

OPTION PRICING USING THE FRACTIONAL FFT

KYRIAKOS CHOURDAKIS

ABSTRACT. This paper shows how the recently developed fractional FFT algorithm (FRFT) can be used to retrieve option prices from the corresponding characteristic functions. The FRFT algorithm has the advantage of using the characteristic function information in a more efficient way than the straight FFT. Therefore less function evaluations are typically needed and substantial savings in computational time can be made. Two experiments, based on the stochastic volatility and the variance-gamma models, illustrate the benefits of using the fractional version of the FFT and show that option prices can be delivered up to forty-five times faster without substantial losses of result accuracy.

1. INTRODUCTION

Using the characteristic function in order to price European option contracts has increased immensely in popularity, since the influential paper of Heston (1993) on stochastic volatility. The subsequent papers of Duffie, Pan, and Singleton (2000) and Bakshi and Madan (2000) shed more light on this relationship, and established that the delta and the price of the option can be retrieved by numerical integration. An increasing number of recently developed models for asset prices utilize the characteristic function to price European option contracts, or for the purpose of bond pricing, including the aforementioned stochastic volatility model of Heston (1993), the jump diffusion stochastic volatility model of Bates (1998), the affine jump diffusions of Duffie et al. (2000), the variance-gamma process of Madan, Carr, and Chang (1998) and the GARCH option pricing model of Heston and Nandi (2000), *inter alia*.

Unfortunately, the functions that need to be integrated are usually not sufficiently well behaved around zero, exhibiting singularities. For this reason, it is not straightforward to use the popular Fast Fourier Transform (FFT) to speed up the computations. Carr and Madan (1999) show how this problem can be overcome, by considering transforms of the option price itself, rather than using the delta/digital decomposition. In that way, the

power of FFT can be unleashed to price options in an accurate and efficient way.

The FFT is approximating the continuous Fourier transform (CFT) with its discrete counterpart (DFT), for a carefully chosen vector $\mathbf{h} = (h_j)_{j=0}^{N-1}$

$$(1.1) \quad \int_0^\infty e^{-ixu} h(u) du \approx \sum_{j=0}^{N-1} e^{-i\frac{2\pi}{N}kj} h_j$$

In this paper we will denote the DFT of a vector \mathbf{h} with $D_k(\mathbf{h})$. Vector \mathbf{h} corresponds to N function evaluations of the function h , at the points $\mathbf{u} = (u_j)_{j=0}^{N-1}$.

The output of the DFT procedure is a new vector $\mathbf{f} = (f_k)_{k=0}^{N-1}$. Each f_k corresponds to an integral of the form (1.1), computed at a *predetermined value* of $x = x_k$. In the context of option pricing, Carr and Madan (1999) show that the vector \mathbf{f} will contain the option prices that correspond to the log-strike prices kept in $\mathbf{x} = (x_k)_{k=0}^{N-1}$. We will denote with δ the grid size of vector \mathbf{u} , and with λ the grid size of \mathbf{x} .

It is well known (Bailey and Swarztrauber, 1991, 1994; Carr and Madan, 1999) that there is an inverse relationship between the grids sizes δ and λ . It follows that to obtain a fine grid across \mathbf{x} , one has to either increase δ thus making the grid across \mathbf{u} coarser, or increase the length of the input vector \mathbf{h} , effectively padding it with zeros. This can lead to a substantial waste of computational time. In addition, the resulting series will extend well beyond the range of log-prices \mathbf{x} that are actually required. As an example, out of the 4096 point FFT employed in Carr and Madan (1999), only 66 of the resulting option prices correspond to strike/spot ratios between 0.80 and 1.20.

In this paper we adapt the methodology of Carr and Madan (1999) in order to use the fractional FFT (FRFT), developed by Bailey and Swarztrauber (1991). The FRFT procedure can be used to rapidly compute sums of the form

$$(1.2) \quad \sum_{j=0}^{N-1} e^{-i2\pi kj\alpha} h_j$$

We will denote these sums with $D_k(\mathbf{h}, \alpha)$. The sums are computed by invoking two normal and one inverse FFT procedures. The benefits of using the FRFT approach is that both grid sizes δ and λ can be chosen independently. Although three FFT's are needed instead of just one, the freedom

of choosing both grid sizes suffices to save substantial computational time, since the input vectors are typically much shorter. This approach could prove beneficial when a large number of characteristic function evaluations are needed (for example when models are calibrated to observed European prices), or when computing the characteristic function is expensive (for example in some affine models where systems of differential equations have to be solved at each function evaluation). As an illustration, to produce 64 options for the strike range 0.80 to 1.20 one should need three $2 \times 64 = 128$ -point FFT's, which decreases the computation time by a factor of about 25, when compared with the Carr and Madan (1999) straight FFT methodology.

Section 2 shows how the approximation procedure and the application of the FRFT can be applied to the option pricing problem. An example, based on the Black-Scholes model is also presented. Section 3 compares the FRFT with the straight FFT of Carr and Madan (1999), in terms of their computation speed and their accuracy. The popular stochastic volatility model of Heston (1993) and the variance-gamma model of Madan et al. (1998) are used as benchmarks for this experiment. Finally, section 4 concludes.

