A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes

Roger Lord and Fang Fang and Frank Bervoets and Kees Oosterlee

Rabobank International, Delft University of Technology and Center for Mathematics and Computer Science (CWI), Amsterdam

28. February 2007

A FAST AND ACCURATE FFT-BASED METHOD FOR PRICING EARLY-EXERCISE OPTIONS UNDER LÉVY PROCESSES

R. LORD*, F. FANG†, F. BERVOETS‡, AND C.W. OOSTERLEE§

Abstract. A fast and accurate method for pricing early exercise and certain exotic options in computational finance is presented. The method is based on a quadrature technique and relies heavily on Fourier transformations. The main idea is to reformulate the well-known risk-neutral valuation formula by recognising that it is a convolution. The resulting convolution is dealt with numerically by using the Fast Fourier Transform (FFT). This novel pricing method, which we dub the Convolution method, CONV for short, is applicable to a wide variety of payoffs and only requires the knowledge of the characteristic function of the model. As such the method is applicable within exponentially Lévy models, including the exponentially affine jump-diffusion models. For an M-times exercisable Bermudan option, the overall complexity is $O(MN\log(N))$ with N grid points used to discretise the price of the underlying asset. It is shown how to price American options efficiently by applying Richardson extrapolation to the prices of Bermudan options.

Key words. option pricing , Bermudan options, American options, convolution, Lévy Processes, Fast Fourier Transform

AMS subject classifications. 65Y20, 65T50, 62P05, 60E10, 91B28

Preferred short title : CONV method for option pricing

1. Introduction. When valuing and risk-managing exotic derivatives, practitioners demand fast and accurate prices and sensitivities. As the financial models and option contracts used in practice are becoming increasingly complex, efficient methods have to be developed to cope with such models. Aside from non-standard exotic derivatives, plain vanilla options in many stock markets are actually of the American type. As any pricing and risk management system has to be able to calibrate to these plain vanilla options, it is of the utmost importance to be able to value these American options quickly and accurately.

By means of the risk-neutral valuation formula the price of any option without early exercise features can be written as an expectation of the discounted payoff of this option. Starting from this representation one can apply several numerical techniques to calculate the price itself. Broadly speaking one can distinguish three types of methods: Monte Carlo simulation, numerical solution of the corresponding partial-(integro) differential equation (P(I)DE) and numerical integration. While the treatment of early exercise features within the first two techniques is relatively standard, the pricing of such contracts via quadrature pricing techniques has not been considered until recently, see [1, 32]. Each of these methods has its merits and demerits, though for the pricing of American options the PIDE approach currently seems to be the clear favorite [19, 34].

In the past couple of years a vast body of literature has considered the modeling of asset returns as infinite activity Lévy processes, due to the ability of such processes to adequately describe the empirical features of asset returns and at the same time provide a reasonable fit to the implied volatility surfaces observed in option markets. Valuing American options in such models is however far from trivial,

*Financial Engineering (UC-R-380), Rabobank International, P.O. Box 17100, Utrecht, The Netherlands, e-mail: roger.lord@rabobank.com

†Delft University of Technology, Delft Institute of Applied Mathematics, Delft, the Netherlands, email: f.fang@ewi.tudelft.nl,

‡Modelling and Research (TC 2121), Rabobank International, Thames Court, 1 Queenhithe, London EC4V 3RL, UK, e-mail: frank.bervoets@rabobank.com,

§CWI – Center for Mathematics and Computer Science, Amsterdam, the Netherlands, email: c.w.oosterlee@cwi.nl, and Delft Institute of Applied Mathematics, Delft University of Technology.

due to the weakly singular kernels of the integral terms appearing in the PIDE, as reported in, e.g., [3, 4, 11, 20, 28, 33].

In this paper we present a novel quadrature-based method for pricing options with early exercise features. The method effectively combines the recent quadrature pricing methods of [1] and [32] with the methods based on Fourier transformation pioneered by [8, 29, 26]. Though the transform methods so far have mainly been used for the pricing of European options, we show how early exercise features can be incorporated naturally. The only requirement of the method is that the conditional characteristic function of the underlying asset is known, which is the case for many exponential Lévy models, with the popular exponentially affine jump-diffusion (EAJD) models of [13] as an important subclass. In contrast to the PIDE methods, processes of infinite activity, such as the Variance Gamma (VG) or CGMY models can be handled with relative ease. In addition to its flexibility, a real benefit of our method is its impressive computational speed, as all integrations can be evaluated using the FFT algorithm.

This paper is organized as follows. We start with an overview of the recent history of the FFT in option pricing. Subsequently we introduce the novel method called Convolution (CONV) method for early exercise options. Its high accuracy and speed are demonstrated by pricing several Bermudan and American options under Geometric Brownian Motion (GBM), VG and CGMY.

2. Overview Transform and Quadrature Pricing Methods. All transform methods depart from the risk-neutral valuation formula that, for a European option, reads:

$$V(t, S(t)) = e^{-r\tau}\mathbb{E}\left[V(T, S(T))\right], \tag{1}$$

where V denotes the value of the option, r is the risk-neutral interest rate, t is the current time point, T is the maturity of the option and $\tau = T - t$. The variable S denotes the asset on which the option contract is based. The expectation is taken with respect to the risk-neutral probability measure. Although we assume throughout the paper that interest rates are deterministic, this assumption can be relaxed at the cost of increasing the dimensionality of some of the methods. As (1) is an expectation, it can be calculated via numerical integration provided that the probability density is known in closed-form.

This is not the case for many models which do however have a characteristic function in closed form.[1] A number of papers starting from Heston [18] have attacked the problem via another route. Focusing on a plain vanilla European call option, note that (1) can be written very generally as:

$$V(t, S(t)) = e^{-r\tau}(F(t,T) \cdot \Delta - K \cdot \mathbb{P}(S(T) > K)), \tag{2}$$

where $F(t,T)$ is the forward price of the underlying asset at time T, as seen from t, $\mathbb{P}(S(T) > K)$ is the risk-neutral probability of ending up in-the-money and Δ is the delta of the option, the sensitivity of the option with respect to changes in the underlying. Note that (2) has the same form as the celebrated Black-Scholes formula. The delta can be interpreted as the probability of ending up in the money under the stock price measure, induced by taking the asset price itself as the numeraire asset. As such, both these cumulative probabilities can be found by inverting the characteristic function, an approach which in the form used here dates back to

[1] Or, the probability density involves complicated special functions whereas the characteristic function is comparatively easier.

Gurland [17] and Gil-Pelaez [16]. We can write:

$$\mathbb{P}(S_T > K) = \frac{1}{2} + \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-iuk}\phi(u)}{iu}du, \tag{3}$$

$$\Delta = \frac{1}{2} + \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-iuk}\phi(u-i)}{iu\phi(-i)}du, \tag{4}$$

where i is the imaginary unit, k is the logarithm of the strike price K and ϕ is the characteristic function of the log-underlying, i.e.,

$$\phi(u) = \mathbb{E}\left[e^{iu \ln S(T)}\right].$$

Carr and Madan [8] considered another approach. Note that L^1-integrability is a sufficient condition for the Fourier transform of a function to exist. A call option is certainly not L^1-integrable with respect to the logarithm of the strike price, as:

$$\lim_{k\to-\infty} V(t, S(t)) = S(t),$$

Damping the option price with $\exp(\alpha k)$ for $\alpha > 0$ solves this however, and Carr and Madan ended up with:

$$\begin{aligned}\mathcal{F}\{e^{\alpha k}V(t,k)\} &= e^{-r\tau}\int_{-\infty}^{\infty} e^{iuk}\mathbb{E}\left[(S(T)-e^k)^+\right]dk \\ &= \frac{e^{-r\tau}\phi(u-(\alpha+1)i)}{-(u-\alpha i)(u-(\alpha+1)i)},\end{aligned} \tag{5}$$

where with abuse of notation we now consider the option price V as a function of time and k. Though this approach was new to mathematical finance, the idea of damping functions on the positive real line in order to be able to find their Fourier transform dates back to at least Dubner and Abate [12].

