]$ in these expressions. With this approximation in place, the price of the call option is approximated by the expression

$$e^{-rT} K \int_0^b (e^y - 1) f_Q(y) dy \quad (25)$$

resulting in an error due to the loss of the contribution to the price from values of y in $[b, \infty]$. A straightforward analysis indicates that the error in replacing the true cost of the call option by Expression (25) is

$$e^{-rT} K \left[\int_0^\infty (e^y - 1) f_Q(y) dy - \int_0^b (e^y - 1) f_Q(y) dy \right]$$

which in turn simplifies to give

$$= e^{-rT} K \left[(e^b - 1)(1 - F_Q(b)) + \int_b^\infty e^y (1 - F_Q(y)) dy \right], \quad (26)$$

where $F_Q(y)$ is the cumulative function of y . Evidently the truncation error in pricing a call option is the sum of two positive contributions, both of which are driven by the extent to which the restriction of the marginal density of Y to the interval $[a, b]$ fails to capture density in the upper tail of the distribution of Y . Thus Expression (25) with a numerically accurate representation of the true marginal density of Y will always underestimate the true value of a call option, although the error can be made arbitrarily small by taking $[a, b]$ to be a suitably large interval.

The inference from this observation is that approximations of $f_Q(y)$ which place unit density in $[a, b]$ may potentially compensate for some of the error incurred in using the truncation procedure to price a call option. The half-range Fourier cosine and full-range Fourier series approximations of transitional probability density both fall into this category, which in turn suggests that the use of the truncation procedure with these approximations of marginal density may well provide rather smaller pricing errors than might casually be anticipated. By contrast, the half-range Fourier sine series approximation of marginal density does not capture unit

density in the interval $[a, b]$, and the corresponding conjecture is that this approximation of marginal density will perform less well than the half-range Fourier cosine and full-range Fourier series approximations with regard to the pricing of a call option. **Table 4** indicates that this conjecture is borne out in practice.

6.2. Approximation Errors

Suppose now that the transitional density $f_Q(y)$ in the interval $[a, b]$ is replaced by the Fourier series

$$f_Q(y) = \frac{a_0}{2} + \sum_{k=1}^{\infty} a_k \cos \omega_k (y - a) + \sum_{k=1}^{\infty} b_k \sin \omega_k (y - a), \tag{27}$$

where the choice of frequencies ω_k and the values of a_k, b_k depend on the choice of Fourier representation. Two errors arise in the calculation of Expression (25). First, the plan is to replace the coefficients a_k, b_k in Equation (27) by the misspecified coefficients A_k, B_k computed directly from the characteristic function $\chi(\omega_k)$ of $f_Q(y)$ via the formula

Table 4. Percentage errors recorded by the half-range cosine series, half-range sine series and full-range Fourier series of pricing a European call option and a binary option with strike $K = 1200$ and maturities one month (panel (a)), three months (panel (b)) and six months (panel (c)). The factor refers to the multiple of $\sigma\sqrt{T}$ used to establish the interval over which to compute the numerical approximation.

(a)						
Factor	Call	Strike = 1200	T = 0.083	Binary	Strike = 1200	T = 0.083
	Cos	Sin	Full	Cos	Sin	Full
5.00	-4.160746	-17.315659	-7.719366	-0.031036	-4.731845	-1.295507
6.00	-0.000537	-0.129887	-0.052720	0.000369	-0.017647	-0.005300
7.00	0.008273	0.018872	0.001766	-0.002995	-0.000040	-0.000010
8.00	0.058873	0.019228	0.002259	0.000286	-0.000010	0.000003
9.00	0.015106	0.018981	0.002546	-0.000038	0.009813	0.000007

(b)						
Factor	Call	Strike = 1200	T = 0.083	Binary	Strike = 1200	T = 0.083
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.006890	-0.084497	-0.035497	0.000000	-0.010412	-0.003842
6.00	-0.000014	-0.000471	-0.000204	0.000000	-0.000041	-0.000015
7.00	0.000000	0.000004	-0.000001	0.000000	-0.000000	0.000001
8.00	0.000000	0.000005	-0.000001	0.000000	-0.000000	0.000000
9.00	0.000000	0.000005	-0.000001	0.000000	0.000000	0.000000