2. OPTION PRICING WITH FRFT

When the characteristic function of the log-price is known explicitly, Carr and Madan (1999) show that the call option price for a log-strike price x can be computed as the transform

$$(2.1) \quad C_T(x) = \frac{e^{-\gamma x}}{\pi} \int_0^\infty e^{-ixu} \psi(u) du$$

where γ is a control parameter¹ and ψ is a function of the characteristic function of the log-price φ , given by

$$\psi(u) = \frac{e^{-rT} \varphi(u - (\gamma + 1)\mathbf{i})}{\gamma^2 + \gamma - u^2 + (2\gamma + 1)u\mathbf{i}}$$

The integral in equation (2.1) can be approximated using an integration rule, such as Simpson's or the trapezoidal rule, as

$$(2.2) \quad \int_0^\infty e^{-ixu} \psi(u) du \approx \sum_{j=0}^{N-1} e^{-ixu_j} \tilde{\psi}_j \delta$$

The points u_j are chosen to be equidistant with grid spacing δ , or $u_j = j\delta$. The value of δ should be sufficiently small to approximate the integral well

¹The choice of γ has to satisfy $ES_T^{\gamma+1} < +\infty$.

enough, while the value of $N\delta$ should be large enough to assume that the characteristic function is equal to zero for $u > \bar{u} = N\delta$. In general, the values $\tilde{\psi}_j$ are set equal to $\tilde{\psi}_j = \psi(u_j)w_j$, with w 's implementing the integration rule. In this paper, the trapezoidal rule is chosen to approximate the integral, and therefore the integration weights are set to $w_j = \frac{1}{2}$ if $j = 0$ or $j = N - 1$, and $w_j = 1$ otherwise.²

2.1. The FFT approach. Application of the FFT will result in a set of integral approximations of the form (2.2), computed for a set $(x_k)_{k=0}^{N-1}$. These values will also be equidistant, with grid spacing equal to λ , and say that we want them to be evenly spread around the at-the-money level of $x = 0$. Then, the values x_k would assume the form $x_k = -\frac{N\lambda}{2} + k\lambda$, for $k = 0 \cdots N - 1$. Now the summation in (2.2) can be written as

$$(2.3) \quad \sum_{j=0}^{N-1} e^{-ix_k u_j} \tilde{\psi}_j \delta = \sum_{j=0}^{N-1} e^{-i(-\frac{N\lambda}{2} + k\lambda)j\delta} \tilde{\psi}_j \delta = \sum_{j=0}^{N-1} e^{-i\lambda\delta jk} h_j$$

where in the last summation $h_j = e^{i\frac{N\lambda}{2}j\delta} \tilde{\psi}_j \delta$. Therefore, in order to apply the FFT, the procedure should be invoked on the vector $\mathbf{h} = (h_j)_{j=0}^{N-1}$. The value of the spacing λ is determined from the restriction (contrast equations 1.1 and 2.3)

$$(2.4) \quad \lambda\delta = \frac{2\pi}{N} \Rightarrow \lambda = \frac{2\pi}{N\delta} = \frac{2\pi}{\bar{u}}$$

It is apparent that out of the three parameters (δ , N and λ) only two can be chosen freely; the third will be determined by the restriction (2.4). Since δ and N are selected to make the integral approximations accurate enough, as discussed above, the value of λ will be inversely proportional to the upper integration bound. The resulting grid spacing might not be sufficiently dense for the purpose of option pricing, with option prices that are computed for log-strike prices x_k that are too large or too small to be of interest. In order to reduce the size of the grid one will have to increase the integration upper bound, padding the vector \mathbf{h} with zeros. Since the FFT is more efficient for values of N that are powers of two, the size of the FFT will have to double every time that a higher accuracy is demanded. For example, in the experiments of Carr and Madan (1999) the upper integration bound

²Other integration schemes are also straightforward to implement. For example, setting $w_j = \frac{1}{3} (3 + (-1)^{j-1} - \mathbf{1}_{j=0})$, with $\mathbf{1}$ being the indicator function, will implement Simpson's approximation. Experiments suggested that, for the majority of cases, the differences between the various schemes are insignificant.

is set equal to 1024 for a 4096 point FFT with $\delta = 0.25$. Equation (2.4) gives the corresponding log-strike grid spacing $\lambda = 0.6\%$. Out of the 4096 option prices that are calculated with this procedure only about 67 will fall within the $\pm 20\%$ log-strike interval, which is relevant for practical applications. This can be very inefficient in terms of computing power.

[Figure 1 somewhere here]

Figure 1 illustrates this point: Part (a) presents the real and imaginary parts of the integrand associated with the Black-Scholes model. Out of the 4096 points that are input to the FFT procedure, only about 80 are greater than 10^{-8} . Therefore a 4096-point grid is used where a 128-point one, integrating over $(0, 32)$, should suffice. Unfortunately, using a 128-point grid would result into $\lambda = 19\%$ which is impractical.