A necessary and sufficient condition for (5) to exist is that

$$\phi(-(\alpha+1)i) = \mathbb{E}[S(T)^{\alpha+1}] < \infty,$$

i.e., that the $(\alpha+1)^{th}$ moment of the asset price exists. The option price can subsequently be recovered by inverting (5) and undamping

$$V(t,k) = \frac{1}{2\pi}e^{-r\tau-\alpha k}\int_{-\infty}^{\infty} e^{-iuk}\frac{\phi(u-(\alpha+1)i)}{-(u-\alpha i)(u-(\alpha+1)i)}du \tag{6}$$

The representation in (6) has two distinct advantages over (3). Firstly, it only requires one numerical integration. Secondly, whereas (2) can suffer from cancellation errors, the numerical stability of (6) can be controlled by means of the damping coefficient α. Finally we note that if we discretise (6) with Newton-Côtes quadrature the option price can very efficiently be evaluated by means of the FFT, yielding option prices over a whole range of strike prices.

The methods considered up till here can only handle the pricing of European options. Before turning to methods that can handle early exercise features, let us introduce some notation. We define the set of exercise dates as $\mathcal{T} = \{t_1, \ldots, t_M\}$ and $0 = t_0 \le t_1$. For ease of exposure we assume the exercise dates are equally spaced, so that $t_{m+1} - t_m = \Delta t$. The best known examples of options with early exercise are American and Bermudan options. American options can be exercised at any time prior to the option's expiry, whereas Bermudan options can only be exercised at certain dates in the future. If the option is exercised at some time $t \in \mathcal{T}$ the holder of the option obtains the exercise payoff $E(t, S(t))$. The Bermudan option price can then be found via backward induction as

$$\begin{cases} V(t_M, S(t_M)) &= E(t_M, S(t_M)) \\ C(t_m, S(t_m)) &= e^{-r\Delta t}\mathbb{E}_{t_m}\left[V(t_{m+1}, S(t_{m+1}))\right] \qquad m = M-1, \ldots, 1, \\ V(t_m, S(t_m)) &= \max\{C(t_m, S(t_m)), E(t_m, S(t_m))\}, \end{cases} \tag{7}$$

with C the continuation value of the option and V the value of the option immediately before the exercise opportunity. Note that we now explicitly attached a subscript to the expectation operator to indicate that the expectation is being taken with respect to all information available at time t_m.

Clearly the dynamic programming problem in (7) is a successive application of the risk-neutral valuation formula, as we can write the continuation value as

$$C(t_m, S(t_m)) = e^{-r\Delta t}\int_{-\infty}^{\infty} V(t_{m+1}, y) f(y|S(t_m))dy, \tag{8}$$

where $f(y|S(t_m))$ represents the probability density describing the transition from $S(t_m)$ at t_m to y at t_{m+1}. Based on (7) and (8) the QUAD method was introduced in [1]. The method requires the transition density to be known in closed-form, which is the case in e.g. the Black-Scholes model and Merton's jump-diffusion model. This requirement is relaxed in [32], where the QUAD-FFT method is introduced. The underlying idea is that the transition density can be recovered by inverting the characteristic function, opening up the QUAD method to a much wider range of models. As such the QUAD-FFT method effectively combines the QUAD method with the early transform methods. The overall complexity of both methods is $O(MN^2)$ for an M-times exercisable Bermudan option with N grid points used to discretise the price of the underlying asset.

The complexity of this method can be improved to $O(MN\log(N))$ if the underlying is a monotone function of a Lévy process. We will demonstrate this shortly. In the remainder we assume, as is common, that the underlying process is modelled as an exponential of a Lévy process. Let $x_1, \ldots, x_N$ be a uniform grid for the log-asset price. If we discretise (8) by the trapezoidal rule we can write the continuation value in matrix form as

$$\mathbf{C}(t_m) \approx e^{-r\Delta t}\Delta x\left[\mathbf{F}\mathbf{V} - \frac{1}{2}\left(V\left(t_{m+1}, x_1\right)\mathbf{f}_1 + V\left(t_{m+1}, x_N\right)\mathbf{f}_N\right)\right], \tag{9}$$

where

$$\mathbf{f}_i = \begin{pmatrix} f(x_i|x_1) \\ \vdots \\ f(x_i|x_N) \end{pmatrix}, \quad \mathbf{F} = (\mathbf{f_1}, \ldots, \mathbf{f_N}), \quad \mathbf{V} = \begin{pmatrix} V(t_{m+1}, x_1) \\ \vdots \\ V(t_{m+1}, x_N) \end{pmatrix},$$

and $f(y|x)$ now denotes the transition density in logarithmic coordinates. The key observation is that the increments of Lévy processes are independent, so that due to the uniform grid

$$\mathbf{F}_{j,\ell} = f(y_j|y_\ell) = f(y_{j+1}|y_{\ell+1}) = \mathbf{F}_{j+1,\ell+1}; \tag{10}$$

The matrix $\mathbf{F}$ is hence a Toeplitz matrix. A Toeplitz matrix can easily be represented as a circulant matrix, which has the convenient property that the FFT algorithm can be employed to efficiently calculate matrix-vector multiplications. Therefore, an overall computational complexity of $O(MN\log(N))$ can be achieved. Though this method is significantly faster than [1] or [32], we do not pursue it in this paper as the method we develop in the next section requires less operations, though the complexity remains the same.

The previous strain of literature does not seem to have picked up on a presentation by Reiner [30], where it was recognised that for the Black-Scholes model the risk-neutral valuation formula in (8) can be seen as a convolution or correlation of the continuation value with the transition density. As convolutions can be handled very efficiently by means of the FFT, an overall complexity of $O(MN \log N)$ can be achieved. By working forward instead of backward in time a number of discrete path-dependent options can also be treated, such as lookbacks, barriers, Asian options and cliquets. Building on Reiner's idea, Broadie and Yamamoto [6] have been able to reduce the complexity to $O(MN)$ for the Black-Scholes model by combining the double-exponential integration formula and the Fast Gauss Transform. Naturally their technique is applicable to any model in which the transition density can be written as a weighted sum of Gaussian densities, which is the case in e.g. Merton's jump-diffusion model.

As one of the defining properties of a Lévy process is that its increments are independent of each other, the insight of Reiner has a much wider applicability than only to the Black-Scholes model. This is especially appealing since the usage of Lévy processes in finance has become more established nowadays. By combining Reiner's ideas with the work of Carr and Madan, we introduce the Convolution method, or CONV method for short. The complexity of the method is $O(MN \log N)$ for an M-times exercisable Bermudan option.