(c)						
Factor	Call	Strike = 1200	T = 0.500	Binary	Strike = 1200	T = 0.500
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.001412	-0.020338	-0.009446	0.000000	-0.002338	-0.000993
6.00	-0.000005	-0.000115	-0.000055	0.000000	-0.000010	-0.000004
7.00	0.000000	-0.000000	-0.000000	0.000000	0.000000	-0.000000
8.00	0.000000	-0.000000	-0.000000	0.000000	0.000000	-0.000000
9.00	0.000000	-0.000000	-0.000000	0.000000	0.000000	-0.000000

$$A_k + iB_k = \frac{2e^{-i\omega_k a}}{b-a} \chi(\omega_k). \tag{28}$$

Consequently a misspecification error arises in the Fourier coefficients. Second, the truncation of the Fourier series (27) to a finite number of terms, say N terms, introduces another approximation error. Therefore the price of a call option based on this strategy is

$$e^{-rT} K \left[\frac{A_0}{2} \int_0^b (e^y - 1) dy + \sum_{k=1}^N A_k \int_0^b (e^y - 1) \cos \omega_k (y - a) dy + \sum_{k=1}^N B_k \int_0^b (e^y - 1) \sin \omega_k (y - a) dy \right] \tag{29}$$

in which $f_Q(y)$ in Expression (25) has been replaced by the first $(N + 1)$ terms of the Fourier series (27) with misspecified coefficients A_k and B_k . The error introduced by this approximation is

$$e^{-rT} K \int_0^b (e^y - 1) f_Q(y) dy - C \tag{30}$$

which has explicit expression

$$e^{-rT} K \left[\frac{(a_0 - A_0)}{2} \int_0^b (e^y - 1) dy + \sum_{k=1}^N (a_k - A_k) \int_0^b (e^y - 1) \cos \omega_k (y - a) dy + \sum_{k=1}^N (b_k - B_k) \int_0^b (e^y - 1) \sin \omega_k (y - a) dy + \sum_{k=N+1}^{\infty} a_k \int_0^b (e^y - 1) \cos \omega_k (y - a) dy + b_k \int_0^b (e^y - 1) \sin \omega_k (y - a) dy \right]. \tag{31}$$

The misspecification error in the coefficients a_k and b_k due to the use of the coefficients A_k and B_k respectively is determined from the identity

$$\begin{aligned} a_k + ib_k &= \frac{2}{b-a} \int_a^b f_Q(y) e^{i\omega_k(y-a)} dy \\ &= A_k + iB_k - 2 \frac{e^{-i\omega_k a}}{b-a} \left[\int_{-\infty}^a f_Q(y) e^{i\omega_k y} dy + \int_b^{\infty} f_Q(y) e^{i\omega_k y} dy \right]. \end{aligned}$$

Standard properties of integral Calculus guarantee that

$$\left| (a_k + ib_k) - (A_k + iB_k) \right| = \frac{2}{b-a} \left| \int_{-\infty}^a f_Q(y) e^{i\omega_k y} dy + \int_b^{\infty} f_Q(y) e^{i\omega_k y} dy \right| \leq \varepsilon_1,$$

where

$$\varepsilon_1 = \frac{2}{b-a} \left[\int_{-\infty}^a f_Q(y) dy + \int_b^{\infty} f_Q(y) dy \right] = \frac{2}{b-a} [F_Q(a) + 1 - F_Q(b)]. \tag{32}$$

It therefore follows directly that

$$|a_k - A_k| \leq \varepsilon_1, \quad |b_k - B_k| \leq \varepsilon_1, \tag{33}$$

for all values of k . Thus the misspecification error in the Fourier coefficients is driven entirely by the choice of interval $[a, b]$ via the magnitude of the cumulative function $F_Q(y)$ at $y = a$ and $y = b$. In practice this error will be significantly smaller than the maximum bound given in Equation (33) as a result of arithmetical cancellation due to the oscillatory nature of the integrand, but as commented previously, the decay of this error as $[a, b] \rightarrow \mathbb{R}$ will be determined by the behaviour of $F_Q(y)$ as $y \rightarrow -\infty$ and as $y \rightarrow \infty$.