2.2. Application of the FRFT. The fractional FFT procedure will rapidly compute sums of the form $\sum_{j=0}^{N-1} e^{-i2\pi kj\alpha} h_j$, for any value of the parameter α . The standard FFT can be seen as a special case for $\alpha = \frac{1}{N}$. Therefore, the sum $\sum_{j=0}^{N-1} e^{-i\lambda\delta jk} h_j$ in (2.3) can be computed using FRFT, without the need of imposing the restriction $\lambda\delta = \frac{2\pi}{N}$. This means that the two grid spacings (among the characteristic function support and among the log-prices) can be chosen independently, with the FRFT parameter $\alpha = \delta\lambda$. For example, a 128-point FRFT will integrate the (modified) characteristic function of figure 1 over the interval $(0, 32)$ using the same spacing as the Carr and Madan 4096-point FFT. The resulting option prices will be computed for any log-price grid δ that we choose.

2.3. Implementation of FRFT. As discussed in Bailey and Swartztrauber (1991, 1994), the fractional transform can be easily implemented by invoking three $2N$ -point FFT procedures. Suppose that we want to compute an N -point FRFT on the vector $\mathbf{h} = (h_j)_{j=0}^{N-1}$. The following $2N$ -point vectors have to be defined

$$\mathbf{y} = \left((h_j e^{-i\pi j^2 \alpha})_{j=0}^{N-1}, (0)_{j=0}^{N-1} \right)$$

$$\mathbf{z} = \left((e^{i\pi j^2 \alpha})_{j=0}^{N-1}, (e^{i\pi(N-j)^2 \alpha})_{j=0}^{N-1} \right)$$

Then, the FRFT is given by

$$D_k(\mathbf{h}, \alpha) = (e^{-i\pi k^2 \alpha})_{k=0}^{N-1} \odot D_k^{-1} (D_j(\mathbf{y}) \odot D_j(\mathbf{z}))$$

where \odot denotes element-by-element vector multiplication. Note that the exponential quantities $(e^{i\pi j^2 \alpha})_{j=0}^{N-1}$ do not depend on the actual function that is integrated, and therefore can be pre-computed and stored.

2.4. Example: the Black-Scholes case. The Black-Scholes (BS) model is based on the assumption that the underlying asset, under risk neutrality, follows a geometric Brownian motion

$$\frac{dS_t}{S_t} = rdt + \sigma dW_t$$

After the normalization $S_0 = 1$, the (time t) log-price will be normally distributed, with corresponding characteristic function

$$\varphi(u) = \exp \left\{ iu \left(r - \frac{1}{2} \sigma^2 \right) t - \frac{1}{2} u^2 \sigma^2 t \right\}$$

Thus, following (2.1) the function to be passed to the FRFT routine will be

$$\psi(u) = \frac{\exp \left\{ \left(r + \frac{1}{2} (1 + \gamma + iu) \sigma^2 \right) (\gamma + iu) t \right\}}{\gamma + \gamma^2 + 2i\gamma u - u(u - i)}$$

In the following example we will assume zero interest rates, volatility $\sigma = 0.30$, time to maturity $t = 0.25$ years, while the dampening parameter is set at $\gamma = 4.00$. The real and imaginary parts of $\psi(u)$ are presented in figure 1. One can observe that since ψ decays exponentially, an upper integration bound of 32 should provide a sufficient approximation.³ A 128-point FRFT will imply a sampling interval $\delta = 0.25$, which is the same used in the 4096-point Carr-Madan straight FFT method. We also request the output to be a vector of option prices computed for strikes between $\exp\{\pm 0.20\}$, by setting the parameters $x_0 = -0.20$ and $\lambda = \frac{2 \times 0.20}{128} = 0.0031$. These choices also imply that the fractional parameter $\alpha = \frac{\delta \lambda}{2\pi} = 1.24 \times 10^{-4}$.

Let \mathbf{u} and \mathbf{x} denote the input and output grids, respectively. In order to implement FRFT, we will need to compute and store the following vectors:

$$\begin{aligned} \psi &= (\psi(u_j) \exp\{-x_0 i u_j\} \delta)_{j=0}^{127}, & \mathbf{a} &= (\exp\{i\alpha \pi j^2\})_{j=0}^{127}, \\ \bar{\mathbf{a}} &= (\exp\{i\alpha \pi (128 - j)^2\})_{j=0}^{127}, & \mathbf{w} &= (0.5, 1, 1, \dots, 1, 1, 0.5), \text{ and} \\ \mathbf{d} &= (\exp\{-\gamma x_j\} / \pi)_{j=0}^{127}. \end{aligned}$$

³Formal methods of estimating the upper integration bound are discussed in section 3 and in Lee (2004). Here, we pick $\bar{u} = 32$ to contrast a 128-point FRFT over $(0, 32)$ with the straight 4096-point FFT over $(0, 1024)$, which exhibits the same sampling interval.