3. The CONV Method. The main premise of the CONV method is that the conditional probability density $f(y|x)$ in (8) only depends on x and y via their difference

$$f(y|x) = f(y - x). \tag{11}$$

Note that x and y do not have to represent the asset price directly, they could be monotone functions of the asset price. The assumption made in (11) therefore certainly holds when the asset price is modelled as a monotone function of a Lévy process, since one of the defining properties of a Lévy process is that its increments are independent of each other. As mentioned earlier, we choose to work with exponential Lévy models in the remainder of this paper. In this case x and y in (11) represent the log-spot price. Let us see what the impact of independent increments is on the continuation value in (8). By including (11) in (8) and changing variables $z = y - x$ the continuation value can be expressed as

$$C(t_m, x) = e^{-r\Delta t} \int_{-\infty}^{\infty} V(t_{m+1}, x + z) f(z) dz, \tag{12}$$

which is a cross-correlation[2] of the option value at time t_{m+1} and the density $f(z)$, or equivalently, a convolution of $V(t_{m+1})$ and the conjugate of $f(z)$. If the density function has an easy closed-form expression, it may be beneficial to proceed along the lines of (9). However, for many exponential Lévy models we either do not have a closed-form expression for the density (e.g. the CGMY/KoBoL model of [5] and [7] and many EAJD models), or if we have, it involves one or more special functions (e.g. the VG model). In contrast, the characteristic function of the log-spot price can typically be found in closed-form or, in case of the EAJD models, via the solution of a system of ODEs.

[2]The cross-correlation of two functions $f(t)$ and $g(t)$, denoted $f \star g$, is defined by

$$f \star g \equiv \bar{f}(-t) * g(t) = \int_{-\infty}^{\infty} f(\tau) g(t + \tau) d\tau,$$

where '$*$' denotes the convolution operator.

Let us therefore take the Fourier transform of (12). The insight that the continuation value can be seen as a convolution is particularly useful here, as the Fourier transform of a convolution is merely the product of the Fourier transforms of the two functions being convolved. In the remainder we will employ the following definitions for the continuous Fourier transform and its inverse,

$$\hat{h}(u) := \mathcal{F}\{h(t)\}(u) \quad = \int_{-\infty}^{\infty} e^{iut} h(t)dt, \tag{13}$$

$$h(t) := \mathcal{F}^{-1}\{\hat{h}(u)\}(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-iut}\hat{h}(u)du. \tag{14}$$

If we dampen the continuation value (12) by a factor $\exp(\alpha x)$ and subsequently take its Fourier transform, we arrive at

$$\begin{aligned} e^{r\Delta t}\mathcal{F}\{c(t_m,x)\}(u) &= \int_{-\infty}^{\infty} e^{iux}e^{\alpha x}\int_{-\infty}^{\infty} V(t_{m+1},x+z)f(z)dzdx \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{iu(x+z)}v(t_{m+1},x+z)e^{-iz(u-i\alpha)}f(z)dzdx. \end{aligned} \tag{15}$$

where in the first step we used the risk-neutral valuation formula from (12). We introduced the convention that small letters indicate damped quantities, i.e., $c(t_m,x) = e^{\alpha x}C(t_m,x)$ and $v(t_m,x+z) = e^{\alpha(x+z)}V(t_m,x+z)$. Changing the order of integration and remembering that $x = y - z$, we obtain

$$\begin{aligned} e^{r\Delta t}\mathcal{F}\{c(t_m,x)\}(u) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{iuy}v(t_{m+1},y)dy\; e^{-i(u-i\alpha)z}f(z)dz \\ &= \int_{-\infty}^{\infty} e^{iuy}v(t_{m+1},y)dy\int_{-\infty}^{\infty} e^{-i(u-i\alpha)z}f(z)dz \\ &= \mathcal{F}\{e^{\alpha y}V(t_{m+1},y)\}(u)\;\phi(-(u-i\alpha)). \end{aligned} \tag{16}$$

In the last step we used the fact that the complex-valued Fourier transform of the density is simply the extended characteristic function

$$\phi(x+yi) = \int_{-\infty}^{\infty} e^{i(x+yi)z}f(z)dz, \tag{17}$$

which is well-defined when $\phi(yi) < \infty$, as $|\phi(x+yi)| \leq |\phi(yi)|$. As such (16) puts a condition on the damping coefficient α, because $\phi(\alpha i)$ must be finite.

The difference with the Carr-Madan approach in (5) is that we take a transform with respect to the log-spot price instead of the log-strike price, something which [26] and [29] also consider for European option prices. The damping factor is again certainly necessary when considering e.g. a Bermudan put, as then $V(t_{m+1},x)$ tends to a constant when $x \to -\infty$, and as such is not L^1-integrable. For the Bermudan put we must choose $\alpha > 0$. Though other values of α are allowed in principle, we need to know the poles of the payoff-transform in order to apply Cauchy's residue theorem, see e.g. [23] and [24]. This restriction on α will disappear when we switch to a discretised version of (16) in the next section. The Fourier transform of the damped continuation value can thus be calculated as the product of two functions, one of which, the extended characteristic function, is readily available in exponential Lévy models. How we proceed should be fairly clear. We recover the continuation value by taking the inverse Fourier transform of the right-hand side of (16), and calculate $V(t_m)$ as the maximum of the continuation and the exercise value at t_m. We repeat (7) recursively until we have obtained the option price at time t_0. In pseudo-code the CONV algorithm is presented in Algorithm 1.

Algorithm 1: The CONV algorithm for Bermudan options

$V(t_M, x) = E(t_M, x)$ for all x
For $m = M - 1$ to 0
  Dampen $V(t_{m+1}, x)$ with $\exp(\alpha x)$ and take its Fourier transform
  Calculate the right-hand side of (16)
  Calculate $C(t_m, x)$ by applying Fourier inversion to (16) and undamping
  $V(t_m, x) = \max\{(E(t_m, x), C(t_m, x)\}$
Next m

In Appendix A we demonstrate how the hedge parameters can be calculated in the CONV method. As differentiation is exact in Fourier space, they will be more stable than when calculated via finite-difference based approximations.

The following section deals with the implementation of the CONV algorithm. In particular we employ the power of the FFT to approximate the continuous Fourier transforms that are involved.

4. Implementation Details of the CONV Method. The very essence of the CONV method is the calculation of a convolution[3]:

$$c(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iux} \hat{v}(u) \phi\left(-(u - i\alpha)\right) du, \tag{18}$$

where $\hat{v}(u)$ is the Fourier transform of v:

$$\hat{v}(u) = \int_{-\infty}^{\infty} e^{iuy} v(y) dy. \tag{19}$$

In the remainder of this section we will just focus on equations (18) and (19) for notational ease. In order to be able to use the FFT for exponentially affine models means that we have to switch to logarithmic coordinates. For this reason the state variables x and y will represent $\ln S(t_m)$ and $\ln S(t_{m+1})$, up to a constant shift. This section is organised as follows. Section 4.1 deals with the discretisation of the convolution in (18) and (19). Section 4.2 analyses the error made by one step of the CONV method and provides guidelines to choosing the grids for u, x and y. Section 4.3 considers the choice of grid further and investigates how to deal with points of discontinuity. This will prove to be very important if we want to guarantee a smooth convergence of the algorithm. Finally, sections 4.4 and 4.5 deal with the pricing of Bermudan and American options with the CONV method.