In conclusion, the total error in pricing a call option, namely

$$\left| e^{-rT} K \int_0^\infty (e^y - 1) f_Q(y) dy - C \right|, \tag{34}$$

is formed by connecting together Equation (26) for the error arising in the truncation of the density $f_Q(y)$ to the finite interval $[a, b]$, Equation (29) for the value of the option in terms of the misspecified Fourier coefficients A_k, B_k , and Equation (31) for the size of the resulting misspecification error in the option price. The result is that the error is

$$\begin{aligned} & e^{-rT} K \left[(e^b - 1)(1 - F_Q(b)) + \int_b^\infty e^y (1 - F_Q(y)) dy \right] + e^{-rT} K \left[\frac{(a_0 - A_0)}{2} \int_0^b (e^y - 1) dy \right. \\ & + \sum_{k=1}^N (a_k - A_k) \int_0^b (e^y - 1) \cos \omega_k (y - a) dy + \sum_{k=1}^N (b_k - B_k) \int_0^b (e^y - 1) \sin \omega_k (y - a) dy \\ & \left. + \sum_{k=N+1}^\infty a_k \int_0^b (e^y - 1) \cos \omega_k (y - a) dy + \sum_{k=N+1}^\infty b_k \int_0^b (e^y - 1) \sin \omega_k (y - a) dy \right]. \end{aligned} \tag{35}$$

Each integral is replaced by its value and the triangle inequality is used to deduce that

$$\begin{aligned} |\text{Error}| & \leq e^{-rT} K \left[(e^b - 1)(1 - F_Q(b)) + \int_b^\infty e^y (1 - F_Q(y)) dy + \frac{\varepsilon_1}{2} (e^b - b - 1) \right. \\ & + \varepsilon_1 \sum_{k=1}^N \left| \frac{(\sin \omega_k a + \omega_k \cos \omega_k a)}{\omega_k (1 + \omega_k^2)} - \frac{e^b \cos \omega_k (b - a)}{1 + \omega_k^2} \right| \\ & + \varepsilon_1 \sum_{k=1}^N \left| \frac{e^b \omega_k \cos \omega_k (b - a) - \sin \omega_k a}{1 + \omega_k^2} + \frac{\cos \omega_k a - \cos \omega_k (b - a)}{\omega_k (1 + \omega_k^2)} \right| \\ & + \sum_{k=N+1}^\infty |a_k| \left| \frac{\sin \omega_k a + \omega_k \cos \omega_k a}{\omega_k (1 + \omega_k^2)} - \frac{e^b \cos \omega_k (b - a)}{1 + \omega_k^2} \right| \\ & \left. + \sum_{k=N+1}^\infty |b_k| \left| \frac{e^b \omega_k \cos \omega_k (b - a) - \sin \omega_k a}{1 + \omega_k^2} + \frac{\cos \omega_k a - \cos \omega_k (b - a)}{\omega_k (1 + \omega_k^2)} \right| \right]. \end{aligned} \tag{36}$$

The contributions to the error from the first, second, third and fourth terms on the right hand side of inequality (36) are dominated by the behaviour of $F_Q(a)$ and $1 - F_Q(b)$. By choosing the interval $[a, b]$ suitably large, $F_Q(a)$ and $1 - F_Q(b)$ can be made arbitrarily small. The behaviour of the fifth and sixth terms on the right hand side of inequality (36) depends on the choice of N . Bearing in mind that $\omega_k = O(k)$, then it straightforward to show that there are positive constants C_1 and C_2 such that

$$\begin{aligned} & \left| \frac{(\sin \omega_k a + \omega_k \cos \omega_k a)}{\omega_k (1 + \omega_k^2)} - \frac{e^b \cos \omega_k (b - a)}{1 + \omega_k^2} \right| \leq \frac{e^b + 1}{\omega_k \sqrt{1 + \omega_k^2}} \leq \frac{C_1}{k^2}, \\ & \left| \frac{e^b \omega_k \cos \omega_k (b - a) - \sin \omega_k a}{1 + \omega_k^2} + \frac{\cos \omega_k a - \cos \omega_k (b - a)}{\omega_k (1 + \omega_k^2)} \right| \leq \frac{e^b \sqrt{1 + \omega_k^2} + 1}{\omega_k \sqrt{1 + \omega_k^2}} \leq \frac{C_2}{k}. \end{aligned}$$

The well known result that if $f_Q(y)$ is a continuous function³ of y , then $|a_k| < K/k$ and $|b_k| < K/k$ indicates that the fourth and fifth terms of inequality (36) are $O(k^{-3})$ and $O(k^{-2})$ respectively. Thus the fifth and sixth terms on the right hand side of inequality (36) can also be made arbitrarily small by a suitably large

³Sharper results can be obtained if $f_Q(y) \in C^p[a, b]$ for $p > 1$. However the function represented by the half-range Fourier cosine series is not generally differentiable at $x = a$, and so the convergence argument is based on the weakest condition satisfied by Fourier coefficients.

choice of N .