Note that $\bar{\mathbf{a}}$ is equal to \mathbf{a} in reverse order, and therefore need not be explicitly computed. The vector of weights \mathbf{w} implements the trapezoidal integration rule, and \mathbf{d} is used to ‘un-dampen’ and recover the option prices in the final stage. We now have to compute the fractional transform (\odot and \oslash denote element-by-element vector multiplication and division, respectively)

$$\begin{aligned} \mathbf{y}_1 &= (\psi \oslash \mathbf{a} \odot \mathbf{w}, (0)_{j=0}^{127}), & \tilde{\mathbf{y}}_1 &= \text{FFT}(\mathbf{y}_1) \\ \mathbf{y}_2 &= (\mathbf{a}, \bar{\mathbf{a}}), & \tilde{\mathbf{y}}_2 &= \text{FFT}(\mathbf{y}_2) \\ \mathbf{y}_3 &= \tilde{\mathbf{y}}_1 \odot \tilde{\mathbf{y}}_2, & \tilde{\mathbf{y}}_3 &= \text{IFFT}(\mathbf{y}_3) \\ \mathbf{y} &= \tilde{\mathbf{y}}_3 \oslash \mathbf{a} \end{aligned}$$

and finally ‘un-dampen’ vector \mathbf{y} to recover the option prices

$$\mathbf{c} = \mathbf{y} \odot \mathbf{d}$$

Figure 1(b) presents the pricing errors of the FRFT and FFT methods. One can readily verify that both methods deliver errors that are virtually zero. It is also worth noting that the speed of the fractional transform is comparable to computing the closed-form BS formula. In particular, the 128-point FRFT implemented here is only about 20% slower than computing the BS formula for 128 strikes. A 64-point transform delivers prices as fast as the closed form, while a 16-point FRFT is actually about 20% faster.

2.5. Number of operations under FRFT. The number of operations needed to perform a standard M -point FFT is routinely estimated as $M \log_2 M$. Since one N -point FRFT will invoke three $2N$ -point FFT procedures, the number of operations will be approximately $6N \log_2(N)$. This implies that there exists a critical point at which the two procedures will have similar operation requirements. Figure 2 illustrates this point: it appears that the ratio of the theoretical FRFT over FFT operations is approximately equal to four. Therefore, option pricing using a 256-point FRFT should demand a similar number of elementary operations as a 1024-point FFT. Alternatively, using the FRFT procedure with vectors smaller than 1024 elements should demand less elementary operations than a 4096-point FFT.

[Figure 2 somewhere here]

Figure 3 shows the theoretical and experimental computational benefits of using the FRFT procedure. The horizontal axis gives the size of an FRFT

procedure, while the vertical axis gives the ratio of FFT/FRFT operations, for three different FFT implementations. The graph also presents the results of the experiments of section 3, where various fractional transform were compared to a 4096-point FFT. It appears that the theoretical values underestimate the performance of FRFT for larger transforms ($N > 64$), and somewhat overestimate the performance for smaller ones ($N < 64$). Our conjecture is that the theoretical estimates ignore memory allocation, which would inevitably slow down large transforms. The estimates also ignore the time required to setup the various FRFT vectors, which could be relatively significant for small transforms.

[Figure 3 somewhere here]

3. COMPARISON WITH FFT

In order to compare the pricing and computational performance of the FRFT to the FFT, we use two popular models that admit characteristic functions in closed form: the stochastic volatility (sv) model of Heston (1993), and the variance-gamma (vg) model of Madan et al. (1998). The characteristic functions of these models serve as examples of functions that decay exponentially (sv), and hyperbolically (vg).

Option prices were computed for a ‘benchmark’ parameter set, and for the sets constructed by changing a single parameter to a substantially lower or higher value, as shown in table 1. For each parameter set we focus on the errors over a grid of 31 equidistant strike prices in the (\$85,\$115) interval around the current price of \$100, where the time value effect is more pronounced.

[Table 1 somewhere here]

Since these models do not admit readily computed prices in closed form,⁴ the ‘true’ prices were computed using either an adaptive quadrature method to integrate the delta and the exercise probability independently, or an FFT

⁴Although the vg model actually offers option prices in closed form, they are represented in terms of various ‘special functions’ of mathematics. Experiments have shown that it is faster and more accurate to compute the option prices via Fourier inversion, rather than computing the corresponding special function analytically.

with a very wide support and a very dense grid.⁵ Three different FRFT procedures and one 4096-point FFT were used to compute option prices, which were then interpolated using a cubic spline, to arrive to the set of strike prices of interest.

3.1. Choice of the auxiliary parameters. In order to implement the transforms, we also have to choose the c.f. grid parameter, δ , and the dampening parameter, γ . In a recent article, Lee (2004) shows how to implement a selection strategy for these parameters, that minimizes the estimated truncation and sampling errors. This strategy is contract dependent, i.e. a different set of parameters is selected for each strike price; in our setting, since we simultaneously compute a *vector* of option prices, this procedure is not readily applicable. When selecting γ , Lee finds that the optimal values for in- and out-of-the-money contracts vary significantly. In particular, for in-the-money options large negative values for γ are typically used⁶, while for out-of-the-money contracts large positive values of γ are preferred. We found that choosing $\gamma = 4.00$ delivers very good results overall.