4.1. Discretising the Convolution. We approximate both integrals in (18) and (19) by a discrete sum, so that the FFT algorithm can be employed for their computation. This necessitates the use of uniform grids for u, x and y:

$$u_j = u_0 + j\Delta u,\; x_j = x_0 + j\Delta x,\; y_j = y_0 + j\Delta y, \tag{20}$$

where $j = 0, \ldots, N-1$. Though they may be centered around a different point, the x- and y-grids have the same mesh size: $\Delta x = \Delta y$. Further, the Nyquist relation must be satisfied, i.e.,

$$\Delta u \cdot \Delta y = \frac{2\pi}{N}. \tag{21}$$

In principle we could use the Fractional FFT algorithm (FrFT) which does not require the Nyquist relation to be satisfied. Numerical tests indicated that the FrFT is on average 4 times slower than the FFT, and that we could obtain a

[3]For notational convenience we have dropped the discounting term out of the equation.

similar accuracy by quadrupling the number of points, so that we opted to use the FFT throughout. Details about the exact location of x_0 and y_0 will be given in Section 4.3, as will details about the range of all grids. Inserting (19) into (18), and approximating (19) with a general Newton-Côtes rule and (18) with the left-rectangle rule yields:

$$c(x_p) \approx \frac{\Delta u \Delta y}{2\pi} \sum_{j=0}^{N-1} e^{-iu_j x_p} \phi\left(-(u_j - i\alpha)\right) \sum_{n=0}^{N-1} w_n e^{iu_j y_n} v(y_n), \tag{22}$$

for $p = 0, \ldots, N-1$. When using the trapezoidal rule we choose the weights w_n as:

$$w_0 = \frac{1}{2}, \;\; w_{N-1} = \frac{1}{2}, \;\; w_n = 1 \;\text{ for } n = 1, \ldots, N-2. \tag{23}$$

Though it may seem that the choice for the left-rectangle rule in (18) would cause the leading error term in (22) to be $O(du)$, the error analysis will show that the Newton-Côtes rule one uses to approximate (19) is one of the main determinants hereof. Inserting the definitions of our grids into (22) yields:

$$c(x_p) \approx \frac{e^{-iu_0(x_0+p\Delta y)}}{2\pi} \Delta u \sum_{j=0}^{N-1} e^{-ijp2\pi/N} e^{ij(y_0-x_0)\Delta u} \phi\left(-(u_j - i\alpha)\right) \hat{v}(u_j), \tag{24}$$

where the Fourier transform of v is approximated by:

$$\hat{v}(u_j) \approx e^{iu_0 y_0} \Delta y \sum_{n=0}^{N-1} e^{ijn2\pi/N} e^{inu_0\Delta y} w_n v(y_n). \tag{25}$$

Let us now define the DFT and its inverse of a sequence $x_p,\ p = 0, \ldots, N-1$, as:

$$\mathcal{D}_j\{x_n\} := \sum_{n=0}^{N-1} e^{ijn2\pi/N} x_n, \;\; \mathcal{D}_n^{-1}\{x_j\} = \frac{1}{N} \sum_{j=0}^{N-1} e^{-ijn2\pi/N} x_j. \tag{26}$$

Though the reason why will become clear later, let us set $u_0 = -N/2\Delta u$. As $e^{inu_0\Delta y} = (-1)^n$ this finally leads us to write (24), (25) as:

$$c(x_p) \approx e^{iu_0(y_0-x_0)} (-1)^p \mathcal{D}_p^{-1}\{e^{ij(y_0-x_0)\Delta u} \phi\left(-(u_j - i\alpha)\right) \mathcal{D}_j\{(-1)^n w_n v(y_n)\}\}. \tag{27}$$

4.2. Error Analysis of the CONV Method. A first inspection of (27) suggests that error will arise from two sources:

- Discretisation of both integrals in (18) and (19);
- Truncation of these integrals.

One way to proceed is to consider both integrals in (18), (19) separately, and estimate both discretisation and truncation errors by applying the error analysis of [2]. [25] recently combined their analysis with sharp upper bounds on European plain vanilla option prices to arrive at a sharp error bound for the discretised Carr-Madan formula. Though it is possible to use parts of their analysis, we found that the resulting error bounds overestimate the true error of the discretised CONV formula. For this reason our analysis will be based on a Fourier series expansion of the damped continuation value $c(x)$. This is quite natural, as the Fourier transform itself is generalised from Fourier series of periodic functions by letting their period approach infinity. We depart from the risk-neutral valuation formula with damping and without discounting:

$$c(x) = \int_{-\infty}^{\infty} v(x+z)e^{-\alpha z} f(z)dz. \tag{28}$$

The first approximation we make is due to replacing v by its Fourier series expansion on $[-L/2, L/2]$, where we have fixed $L > 0$:

$$\widetilde{c}_1(x) \approx \int_{-\infty}^{\infty} \sum_{j=-\infty}^{\infty} v_j e^{ij2\pi(x+z)/L} e^{-\alpha z} f(z)dz = \sum_{j=-\infty}^{\infty} v_j e^{ij2\pi x/L} \phi\left(-(j\frac{2\pi}{L} - i\alpha)\right), \tag{29}$$

using the dominated convergence theorem. The Fourier series coefficients of v are given by:

$$v_j = \frac{1}{L}\int_{-L/2}^{L/2} v(y)e^{ij2\pi y/L}dy. \tag{30}$$

As the Fourier series expansion of v is a periodic function with period L, only agreeing with v on $[-L/2, L/2]$, the error from this approximation equals:

$$\begin{aligned} e_1(L) &= \widetilde{c}_1(x) - c(x) \\ &= \int_{\mathbb{R}\setminus[-L/2,L/2]} \left(v(x+z) - \sum_{j=-\infty}^{\infty} v_j e^{ij2\pi(x+z)/L} \right) e^{-\alpha z} f(z)dz. \end{aligned} \tag{31}$$

A general guideline for choosing L is to ensure that the mass of the density outside $[-L/2, L/2]$ is negligible. The function $\widetilde{c}_1$ can, at least on this interval, be interpreted as an approximate Fourier series expansion of $c(x)$.

The second error arises by truncating the infinite summation from $-N/2$ to $N/2-1$, leading to $\widetilde{c}_2$ and its associated error e_2:

$$\begin{aligned} \widetilde{c}_2 &= \sum_{j=-N/2}^{N/2-1} v_j e^{-ij2\pi x/L} \phi\left(-(j\frac{2\pi}{L} - i\alpha)\right), \\ |e_2(L,N)| = |\widetilde{c}_1(x) - \widetilde{c}_2(x)| &\le \sum_{|j|=N/2}^{\infty} |v_j||\phi\left(-(j\frac{2\pi}{L} - i\alpha)\right)|. \end{aligned} \tag{32}$$

To further bound this error we require knowledge about the rate of decay of Fourier coefficients. It is well known that even if v is only piecewise C^1 on $[-L/2, L/2]$ its Fourier series coefficients v_j tend to zero as $j \to \pm\infty$. The modulus of v_j can therefore be bounded as:

$$|v_j| \le \frac{\eta_1(L)}{|j|^{\beta_1}}. \tag{33}$$

By $\eta_i(\cdot)$ we denote a bounding constant. The quantities it depends on are between brackets. For functions that are piecewise continuous on $[-L/2, L/2]$ but whose L-periodic extension is discontinuous, we typically have $\beta_1 = 1$ as the following example demonstrates.