In conclusion, the error in pricing a call option can be made arbitrarily small by restricting the marginal probability density to a finite interval $[a, b]$ and approximating the density in that interval by a misspecified Fourier series constructed from the characteristic function at the appropriate frequencies.

7. Performance under Simulation

A series of simulation experiments was undertaken in order to examine the efficacy of the half-range Fourier cosine series, half-range Fourier sine series and full-range Fourier series in respect of how accurately these approximations price European call options and binary options. The first experiment prices options in a Black-Scholes world so that a closed-form solution may be used to assess the pricing error, whereas the second experiment prices the same options using Heston's model of stochastic volatility.

7.1. Black-Scholes Pricing

Assume that the asset price, S , follows a geometric Brownian motion

$$dS = S(rdt + \sigma dW)$$

in which dW is the increment of a Wiener process. The advantage of using this specification is that, following the seminal work of Black and Scholes [13], exact prices are known for each type of option for all combinations of spot price S_0 , strike price K , and maturity T . It follows, therefore, that the relative percentage pricing error incurred using numerical methods based on Fourier series for a variety of different strike prices and maturities can be identified exactly.

Three major experiments are performed. In each of these experiments $r = 0.05$, $\sigma = 0.15$ per year, $S_0 = 1000$ and the strikes K are taken to be 1200, 1000 and 800 respectively so as to examine the performance of the algorithms when the options are out-of-the-money, at-the-money and in-the-money. In each simulation the accuracy of the various approximations of transitional density is assessed for prescribed values of a and b in which the size of the interval $[a, b]$ is expressed as a multiple of $\sigma\sqrt{T}$ ranging from 5 to 9. The results of these exercises are reported in **Tables 4-6** with each table reporting the results for maturities of one month (panel (a)), three months (panel (b)) and six months (panel (c)).

Two very clear general conclusions emerge from these results.

1) For options that are deep in-the-money (**Table 6**) in which $S_0 = 1000$ and $K = 800$, or at-the-money (**Table 5**) in which $S_0 = 1000$ and $K = 1000$, it matters little (if at all) which of these three pricing algorithms are chosen irrespective of the maturity of the option. Indeed the relative pricing error is zero for all practical purposes.

2) For options that are deep out-of-the-money (**Table 4**) in which $S_0 = 1000$ and $K = 1200$, the choice of algorithm is more important. The results may be summarized succinctly as follows.

a) The half-range Fourier sine series does not perform as well as the other two approximations and its use is therefore not recommended. This finding accords well with our previous intuition based on a consideration of the contribution to the price of a call option lost as a result of truncating marginal density.

b) The half-range Fourier cosine series and the full-range Fourier series both perform relatively well. When the size of the interval of approximation is a relatively small multiple of $\sigma\sqrt{T}$, namely either 5 or 6, then the half-range Fourier cosine series performs better than the full-range Fourier series. As the multiple of $\sigma\sqrt{T}$ increases and the size of the interval of approximation becomes larger, the full-range Fourier series begins to dominate. When the interval of approximation has size $10\sigma\sqrt{T}$, then the full-range Fourier series is unambiguously superior to the half-range Fourier cosine series, particularly for options of short maturity.

On the basis of this analysis and on accuracy grounds, it is hard to ignore the claim that the full-range Fourier series is the algorithm of choice when using Fourier methods to price options. Moreover, the full-range Fourier series converges faster than either the half-range sine or cosine series and is therefore likely to price options more rapidly. These themes are explored in more detail in the pricing of call options under Heston's model of stochastic volatility.

7.2. Heston's Model

A total of approximately 40,000 options over ten years were generated by simulation of Heston's model. Sixteen

Table 5. Percentage errors recorded by the half-range cosine series, half-range sine series and full-range Fourier series of pricing a European call option and a binary option with strike $K = 1000$ and maturities one month (panel (a)), three months (panel (b)) and six months (panel (c)). The factor refers to the multiple of $\sigma\sqrt{T}$ used to establish the interval over which to compute the numerical approximation.