To select the upper integration bound, one approach would be to use Lee's procedure with respect to the truncation error alone, for different strike prices, and pick the most conservative estimate. We found that since we use more sampling points, a much simpler strategy can deliver faster and accurate results. As with any complex number, both absolute real and imaginary parts of the modified transform $\psi(u)$ are less or equal than its modulus

$$|\operatorname{Re}(\psi(u))| \leq |\psi(u)|, \text{ and } |\operatorname{Im}(\psi(u))| \leq |\psi(u)|.$$

Since ψ decays for large u , we can use the modulus $|\psi|$ to find a point \bar{u} , such that

$$(3.1) \quad |\psi(u)| \leq 10^{-m}, \text{ for all } u > \bar{u}$$

Since we are not interested in determining \bar{u} at a high precision level, we constrain ourselves to integer values, and employ a simple search method that finds the first integer that satisfies (3.1). Apparently, the threshold parameter m could increase as the sampling rate increases. The choices for the pairs (N, m) that we use in section 3 are reported in table 2. As an

⁵Where FFT was used, a 131072-point FFT over the interval $[0, 8192]$ was employed, implying a grid spacing of $1/16$ over its support. Where an adaptive quadrature was used, it was set to achieve an estimated accuracy of 10^{-7} .

⁶Note that when we set $\gamma < 0$, we are actually pricing the corresponding put.

example, table 2 also gives the resulting upper integration bounds and grid sizes for the benchmark parameter sets.

[Table 2 somewhere here]

One final important observation can be made regarding table 2: In the sv model, the joint increase of the required accuracy and the number of points, result in smaller grid sizes with roughly the same upper integration bound; in contrast, in the vg case we observe higher upper integration bounds and roughly the same sampling grid. This can be intuitively explained by contrasting the exponential decay of sv against the hyperbolic decay of vg. For this reason we chose to increase the parameter m for the sv model at a higher rate, as the transform points increase, relative to vg.

3.2. Pricing errors. The pricing errors for the 4096-point FFT and the corresponding FRFT procedures are reported in tables 3 and 4, for the sv and vg model, respectively. Since every numerical inversion typically takes a fraction of a second, and in order to increase the stability of the reported elapsed time, for each approximation and for each set of parameters 200 Fourier inversions were performed. The following statistics are reported:

$$\begin{aligned} \text{mean} &= \frac{1}{K} \sum_{k=1}^K C_k^{\text{CF}} - C_k^{\text{T}} \\ \text{RMSE} &= \sqrt{\frac{1}{K} \sum_{k=1}^K (C_k^{\text{CF}} - C_k^{\text{T}})^2} \\ \text{abs err} &= \max_{k=1 \dots K} |C_k^{\text{CF}} - C_k^{\text{T}}| \\ \% \text{ err} &= \max_{k=1 \dots K} \left| \frac{C_k^{\text{CF}} - C_k^{\text{T}}}{C_k^{\text{CF}}} \right| \end{aligned}$$

K denotes the number of call option prices considered, namely $K = 31$ for strikes ranging from \$85 to \$115 in increments of \$1. The quantities C^{CF} denote the ‘exact’ call price computed using the quadrature, while C^{T} denote the call prices computed using the corresponding transform (FFT or FRFT). All errors are reported in terms of 10^{-5} (for example 1.23 indicates 1.23×10^{-5}), and all relative errors are given in percentage terms (for example 0.010 indicates 0.010%). The relative gains in computational times are also reported in terms of the FFT elapsed time (for example a value of 35 indicates that the FRFT transform was 35 times faster than the corresponding FFT).

3.3. The stochastic volatility model. In this model the dynamics are described by the system of stochastic differential equations

$$\begin{aligned} \frac{dS_t}{S_t} &= rdt + \sigma_t dW_t \\ d\sigma_t^2 &= \theta(\bar{\sigma}^2 - \sigma_t^2)dt + \phi\sigma_t dZ_t \\ dW_t dZ_t &= \rho dt \end{aligned}$$

The characteristic function of the SV model takes the exponential affine form

$$\varphi(u) = \exp \{ C(u) + D(u)\sigma_0^2 + iu \log S_0 \}$$

with the functions C and D given in Heston (1993). The European call option with strike price K and maturity T can then be computed as

$$C = S_0\Pi_1 - Ke^{-rt}\Pi_2$$

where the Delta of the option Π_1 , and the risk neutral exercise probability Π_2 , are integrals of functions based on the characteristic function φ .

The parameters used in the calibrations are presented in table 1. Following Carr and Madan (1999), the grid spacing when using the FFT approach is set to 0.25, resulting in a log-strike grid of approximately 0.6. In the FRFT procedures the two grids can be chosen independently. Therefore, for an N -point FRFT the grid across log-strike was chosen to be $\lambda = \frac{2 \times 0.20}{N}$, implying a strike grid between $\exp\{\pm 0.20\}$, which is relatively tight around the strike prices of interest.