Example 4.2.1 (European Put). *Suppose that we have a European put payoff and that $y = \ln S(t) - \ln K$. Then the payoff function equals $v(y) = e^{\alpha y}K(1-e^y)^+$ and its Fourier series coefficients equal:*

$$v_j = K\left(e^{-L\alpha/2}(-1)^j \frac{e^{-L/2}-1}{L(\alpha+1)+2\pi ij} - L\frac{e^{-L\alpha/2}(-1)^j - 1}{(L(\alpha+1)+2\pi ij)(L\alpha+2\pi ij)}\right). \tag{34}$$

Clearly, $\beta_1 = 1$ in (33), though when $L \to \infty$ and $j2\pi/L \to u$ it can be shown that the Fourier series coefficient converges to the Fourier transform of the payoff function, which can be seen to be $O(u^{-2})$ from (5).

The characteristic function can always be assumed to have power decay:

$$|\phi(x+yi)| \leq \frac{\eta_2(y)}{|x|^{\beta_2}}. \tag{35}$$

This is overly conservative for e.g. the Black-Scholes model, where the characteristic function of the log-underlying $\phi(x+yi)$ decays as $\exp(-cx^2)$, or the Heston model where the characteristic function has exponential decay. For the most popular Lévy models however the power decay assumption is appropriate. The VG model for example has $\beta_2 = 2\tau/\nu$ with τ being the time step. Using (33) and (35) yields:

$$\begin{aligned} |e_2(L,N)| &\leq \sum_{|j|=N/2}^{\infty} \frac{\eta_1(L)}{|j|^{\beta_1}} \frac{\eta_2(\alpha)}{\left(\frac{2\pi}{L}\right)^{\beta_2} |j|^{\beta_2}} \leq \eta_3(\alpha,L) \int_{N/2-1}^{\infty} x^{-\beta_1-\beta_2} dx \\ &= \eta_3(\alpha,L) \frac{(N/2-1)^{1-\beta_1-\beta_2}}{\beta_1+\beta_2-1}, \end{aligned} \tag{36}$$

where $\eta_3(\alpha,L) = 2\eta_1(L)\eta_2(\alpha)(2\pi/L)^{-\beta_2}$. We finally arrive at the discretised CONV formula in (27) by approximating the Fourier series coefficients of v in (32) with a Newton-Côtes rule:

$$\widetilde{v}(u_j) = \frac{1}{L}\Delta y \sum_{n=0}^{N-1} w_n e^{iu_j y_n} v(y_n). \tag{37}$$

This is equal to the right-hand side of (24) multiplied by $1/L$. It becomes clear that we can set $\Delta y = L/N$ and $y_0 = -L/2$.

Inserting (37) in $\widetilde{c}_2$ results in the third and final approximation:

$$\widetilde{c}_3(x) = \sum_{j=-N/2}^{N/2-1} \widetilde{v}(u_j) e^{-ij2\pi/Lx} \phi\left(-(j\frac{2\pi}{L} - i\alpha)\right). \tag{38}$$

Assuming that the chosen Newton-Côtes rule is of $O(N^{-\beta_3})$, one can bound:

$$|v_j - \widetilde{v}(u_j)| \leq \frac{\eta_4(\alpha,L)}{N^{\beta_3}}, \tag{39}$$

leading to the following error estimate for $\beta_2 \neq 1$:

$$\begin{aligned} |e_3(L,N)| = |\widetilde{c}_2(x) - \widetilde{c}_3(x)| &\leq \frac{\eta_4(\alpha,L)}{N^{\beta_3}} \sum_{j=-N/2}^{N/2-1} |\phi\left(-(j\frac{2\pi}{L} - i\alpha)\right)| \\ &\leq \frac{\eta_4(\alpha,L)}{N^{\beta_3}} \left(3\phi(i\alpha) + 2\eta_2(\alpha)\left(\frac{2\pi}{L}\right)^{-\beta_2} \sum_{j=2}^{N/2} \frac{1}{|j|^{\beta_2}}\right) \\ &= \frac{\eta_5(\alpha,L)}{N^{\beta_3}} + \frac{\eta_6(\alpha,L)}{(1-\beta_2)N^{\beta_3}} \left(\frac{2^{\beta_2-1}}{N^{\beta_2-1}-1}\right). \end{aligned} \tag{40}$$

with $\eta_5(\alpha,L) = 3\eta_4(\alpha,L)\phi(i\alpha)$ and $\eta_6(\alpha,L) = 2\eta_2(\alpha)\eta_4(\alpha,L)(2\pi/L)^{-\beta_2}$. For $\beta_2 = 1$ the second error term should of course be $\eta_6(\alpha,L)\ln N/2/N^{\beta_3}$.

Summarising, if we use a Newton-Côtes rule to discretise the Fourier transform of the (continuous) function $v(y)$, the error in the discretised CONV formula can be bounded as:

$$\begin{aligned}|c(x)-\widetilde{c}_3(x)| &\le |c(x)-\widetilde{c}_1(x)|+|\widetilde{c}_1(x)-\widetilde{c}_2(x)|+|\widetilde{c}_2(x)-\widetilde{c}_3(x)| \\ &\le e_1(L)+e_2(L,N)+e_3(L,N) \\ &= e_1(L)+O(N^{-\min(\beta_3,\beta_2+\beta_3-1,\beta_1+\beta_2-1)}) \end{aligned} \tag{41}$$

As demonstrated, in most applications $\beta_1 = 1$. This implies that, aside from the truncation error, the order of convergence will be:

- $O(N^{-\beta_3})$ for characteristic functions decaying faster than a polynomial;
- $O(N^{\min(\beta_3,\beta_2+\beta_3-1,\beta_1+\beta_2-1)})$ for characteristic functions with power decay.

The magnitude of β_3 will depend on the interplay between the chosen Newton-Côtes rule and the nature of the payoff function. One final word should be mentioned on the damping coefficient α. In the continuous version of the algorithm in Section 3 α was chosen such that the damped continuation value was L^1-integrable. The direct construction of the discretised CONV formula in Section 4.2 via a Fourier series expansion of the continuation value replaces L^1-integrability on $(-\infty,\infty)$ with L^1-summability on $[-L/2, L/2]$, so that the restriction on α is removed. In principle any value of α is allowed as long as $\phi(i\alpha)$ is finite. Nevertheless it seems sensible to adhere to the guidelines stated before, as the function will resemble its continuous counterpart more and more as L increases. The impact of α on the accuracy of the CONV algorithm is investigated in Section 5.1.

This concludes the error analysis of one step of the CONV algorithm. It is easy to show that the error is not magnified further in the remaining time steps. The leading error of our algorithm is therefore dictated by the time step where the order of convergence in (41) is the smallest.

4.3. Dealing with Discontinuities. Our focus in this section lies on achieving smooth convergence for the CONV algorithm. As numerical experiments have shown that it is difficult to achieve smooth convergence with high order Newton-Côtes rules, we will from here on focus on the second order trapezoidal rule in (23). Smooth convergence is desirable as we will be using extrapolation techniques later on to price American options in Section 4.5.