(a)						
Factor	Call	Strike = 1000	T = 0.083	Binary	Strike = 1000	T = 0.083
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000044	-0.000313	-0.001045	0.000000	-0.000159	-0.000043
6.00	-0.000000	-0.000002	-0.000005	0.000000	-0.000001	-0.000000
7.00	0.000000	0.000000	0.000000	0.000000	-0.000000	-0.000000
8.00	-0.000000	0.000000	0.000000	0.000000	-0.000000	-0.000000
9.00	-0.000000	0.000000	0.000000	0.000000	0.000000	0.000000

(b)						
Factor	Call	Strike = 1000	T = 0.250	Binary	Strike = 1000	T = 0.250
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000063	-0.001375	-0.000547	0.000000	-0.000202	-0.000074
6.00	-0.000000	-0.000007	-0.000003	0.000000	-0.000001	-0.000000
7.00	0.000000	-0.000000	-0.000000	0.000000	-0.000000	-0.000000
8.00	-0.000000	-0.000000	-0.000000	0.000000	-0.000000	-0.000000
9.00	-0.000000	-0.000000	-0.000000	0.000000	0.000000	0.000000

(c)						
Factor	Call	Strike = 1000	T = 0.500	Binary	Strike = 1000	T = 0.500
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000116	-0.002208	-0.001059	-0.000000	-0.000254	-0.000108
6.00	-0.000000	-0.000012	-0.000006	-0.000000	-0.000001	-0.000000
7.00	0.000000	-0.000000	-0.000000	-0.000000	-0.000000	-0.000000
8.00	0.000000	-0.000000	0.000000	-0.000000	-0.000000	-0.000000
9.00	0.000000	-0.000000	0.000000	-0.000000	-0.000000	-0.000000

options were generated each day spread over 4 maturities ranging from 92 to 5 days and 4 strikes two of which are initialised at 20% out of and into the money and two of which are initialised at 7% out of and into the money. The half-range Fourier cosine series and the full-range Fourier series are then used to price call options and binary options with these strikes. In this instance no exact solutions are available, and so the accuracy of each method in respect of each type of option is gauged by comparison against values calculated using a large interval. The left hand and middle columns of **Table 7** show respectively the L_2 and L_1 relative pricing errors for out-of-the-money call options calculated using the half-range Fourier cosine series and the full-range Fourier series. The right hand column of **Table 7** shows the CPU time (secs) needed to perform the calculations. The left hand column specifies the length of the interval $[a, b]$ in multiples of $\sigma_Y = \sqrt{\gamma_Q DT}$, the unconditional standard deviation of the distribution of Y about $\log(S_0/K)$ after an interval of duration DT , taken to be a trading day in the simulation study, *i.e.* $DT = 1/252$.

A clear finding from **Table 7** is that the half-range Fourier cosine series performs well in respect of the L_2 measure for shorter intervals and in the L_1 measure for almost all choices of interval length. By contrast the

Table 6. Percentage errors recorded by the half-range cosine series, half-range sine series and full-range Fourier series of pricing a European call option and a binary option with strike $K = 800$ and maturities one month (panel (a)), three months (panel (b)) and six months (panel (c)). The factor refers to the multiple of $\sigma\sqrt{T}$ used to establish the interval over which the numerical approximation is computed.

(a)						
Factor	Call	Strike = 800		Binary	Strike = 800	
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000003	-0.000184	-0.000049	-0.000043	-0.000127	-0.000068
6.00	-0.000000	-0.000001	-0.000000	-0.000000	-0.000000	-0.000000
7.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
8.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
9.00	-0.000000	0.000000	0.000000	-0.000000	-0.000010	0.000000

(b)						
Factor	Call	Strike = 800		Binary	Strike = 800	
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000011	-0.000343	-0.000133	-0.000000	-0.000111	-0.000041
6.00	-0.000000	-0.000002	-0.000001	-0.000000	-0.000000	-0.000000
7.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
8.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
9.00	-0.000000	0.000000	0.000000	-0.000000	-0.000000	0.000000

(c)						
Factor	Call	Strike = 800		Binary	Strike = 800	
	Cos	Sin	Full	Cos	Sin	Full
5.00	-0.000022	-0.000580	-0.000259	-0.000008	-0.000147	-0.000063
6.00	-0.000000	-0.000003	-0.000001	-0.000000	-0.000001	-0.000000
7.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
8.00	-0.000000	0.000000	-0.000000	-0.000000	-0.000000	0.000000
9.00	-0.000000	0.000000	0.000000	-0.000000	-0.000000	0.000000

performance of the full-range Fourier series is poor with regard to both measures for shorter intervals. However, the quality of approximation provided by the half-range cosine series is erratic as the size of the Fourier window increases whereas that provided by the full-range Fourier series improves systematically to the extent that its performance surpasses that of the half-range Fourier cosine series for intervals of length 24 standard deviations or more. Furthermore, this level of accuracy is achieved by the full-range Fourier series in approximately 25% quicker than that required by the half-range Fourier cosine series.