[Table 3 somewhere here]

Table 3 presents the pricing errors for the stochastic volatility case. We find that even a very small 16-point FRFT performs surprising well, delivering prices, for most parameter values, that lie within 0.1% of the theoretical ones. Two notable exceptions occur, where the maturity or the initial volatility is very small. The maximum absolute error reported is \$0.0034, which would be sufficient for most practical applications (note that the initial asset price is $S_0 = \$100$, and that the minimum possible bid-ask spread is \$0.05). These FRFT prices are delivered about 45 times faster than their FFT counterparts.

As the number of FRFT points increases, so does the accuracy of the results. A 32-point FRFT will deliver prices that are of the same order of accuracy as the FFT method, in roughly a 1/35 of the time. The 64-point

FRFT produces prices that are virtually indistinguishable from the 4096-point FFT ones.

3.4. The variance-gamma model. The variance-gamma model of Madan et al. (1998) serves as an example of a model that exhibits a hyperbolically decaying characteristic function. The c.f. is given by

$$\varphi(u) = \exp\{iu(\omega + r)t + iu \log S_0\} \left(1 - iu\theta\nu - \frac{1}{2}u^2\sigma^2\nu\right)^{-\frac{t}{\nu}}$$

where

$$\omega = \frac{1}{\nu} \log\left\{1 - \theta\nu - \frac{1}{2}\sigma^2\nu\right\}$$

Parameter ν is a measure of trading intensity, which controls for the kurtosis of the risk neutral density, while parameter θ controls for the skewness. The parameter values used to assess the accuracy of the FRFT are given in table 1. As Lee (2004, equation A2, and section 6.5) shows, the c.f. decays hyperbolically at a rate $u^{-\frac{2t}{\nu}}$. This slow decay has an adverse effect when the time-to-maturity is small ($t = 0.10$), with the upper integration bound increasing rapidly. In these cases, the threshold parameters m were set at values which are lower than the ones presented in table 2, effectively reducing the integration interval. The pairs (N, m) chosen in these cases were $(32, 4)$, $(64, 5)$ and $(128, 6)$.

[Table 4 somewhere here]

Table 4 presents the pricing errors for the variance-gamma specification. Due to the hyperbolic decay of the modified characteristic function ψ , we need to integrate over an interval which is typically wider than the SV ones. In addition, we expect the truncation errors to be somewhat higher. For medium and long dated options, the 32-point FFT produces a maximum absolute error of \$0.0070, which is still within the one penny bound.⁷ For shorter times to maturity, the 32-point FRFT produces a maximum error of \$0.06, which is larger than the minimum \$0.05 bid-ask spread. In these cases, and if higher accuracy is required, the 64-point FRFT will deliver prices that are comparable to the FFT ones, in about 1/25 of the time. The maximum absolute difference between the 64-point FRFT and the FFT is about \$0.0070.

⁷This error is reported where the volatility is low, together with a relative error of 3%, indicating that it is associated with a contract which is far-out-of-the-money.

A 128-point FRFT produces prices that are equal to their FFT counterparts, for most practical applications.

4. CONCLUSIONS AND EXTENSIONS


This paper proposes a new algorithm that inverts the modified characteristic function in order to compute option prices. This algorithm, coined the fractional FFT has the advantage of using the characteristic function information in a more efficient way than the straight FFT approach. Therefore less function evaluations are typically needed. Experiments, based on the stochastic volatility model of Heston (1993) and the variance-gamma model of Madan et al. (1998), are used to illustrate the benefits of using the fractional version of the FFT. It is found that very small transforms can compute call options up to forty-five times faster than the straight FFT of Carr and Madan (1999), without substantial losses in the accuracy of the results.

The methodology presented here can be modified in order to create fast adaptive procedures that compute Fourier transforms. For example instead of choosing a single value \bar{u} that bounds $|\varphi(u)| < 10^{-m}$ for $u > \bar{u}$, one can choose successive bounds $\bar{u}_1 \cdots \bar{u}_n$ with estimated accuracies $m_1 \cdots m_n$ and compute the corresponding areas piecewise. Another interesting extension would be the modification of the algorithm to cope with two- or multi-dimensional Fourier transforms, used to value spread or correlation contracts; see for example Dempster and Hong (2000).


REFERENCES

- Bailey, D. H. and P. N. Swarztrauber (1991). The Fractional Fourier Transform and applications. *SIAM Review* 33(3), 389–404.
- Bailey, D. H. and P. N. Swarztrauber (1994). A fast method for the numerical evaluation of continuous Fourier and Laplace transforms. *SIAM Journal on Scientific Computing* 15(5), 1105–1110.
- Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics* 55, 205–238.
- Bates, D. S. (1998). Pricing options under jump diffusion processes. Technical Report 37/88, The Wharton School, University of Pennsylvania.
- Carr, P. and D. Madan (1999). Option valuation using the Fast Fourier Transform. *Journal of Computational Finance* 3, 463–520.
- Dempster, M. A. H. and S. S. G. Hong (2000). Spread option valuation and the Fast Fourier Transform. Technical Report WP 26/2000, The Judge Institute of Management Studies, University of Cambridge.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump–diffusions. *Econometrica* 68, 1343–1376.
- 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, 327–344.
- Heston, S. L. and S. Nandi (2000). A closed-form GARCH option pricing model. *Review of Financial Studies* 13(3), 585–625.
- Lee, R. (2004). Option pricing by transform methods: Extensions, unification and error control. *Journal of Computational Finance* 7(3), 51–86.
- Madan, D., P. Carr, and E. Chang (1998). The variance-gamma process and option pricing. *European Finance Review* 2, 79–105.