The previous section analysed the error in the discretised CONV formula when we use a Newton-Côtes rule to integrate the function V, the maximum of the continuation value and the exercise value. If we focus on a simple Bermudan put it is clear that already at the last time step this function will have a discontinuous first derivative. Certainly it is also possible that V itself is discontinuous, think of contracts with a barrier clause. This will affect the order of convergence.

It is well-known that if we want to numerically integrate a function with (a finite number of) discontinuities, we should split up the integration domain such that we are only integrating continuous functions. Appendix B demonstrates this for the trapezoidal rule. In particular, we show that the trapezoidal rule remains second-order if only the first derivative of the integrand is discontinuous, at the cost of non-smooth convergence. If the integrand itself is discontinuous, the trapezoidal rule loses an order. Smooth second-order convergence can be restored by placing the discontinuities on the grid. This notion has often been utilised in lattice-based techniques, though the solutions have more often than not been payoff-specific. An approach that is more or less payoff-independent was recently proposed in [22], generalising previous work by [23], which essentially places discontinuities on the grid. Unfortunately, we cannot use their methodology here, as our desire to use the FFT binds us to a uniform grid.

Before investigating how to handle discontinuities in the CONV algorithm, we collect the results from the previous sections and restate the grid choice for the basic CONV algorithm. Equating the grids for x and y for now we have:

$$u_j = (j - \frac{n}{2})\Delta u,\ x_j = y_j = (j - \frac{1}{2})\Delta y, \quad j = 0, \ldots, N-1.$$

Here x and y represent, up to a constant shift, $\ln S(t_m)$ and $\ln S(t_{m+1})$, respectively. If in particular $x = \ln S(t_m) - \ln S(0)$ and $y = \ln S(t_{m+1}) - \ln S(0)$, so that x and y represent total log-returns, we will refer to this discretisation as Discretisation I. A convenient property of this discretisation is that the spot price always lies on the grid, so that no costly interpolation is required to back out the desired option value. Note that we need to ensure that the mass of the density of x and y outside $[-L/2, L/2]$ is negligible. Though more sophisticated approximations can be devised, we use a rule of thumb from [32] which chooses L as a multiple of the standard deviation of $\ln S(t_m)$, i.e.,

$$L = \delta \cdot \sqrt{- \left.\frac{\partial^2 \phi(t_m, u)}{\partial u^2}\right|_{u=0} + \left(\left.\frac{\partial \phi(t_m, u)}{\partial u}\right|_{u=0} \right)^2} \tag{42}$$

where $\phi(t_m, u)$ is the characteristic function of $\ln S(t_m)$ conditional upon $\ln S(0)$, and δ is a proportionality constant. Note that there is a trade-off in the choice of L: as we set $\Delta y = L/N$, the Nyquist relation implies $\Delta u = 2\pi/L$ and hence $[u_0, u_{N-1}] = [-N\pi/L, (N-2)\pi/L]$. Though larger values of L imply smaller truncation errors, they also cause the range of the grid in the Fourier domain to be smaller, so that the error in turn will be larger initially.

It is easy to come up with a choice of grid that allows us to place one discontinuity on the grid. Suppose that at time t_m the discontinuity we would like to place on the grid is d_m. We can then shift our grid by a small amount to arrive at:

$$x_j = \epsilon_x + (j - \frac{L}{2})\Delta y,\ y_j = \epsilon_y + (j - \frac{L}{2})\Delta y, \tag{43}$$

where $\epsilon_x = d_m - \lceil d_m/\Delta x \rfloor \cdot \Delta x$ and ϵ_y is chosen in a similar fashion. This discretisation will be referred to as Discretisation II. Even for plain vanilla European options where only one time step is required this is very useful. By choosing $\epsilon_y = \ln K/S(0)$ and $\epsilon_x = 0$ we ensure that the discontinuity of the call or put payoff lies on the y-grid, and the spot price lies on the x-grid. When more discontinuities are present it seems impossible to guarantee smooth convergence without abandoning the restriction of a uniform grid. In order to still be able to use the computational speed of the FFT we will then have to resort to e.g. the discontinuous FFT algorithm of [14] or a recent transform inversion technique in [24]. These directions are left for further research. Luckily, Discretisation II is well-suited for the pricing of Bermudan and American options, as we will show in the following sections.

4.4. Pricing Bermudan Options. As mentioned, when pricing Bermudan options the function V in (7) will have a discontinuous first derivative. Though at the final exercise time t_M the location of this discontinuity is known, this is not the case at previous exercise times. All we know after calculating V by equation (7) is that the discontinuity is contained in an interval of width Δx, say $[x_\ell, x_{\ell+1}]$.

If we proceed with the CONV algorithm without placing the discontinuity on the grid, the algorithm will display a non-smooth convergence. Andricopoulos et al. [1] overcome this problem by equating the exercise payoff and the continuation value, and solving numerically for the location of the discontinuity. In our framework this can be quite costly, so that we propose an effective alternative. We can use a simple linear interpolation to locate the discontinuity, say d_m:

$$d_m \approx \frac{x_{\ell+1}(C(t_m, x_\ell) - E(t_m, x_\ell)) - x_\ell(C(t_m, x_{\ell+1}) - E(t_m, x_{\ell+1}))}{(C(t_m, x_\ell) - E(t_m, x_\ell)) - (C(t_m, x_{\ell+1}) - E(t_m, x_{\ell+1}))}. \tag{44}$$

In the actual implementation we will use a cubic, instead of a linear interpolation. As in Discretisation II we can now shift the grid such that d_m lies on it, and recalculate both the continuation and the exercise value. In particular, note that the inner DFT of (27) does not have to be recalculated, the only term that is affected is the outer inverse DFT. As a by-product, calculating d_m automatically gives us an approximation of the exercise boundary.

It is demonstrated in Appendix B that if we opt for the trapezoidal rule a linear interpolation is sufficient to guarantee a smooth convergence. Obviously, if higher-order Newton-Côtes rules are used, higher order interpolation schemes will have to be employed to locate the discontinuity. The resulting algorithm we use to value Bermudan call or put options with a fixed strike K is presented below in pseudo-code.

Algorithm 2: Details of the algorithm for valuing Bermudan options.

```
Ensure that the strike K lies on the grid by setting ε_y = ln K/S(0)
For m = M − 1 to 1
    Equate the x-grid at t_m to the y-grid at t_{m+1}
    Compute C(t_m, x) through (27)
    Locate x_ℓ and x_{ℓ+1} and approximate d_m, e.g. via (44)
    Set ε_x = d_m and recompute C(t_m, x)
    Calculate V(t_m, x) = max (E(t_m, x), C(t_m, x))
    Set the y-grid at t_m to be equal to the x-grid at t_m
Next m
Set ε_x = 0 such that the initial spot price lies on the grid
Compute V(0, x) = C(0, x) using (27)
```

4.5. Pricing American Options. Within the CONV algorithm there are two ways to value an American option. One way is to approximate an American option by a Bermudan option with many exercise opportunities, the other is to use Richardson extrapolation on a series of Bermudan options with an increasing number of exercise opportunities. The method we use has been described in detail by Chang, Chung, and Stapleton [10], though the approach in finance dates back to Geske and Johnson [15]. The QUAD method in [1] also uses the same technique to price American options. We restrict ourselves to the essentials here. Let $V(\Delta t)$ be the price of a Bermudan option with a maturity of T years where the exercise dates are Δt years apart. It is assumed that $V(\Delta t)$ can be expanded as

$$V(\Delta t) = V(0) + \sum_{i=1}^{\infty} a_i(\Delta t)^{\gamma_i}, \tag{45}$$

with $0 < \gamma_i < \gamma_{i+1}$. $V(0)$ is the price of the American option. Classical extrapolation procedures assume that the exponents γ_i are known, which means that we can use $n+1$ Bermudan prices with varying Δt in order to eliminate n of the leading order terms in (45). The only paper considering an expansion of the Bermudan option price in terms of Δt we are aware of is of Howison [21], who shows that $\gamma_1 = 1$ for the Black-Scholes model. Nevertheless, numerical tests indicate that the assumption $\gamma_i = i$ produces satisfactory results for the Lévy models we consider.