Fang and Oosterlee [10] suggest choosing intervals of length 20 standard deviations. In this problem, it is evident that the half-range Fourier cosine series still enjoys an advantage over the full-range Fourier series for intervals of this length. The suggestion of this investigation is that 20 standard deviations should be regarded as a minimum length of interval. Table 8 shows the L_2 and L_1 relative pricing errors calculated from the half-range Fourier cosine series and the full-range Fourier series in respect of in-the-money call options based on Heston's model.

Table 7. The L_2 and L_1 relative pricing errors calculated from the half-range Fourier cosine series and the full-range Fourier series for out-of-the-money call options based on Heston’s model. The timings refer to the times (seconds) needed to price 12 daily call options over 10 years with maturities ranging from 5 to 92 days. The factor refers to the multiple of σ_Y used to establish the interval over which the numerical approximation is computed.

Factor	Out-of-the-money call		Out-of-the-money call		Timings (sec)	
	$(L_2$ Error)	$(L_2$ Error)	$(L_1$ Error)	$(L_1$ Error)	Cosine	Full
	Cos Series	Full Series	Cos Series	Full Series		
10	0.000228	0.172011	0.000289	0.208449	0.140	0.109
12	0.000523	0.056627	0.000045	0.063078	0.156	0.109
14	0.000053	0.016726	0.000029	0.018403	0.187	0.141
16	0.000242	0.004654	0.000030	0.005235	0.203	0.172
18	0.000374	0.001251	0.000030	0.001468	0.234	0.171
20	0.000190	0.000330	0.000035	0.000417	0.265	0.187
22	0.000064	0.000086	0.000028	0.000129	0.281	0.218
24	0.000460	0.000023	0.000030	0.000052	0.296	0.218
26	0.000058	0.000010	0.000027	0.000031	0.328	0.249
28	0.000315	0.000008	0.000029	0.000025	0.343	0.265
30	0.000219	0.000007	0.000035	0.000021	0.359	0.280

Table 8. The L_2 and L_1 relative pricing errors calculated from the half-range cosine series and the full-range Fourier series for in-the-money call options based on Heston’s model. The timings refer to the times (seconds) needed to price 12 daily call options over 10 years with maturities ranging from 5 to 92 days. The factor refers to the multiple of σ_Y used to establish the interval over which the numerical approximation is computed.

Factor	In-the-money call		In-the-money call		Timings (sec)	
	$(L_2$ Error)	$(L_2$ Error)	$(L_1$ Error)	$(L_1$ Error)	Cosine	Full
	Cos Series	Full Series	Cos Series	Full Series		
10	0.000027	0.002421	0.000362	0.215297	0.140	0.109
12	0.000008	0.000845	0.000063	0.061658	0.156	0.109
14	0.000008	0.000281	0.000043	0.017301	0.187	0.141
16	0.000008	0.000090	0.000043	0.004790	0.203	0.172
18	0.000008	0.000029	0.000043	0.001326	0.234	0.171
20	0.000024	0.000011	0.000052	0.000382	0.265	0.187
22	0.000008	0.000008	0.000041	0.000130	0.281	0.218
24	0.000008	0.000008	0.000043	0.000062	0.296	0.218
26	0.000007	0.000007	0.000041	0.000044	0.328	0.249
28	0.000008	0.000007	0.000042	0.000038	0.343	0.265
30	0.000025	0.000007	0.000053	0.000034	0.359	0.280

The results reported in **Table 8** exhibit a similar pattern of behaviour to those reported in **Table 7**, namely that for intervals of short length the half-range Fourier cosine series outperforms the full-range Fourier series with respect to both the L_2 and L_1 measures of accuracy, but that this dominance vanishes when the length of the interval of approximation is 24 standard deviations or more. In particular, both approaches generate noticeably smaller relative errors for in-the-money options than for out-of-the-money options. Because in-the-money call options carry a higher price, one inference of this observed reduction in relative error is that the absolute pricing error in using each Fourier representation is insensitive to the moneyness of the option. **Table 8** and **Table 9** report the performance of the half-range Fourier cosine series and the full-range Fourier series when used to price binary options.