FIGURE 1. (a) The real (top) and imaginary (bottom) parts of a function $\psi(u)$ which is passed to the FRFT routine. The Black-Scholes model with parameters $\{\sigma, r, t\} = \{0.30, 0.00, 0.25\}$ is used. The current price is $S_0 = \$1$ and the dampening parameter $\gamma = 4.00$. The function is sampled at intervals $\delta = 0.25$, for a 128-point FRFT. (b) The pricing errors of the FFT and the FRFT methods for the same BS specification.



(a) Sampled function ψ



(b) Pricing Errors

FIGURE 2. Number of elementary operations needed to invoke one N -point FFT or FRFT procedure.




FIGURE 3. Speed increase when using the FRFT procedure. The curves labeled (th) represent the theoretical ratio of the number of elementary operations with respect to three different FFT sizes, while the curve labeled (act) represents the actual speed increases as computed in the experiments of section 3 for the 4096-point FFT.




TABLE 1. Calibration parameters for the SV and the VG models.

Stochastic volatility				Variance gamma			
	Bench	Low	High		Bench	Low	High
S_0	\$100·			S_0	\$100·		
r	0·00			r	0·00		
t	0·25	0·10	1·00	t	0·25	0·10	1·00
σ_0	0·30	0·10	0·90	σ	0·30	0·15	0·45
θ	3·00	1·00	9·00	θ	-0·20	-0·10	-0·30
$\bar{\sigma}$	0·30			ν	0·20	0·10	0·30
ϕ	0·15	0·05	0·45				
ρ	-0·50	0·00	-0·90				

TABLE 2. Typical choices for the integration accuracy m (significant figures), the upper integration bound \bar{u} , and the corresponding grid δ , for different FRFT procedures. The benchmark parameter sets of table 1 were used.

	Stochastic volatility			Variance gamma		
	m	\bar{u}	δ	m	\bar{u}	δ
4096-FFT		1024	0.25		1024	0.25
16-FRFT	5	24	1.50			
32-FRFT	7	32	1.00	5	50	1.56
64-FRFT	9	39	0.61	6	83	1.30
128-FRFT				7	140	1.09

TABLE 3: Comparison of the FRFT and FFT procedures for the stochastic volatility model. The straight FFT is compared with a 16-, 32- and 64-point FRFT. The numbers in brackets show the average relative speed increases with respect to FFT (for example a value $\times 10$ indicates that the FRFT performed ten times faster than FFT). The mean difference between all approximate prices and the corresponding closed form solutions are presented, as well as the root mean square error (RMSE), the maximum absolute error and the maximum relative error. The first three are given in multiples of 10^{-5} , while the relative error is in percentage terms (a value of 0.01 indicates 0.01%). Strike prices between \$85 and \$115 in increments of \$1 were used for this calculation.

		bench	$\sigma_0 = 0.10$	$\sigma_0 = 0.90$	$\theta = 1.00$
FFT	mean	-\$1.10	-0.749	-4.98	-0.703
	RMSE	\$1.34	1.98	4.99	1.66
	abs err	\$4.09	7.38	5.75	7.27
	% err	0.0010%	0.0044	0.0006	0.0021
FRFT 16 ($\times 44$)	mean	-19.7	-151.	20.5	-22.1
	RMSE	81.0	187.	48.9	104.
	abs err	15.1	340.	85.9	198.
	% err	0.0422	0.2829	0.0040	0.0505
FRFT 32 ($\times 35$)	mean	-0.988	-1.79	5.00	-0.607
	RMSE	1.25	2.83	5.02	1.74
	abs err	2.99	8.58	5.67	7.67
	% err	0.0012	0.0048	0.0006	0.0022
FRFT 64 ($\times 26$)	mean	-1.10	-0.748	-4.98	-0.704
	RMSE	1.34	1.97	4.99	1.66
	abs err	4.10	7.37	5.75	7.27
	% err	0.0010	0.0043	0.0006	0.0021