5. Numerical Experiments. By various experiments we show the accuracy and speed of the CONV method. The method's flexibility is presented by showing results for three asset price processes, GBM, VG, and CGMY. In addition, we value a multi-asset option to give an impression of the CPU times required to value a basket option of moderate dimension. The pricing problems considered are of European, Bermudan and American style. We typically present the (positive or negative) error $V(0, S(0)) - V_{ref}(0, S(0))$, where the reference value $V_{ref}(0, S(0))$

is either obtained via another numerical scheme, or via the CONV algorithm with 2^{20} grid points. In the tables to follow we will also present the error convergence defined as the absolute value of the ratio between two consecutive errors. A factor of 4 then denotes second order convergence. All single-asset tests were performed in Matlab 7.0.1 on an Intel Xeon CPU 5160, 3.00GHz with 2GB RAM. The multi-asset calculations were done in C on a 64-bit machine, with a 1 GHz Bus frequency and 8GB RAM.

5.1. Characteristic Function for Lévy Price Processes. The CONV method, as outlined in Section 3, is particularly well-suited for exponential Lévy models whose characteristic functions are available in closed-form. We will briefly review some defining properties of these models before turning to the extended CGMY/KoBoL model (from hereon extended CGMY model) of [5] and [7] that will be used to access the performance of the CONV method. For more background we refer you to [11] for the usage of Lévy processes in a financial context and to [31] for a detailed analysis of Lévy processes in general.

In exponential Lévy models the asset price is modelled as an exponential function of a Lévy process $L(t)$:

$$S(t) = S_0 \exp(L(t)). \tag{46}$$

Though the CONV method can be adapted to cope with discrete dividend payments, for ease of exposure we assume the asset pays a continuous stream of dividends, measured by the dividend rate q. In addition, we assume the existence of a bank account $B(t)$ which evolves according to $dB(t) = rB(t)dt$, r being the risk-free rate. Recall that a process $L(t)$ on $(\Omega, \mathcal{J}, P)$, with $L(0) = 0$, is a Lévy process if:

1 it has independent increments;
2 it has stationary increments;
3 it is stochastically continuous, i.e., for any $t \geq 0$ and $\epsilon > 0$ we have

$$\lim_{s \to t} \mathbb{P}(|L(t) - L(s)| > \epsilon) = 0. \tag{47}$$

The first property (cf. (11)) is exactly the property we required to be able to recognise a cross-correlation in the risk-neutral valuation formula. Each Lévy process can be characterised by a triplet (μ, σ, ν) with $\mu \in \mathbb{R}, \sigma \geq 0$ and ν a measure satisfying $\nu(0) = 0$ and

$$\int_{\mathbb{R}} \min\,(1, |x|^2)\nu(dx) < \infty. \tag{48}$$

In terms of this triplet the characteristic function of the Lévy process equals:

$$\begin{aligned} \phi(u) &= \mathbb{E}[\exp\,(iuL(t))] \\ &= \exp\,(t(i\mu u - \frac{1}{2}\sigma^2 u^2 + \int_{\mathbb{R}} (e^{iux} - 1 - iux\mathbf{1}_{[|\mathbf{x}|<\mathbf{1}]}\nu(dx))), \end{aligned} \tag{49}$$

the celebrated Lévy-Khinchine formula. As is common in most models nowadays we assume that (46) is formulated directly under the risk-neutral measure. To ensure that the reinvested relative price $e^{qt}S(t)/B(t)$ is a martingale under the risk-neutral measure, we need to ensure that

$$\phi(-i) = \mathbb{E}[\exp\,(L(t))] = e^{(r-q)t}, \tag{50}$$

which is satisfied if we choose the drift μ as:

$$\mu = r - q - \frac{1}{2}\sigma^2 - \int_{\mathbb{R}} (e^x - 1 - x\mathbf{1}_{[|x|<1]})\nu(dx) \tag{51}$$

The motivation behind using more general Lévy processes than the Brownian motion with drift is the simple fact that the Black-Scholes model is not able to reproduce the volatility skew or smile present in most financial markets. Over the past few years it has been shown that several exponential Lévy models are, at least to some extent, able to reproduce the skew or smile. The particular model we will consider is the extended CGMY model. The underlying Lévy process is characterised by the triple $(\mu, \sigma, \nu_{CGMY})$, where the Lévy density is specified as:

$$\nu_{CGMY}(x) = \begin{cases} C\dfrac{\exp\left(-G|x|\right)}{|x|^{1+Y}} & \text{if} \quad x < 0 \\ C\dfrac{\exp\left(-M|x|\right)}{|x|^{1+Y}} & \text{if} \quad x > 0. \end{cases} \tag{52}$$

The parameters satisfy $C \geq 0$, $G \geq 0$, $M \geq 0$, and $Y < 2$. The condition $Y < 2$ is induced by the requirement that Lévy densities integrate x^2 in the neighbourhood of 0. Conveniently, the characteristic function of the log-asset price can be found in closed-form as:

$$\begin{aligned} &\phi(u) = \qquad\qquad (53) \\ &\quad S(0)^{iu} \exp\left(iu\mu t - \frac{1}{2}u^2\sigma^2 t + tC\Gamma(-Y)[(M-iu)^Y - M^Y + (G+iu)^Y - G^Y]\right), \end{aligned}$$

where $\Gamma(x)$ is the gamma function. One can verify that the parameters G and M represent respectively the smallest and largest finite moment in the model, as $\phi(-iu) = \mathbb{E}[S(t)^u]$ is infinite for $u < -G$ and for $u > M$. The model encompasses several models. When $\sigma = 0$ and $Y = 0$ we obtain the Variance Gamma (VG) model, which is often parameterised slightly differently with parameters [4] σ, θ and ν, related to C, G and M through:

$$C = \frac{1}{\nu},\ G = \frac{1}{\sqrt{\frac{1}{4}\theta^2\nu^2 + \frac{1}{2}\sigma^2\nu - \frac{1}{2}\theta\nu}},\ M = \frac{1}{\sqrt{\frac{1}{4}\theta^2\nu^2 + \frac{1}{2}\sigma^2\nu + \frac{1}{2}\theta\nu}}. \tag{54}$$

Finally, when $C = 0$ the model collapses to the Black-Scholes model.