The results presented in **Table 9** and **Table 10** demonstrate that both the half-range Fourier cosine series and the full-range Fourier series both generate good estimates of the price of binary options with the former performing better than the latter for intervals of short length, but with this advantage disappearing when the length of the interval is increased.

Table 9. The L_2 and L_1 relative pricing errors calculated from the half-range cosine series and the full-range Fourier series for out-of-the-money binary options based on Heston's model. The timings refer to the times (seconds) needed to price 12 daily call options over 10 years with maturities ranging from 5 to 92 days. The factor refers to the multiple of σ_Y used to establish the interval over which the numerical approximation is computed.

Factor	Binary out-of-the-money		Binary out-of-the-money		Timings (sec)	
	$(L_2 \text{ Error})$	$(L_2 \text{ Error})$	$(L_1 \text{ Error})$	$(L_1 \text{ Error})$	Cosine	Full
	Cos Series	Full Series	Cos Series	Full Series		
10	0.000020	0.005846	0.000000	0.000328	0.172	0.094
12	0.000006	0.001753	0.000000	0.000081	0.203	0.125
14	0.000020	0.000504	0.000000	0.000020	0.234	0.125
16	0.000007	0.000141	0.000000	0.000005	0.265	0.141
18	0.000028	0.000047	0.000000	0.000001	0.281	0.156
20	0.000004	0.000022	0.000000	0.000000	0.312	0.172
22	0.000020	0.000020	0.000000	0.000000	0.343	0.187
26	0.000020	0.000020	0.000000	0.000000	0.390	0.203
30	0.000022	0.000020	0.000000	0.000000	0.437	0.250

Table 10. The L_2 and L_1 relative pricing errors calculated from the half-range cosine series and the full-range Fourier series for in-the-money binary options based on Heston's model. The timings refer to the times (seconds) needed to price 12 daily call options over 10 years with maturities ranging from 5 to 92 days. The factor refers to the multiple of σ_Y used to establish the interval over which the numerical approximation is computed.

Factor	Binary in-the-money		Binary in-the-money		Timings (sec)	
	$(L_2 \text{ Error})$	$(L_2 \text{ Error})$	$(L_1 \text{ Error})$	$(L_1 \text{ Error})$	Cosine	Full
	Cos Series	Full Series	Cos Series	Full Series		
10	0.000025	0.000678	0.000004	0.000186	0.172	0.109
12	0.000007	0.000193	0.000000	0.000044	0.203	0.125
14	0.000007	0.000054	0.000000	0.000011	0.234	0.125
16	0.000007	0.000016	0.000000	0.000003	0.265	0.156
18	0.000007	0.000008	0.000000	0.000001	0.281	0.156
20	0.000006	0.000007	0.000000	0.000000	0.312	0.171
22	0.000009	0.000007	0.000000	0.000000	0.343	0.187
26	0.000016	0.000007	0.000000	0.000000	0.390	0.219
30	0.000012	0.000007	0.000000	0.000000	0.437	0.234

The results reported in **Tables 7-10** are characterised by two common denominators. First, the computational time required by the half-range Fourier cosine series is uniformly longer than that required by the full-range Fourier series for an interval of fixed length. The simple explanation for this observation is that the full-range Fourier series uses larger frequencies for a given length of interval, and therefore the full-range Fourier series converges more rapidly. Second, the pricing of call options and binary options using the half-range Fourier cosine series representation of transitional density is noticeably better than the corresponding pricing using the full-range Fourier series for short intervals $[a, b]$, but this advantage vanishes in the case of the binary option when the interval becomes suitably large and is reversed in the case of a call option. The primary explanation for this behaviour lies in the realisation that tail density is characterised by long wavelengths, or equivalently by the presence of low frequency terms, whereas peaky density is characterised by short wavelengths, or equivalently high frequency terms.

For an interval of given length, the frequencies present in the half-range Fourier cosine expansion are smaller than the frequencies present in the full-range Fourier series, and therefore when the length of the interval $[a, b]$ is small, the half-range Fourier cosine series better captures tail behaviour. Increasing the length of the interval lowers the range of frequencies admitted to the full-range Fourier series expansion, which in turn improves its ability to handle tail density. However, the same increase in length will benefit the half-range Fourier cosine series representation of transitional density only provided there remains significant tail density to be characterised by these new frequencies. On the other hand, the half-range Fourier cosine series will always suffer from the drawback that the (periodic) function it represents has gradient zero at $x = a$ and $x = b$ (because the series is not generally differentiable at these points). Consequently, the half-range Fourier cosine series approximation will always misrepresent the gradient of the true density function at $x = a$ and $x = b$ in contrast to the behaviour of the approximation provided by the full-range Fourier series.