Continued on the next page

TABLE 3 : Continued from the previous page

		$\theta = 9.00$	$\phi = 0.05$	$\phi = 0.45$	$\rho = 0.00$
FFT	mean	-\$1.02	-0.974	-1.12	-0.967
	RMSE	\$1.07	1.01	1.34	1.00
	abs err	\$1.75	1.60	3.56	1.47
	% err	0.0010%	0.0011	0.0013	0.0010
FRFT 16	mean	-21.5	-15.7	-14.7	-11.3
	RMSE	104.	80.7	86.9	75.3
	abs err	198.	160.	152.	165.
	% err	0.0505	0.0363	0.0703	0.0299
FRFT 32	mean	-0.965	-1.04	-1.27	-1.17
	RMSE	1.15	1.19	1.71	1.38
	abs err	2.55	2.00	3.12	2.41
	% err	0.0013	0.0013	0.0019	0.0011
FRFT 64	mean	-1.02	-0.973	-1.12	-0.967
	RMSE	1.07	1.01	1.34	0.999
	abs err	1.74	1.60	3.55	1.46
	% err	0.0011	0.0011	0.0013	0.0010

		$\rho = -0.90$	$t = 0.10$	$t = 1.00$
FFT	mean	-\$0.245	-0.591	-3.08
	RMSE	\$3.30	0.797	3.09
	abs err	\$17.4	2.79	3.69
	% err	0.0075%	0.0032	0.0006
FRFT 16	mean	-10.4	-92.1	-4.10
	RMSE	82.8	138.	27.3
	abs err	151.	268.	53.3
	% err	0.0388	0.1482	0.0031
FRFT 32	mean	-0.310	-0.756	-3.09
	RMSE	3.36	1.25	3.11
	abs err	17.5	2.80	3.78
	% err	0.0075	0.0015	0.0006
FRFT 64	mean	-0.245	-0.592	-3.08
	RMSE	3.30	0.798	3.09
	abs err	17.4	2.79	3.69
	% err	0.0075	0.0032	0.0006

TABLE 4: Comparison of the FRFT and FFT procedures for the variance gamma model. The straight FFT is compared with a 32-, 64- and 128-point FRFT. The numbers in brackets show the average relative speed increases with respect to FFT (for example a value $\times 10$ indicates that the FRFT performed ten times faster than FFT). The mean difference between all approximate prices and the corresponding closed form solutions are presented as, well as the root mean square error (RMSE), the maximum absolute error and the maximum relative error. The first three are given in multiples of 10^{-5} , while the relative error is in percentage terms (a value of 0.01 indicates 0.01%). Strike prices between \$85 and \$115 in increments of \$1 were used for this calculation.

		bench	$\sigma_0 = 0.15$	$\sigma_0 = 0.45$	$\theta = -0.10$
FFT	mean	\$15.9	20.0	12.8	16.4
	RMSE	\$19.8	28.8	15.4	20.6
	abs err	\$44.2	75.1	31.3	46.3
	% err	0.0153%	0.1029	0.0054	0.0124
FRFT 32 ($\times 33$)	mean	34.9	157.	15.6	23.6
	RMSE	218.	305.	168.	210.
	abs err	405.	696.	296.	382.
	% err	0.1103	3.0449	0.0377	0.0779
FRFT 64 ($\times 22$)	mean	16.4	-5.37	11.3	17.4
	RMSE	34.5	41.7	24.6	33.7
	abs err	88.6	116.	65.4	85.9
	% err	0.0279	0.4554	0.0076	0.0169
FRFT 128 ($\times 15$)	mean	16.0	18.3	12.8	16.3
	RMSE	20.3	27.6	15.0	20.8
	abs err	52.5	78.7	27.4	48.0
	% err	0.0165	0.0818	0.0052	0.0117

Continued on the next page

TABLE 4 : Continued from the previous page

		$\theta = -0.30$	$\nu = 0.10$	$\nu = 0.30$	$t = 0.10$
FFT	mean	\$15.2	15.5	16.3	23.8
	RMSE	\$18.4	18.5	21.3	45.3
	abs err	\$39.1	36.8	53.2	167.
	% err	0.0187%	0.0154	0.0149	0.0780
FRFT 32	mean	13.6	3.16	-44.7	104.
	RMSE	193.	137.	295.	2818.
	abs err	348.	245.	691.	6314.
	% err	0.1232	0.0813	0.1974	3.1797
FRFT 64	mean	13.7	15.5	-2.92	8.53
	RMSE	34.0	24.1	41.7	401.
	abs err	83.3	53.8	123.	871.
	% err	0.0222	0.0116	0.0282	0.4779
FRFT 128	mean	15.2	15.4	8.23	17.1
	RMSE	18.6	18.5	16.7	56.7
	abs err	35.3	37.7	59.8	147.
	% err	0.0171	0.0152	0.0153	0.0811
$t = 1.00$					
FFT	mean	8.56			
	RMSE	9.95			
	abs err	16.4			
	% err	0.0025			
FRFT 32	mean	8.85			
	RMSE	10.2			
	abs err	17.3			
	% err	0.0026			
FRFT 64	mean	8.91			
	RMSE	10.2			
	abs err	16.9			
	% err	0.0026			
FRFT 128	mean	8.66			
	RMSE	9.99			
	abs err	16.2			
	% err	0.0025			

DEPT OF ECONOMICS, QUEEN MARY, UNIVERSITY OF LONDON, LONDON E1 4NS,
UNITED KINGDOM

E-mail address: k.chourdakis@qmul.ac.uk, www.theponytail.net