8. Conclusion

One clear conclusion from these calculations is that the half-range Fourier cosine series and the full-range Fourier series perform uniformly better than the half-range Fourier sine series. The half-range Fourier cosine series and the full-range Fourier series both perform with credit. When the length of the interval $[a, b]$ is relatively small, say ten or so standard deviations, it is clear that the half-range Fourier cosine series outperforms the full-range Fourier series over the same interval. On the other hand for intervals of larger length the full-range Fourier series is at least as good as the half-range Fourier cosine series, and outperforms the latter in pricing out-of-the-money call options, in particular, with maturities of three months or less. In general the full-range Fourier series outperforms the half-range Fourier cosine series in all circumstances provided that the interval $[a, b]$ is suitably large, although the effect is so small as not to be significant in practice. The explanation for this behaviour lies in the fact that the half-range Fourier cosine series, although not representing a differentiable function, always uses a lower spectrum of frequencies than the full-range Fourier series and therefore enjoys an initial advantage in describing tail density. As the interval length is increased, this advantage vanishes. On the other hand the larger spectrum of the full-range Fourier series guarantees more rapid convergence. Of course, timing issues may not be important if small numbers of model call prices are to be calculated, but when facing calibration problems or estimation problems as described in the introductory remarks, timing issues are indeed significant once adequate numerical accuracy is assured.

References

- [1] Johannes, M.S., Polson, N.G. and Stroud, J.R. (2009) Optimal Filtering of Jump Diffusions: Extracting Latent States from Asset Prices. *Review of Financial Studies*, **22**, 2759-2799. <http://dx.doi.org/10.1093/rfs/hhn110>
- [2] Broadie, M., Chernov, M. and Johannes, M. (2007) Model Specification and Risk Premia: Evidence from Futures Options. *Journal of Finance*, **62**, 1453-1490. <http://dx.doi.org/10.1111/j.1540-6261.2007.01241.x>
- [3] Christoffersen, P., Jacobs, K. and Mimouni, K. (2010) Volatility Dynamics for the S&P500: Evidence from Realized Volatility, Daily Returns and Option Prices. *Review of Financial Studies*, **23**, 3141-3189. <http://dx.doi.org/10.1093/rfs/hhq032>
- [4] Hurn, A.S., Lindsay, K.A. and McClelland, A.J. (2012) Estimating the Parameters of Stochastic Volatility Models Using Option Price Data. Unpublished Working Paper, NCER.
- [5] Andersen, T.G., Fusari, N. and Todorov, V. (2012) Parametric Inference and Dynamic State Recovery from Option Pa-

nels. NBER Working Paper Series.

- [6] Carr, P.P. and Madan, D.B. (1999) Option Evaluation Using the Fast Fourier Transform. *Journal of Computational Finance*, **2**, 61-73.
- [7] Borak, S., Detlefsen, K. and Hardle, W. (2005) FFT Based Option Pricing. SFB Discussion Paper 649.
- [8] Lord, R., Fang, F. Bervoets, F. and Oosterlee, C.W. (2007) A Fast and Accurate FFT-Based Methodology for Pricing Early-Exercise Options under Levy Processes. *SIAM Journal of Scientific Computing*, **20**, 1678-1705.
- [9] Kwok, Y.K., Leung, K.S. and Wong, H.Y. (2012) Efficient Options Pricing Using the Fast Fourier Transform. In: Duan, J.C., Ed., *Handbook of Computational Finance*, Springer, Berlin, 579-604. http://dx.doi.org/10.1007/978-3-642-17254-0_21
- [10] Fang, F. and Oosterlee, C.W. (2008) A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, **31**, 826-848. <http://dx.doi.org/10.1137/080718061>
- [11] Zhang, B., Grzelak, L.A. and Oosterlee, C.W. (2012) Efficient Pricing of Commodity Options with Earlyexercise under the Ornstein-Uhlenbeck Process. *Applied Numerical Mathematics*, **62**, 91-111. <http://dx.doi.org/10.1016/j.apnum.2011.10.005>
- [12] Heston, S.L. (1993) A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, **6**, 327-343.
- [13] Black, F. and Scholes, M. (1973) The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, **81**, 637-654. <http://dx.doi.org/10.1086/260062>