To conclude this section, Table 1 contains a set of five parameter sets which will be used in various tests throughout this section. The only two parameters we have not specified yet are δ from (42), which determines the range of the grid, and the damping coefficient α. For all GBM tests we set $\delta = 20$; for the other Lévy models, which have fatter tails, we use $\delta = 40$.

Regarding the choice of α, Lord and Kahl [27] have demonstrated recently how to approximate the optimal damping coefficient when the payoff-transform is known, which increases the numerical stability of the Carr-Madan formula. This is particularly effective for in/out-of-the-money options and options with short maturities. Though their rationale can to some extent be carried over to the pricing of European plain vanilla options (the difference being that now the payoff-transform is also approximated numerically), the problem becomes much more opaque when dealing with Bermudan options. To see this, note that the continuation value of the Bermudan option at the penultimate exercise date equals that of a European option. In each grid point, the European option will have a different degree of moneyness, calling for a different value of α per grid point. The situation worsens as the number of exercise dates increases so that it is hard to say what the overall optimal value of α will be. What is evident from Figure 1, where we graph the error of the CONV algorithm as a function of α for a European and a Bermudan put

[4] The parameters σ and ν should not be confused with the volatility and Lévy density of the Lévy triplet.

under T2-VG, is that there is a relatively large range for which the error is stable. In all numerical experiments we will set $\alpha = 0$ which, at least for our examples, produces satisfactory results.

T1-GBM:	$S_0 = 100,$	$r = 0.1,$	$q = 0,$	$\sigma = 0.25;$
T2-VG:	$S_0 = 100,$ $\theta = -0.14,$	$r = 0.1,$ $\nu = 0.2;$	$q = 0,$	$\sigma = 0.12,$
T3-CGMY:	$S_0 = 1,$ $C = 1,$	$r = 0.1,$ $G = 5,$	$q = 0,$ $M = 5,$	$\sigma = 0,$ $Y = 0.5;$
T4-CGMY:	$S_0 = 90,$ $C = 0.42,$	$r = 0.06,$ $G = 4.37,$	$q = 0,$ $M = 191.2,$	$\sigma = 0$ $Y = 1.0102;$
T5-GBM:	$S_0 = 40,$ $\rho_{ij} = 0.25.$	$r = 0.06,$	$q = 0.04,$	$\sigma_i = 0.2,$

TABLE 1
Parameter sets in the numerical experiments

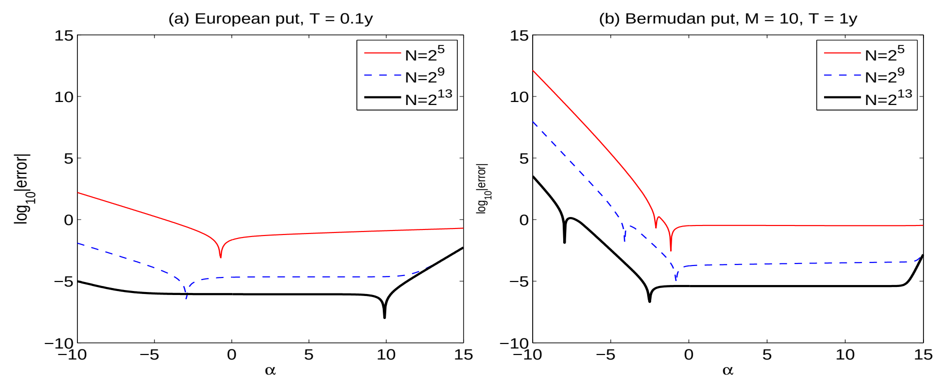

FIG. 1. *Error of CONV method under T2-VG and $K = 110$ for a European and Bermudan put in dependence of parameter α.*

5.2. European Call under GBM and VG. First of all, we evaluate the CONV method for pricing European options under VG. The parameters for the first test are from T2-VG with $T = 1$. Figure 2 shows that Discretisations I and II generate results of similar accuracy. What we notice from Figure 2 is that the only option with a stable convergence in Discretisation I is the at-the-money option with $K = 100$. It is clear that placing the strike on the y-grid in Discretisation II ensures a regular second order convergence. The results are obtained in comparable CPU time. From the error analysis in Section 4.2 it became clear that for short maturities in the VG model, the slow decay of the characteristic function ($\beta_3 = 2\tau/\nu$) might impair the second order convergence. To demonstrate this, we choose a call option with a maturity of 0.1 years, and $K = 90$. Table 2 presents the error of Discretisation II for this option in models T1-GBM and T2-VG. The convergence under GBM is clearly of a regular second order. From the error analysis we expect the convergence under VG to be of first order. Most probably the highly oscillatory integrand causes the non-smooth behaviour observed in Table 2. Note that all reference values are based on an adaptive integration of the Carr-Madan formula; all CPU times are determined after averaging the times of 1000 experiments.

In Appendix A the Greeks of the GBM call from Table 2 are computed.

5.3. Bermudan Option under GBM and VG. Turning to Bermudan options, we compare Discretisations I and II for 10-times exercisable Bermudan put

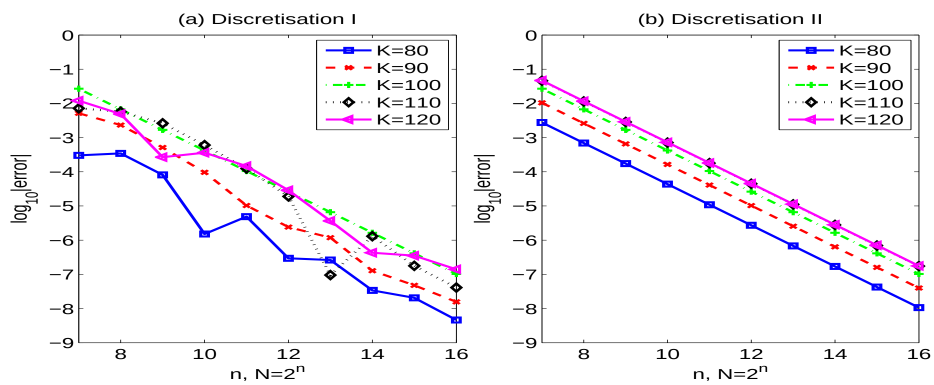

FIG. 2. *Convergence of the two discretisation methods for pricing European call options at various K under T2-VG; left: Discretisation I, right: Discretisation II.*

TABLE 2

CPU time, error and convergence rate for European call options under T1-GBM and T2-VG, $K = 90, T = 0.1$ *(using Discretisation II)*

$(N = 2^n)$	GBM: $V_{ref}(0, S_0) = 11.1352431$;			VG: $V_{ref}(0, S_0) = 10.9937032$;		
n	time(sec)	error	conv.	time(sec)	error	conv.
7	0.0001	-2.08e-3	–	0.0002	-2.92e-4	–
8	0.0002	-5.22e-4	4.0	0.0003	-1.42e-4	2.1
9	0.0003	-1.30e-4	4.0	0.0006	-4.61e-5	3.1
10	0.0006	-3.26e-5	4.0	0.0011	-9.48e-6	4.9
11	0.0012	-8.15e-6	4.0	0.0023	-8.41e-7	11.3
12	0.0023	-2.04e-6	4.0	0.0045	8.10e-7	1.0

options under both T1-GBM and T2-VG. The reference values reported in Table 3 and 4 are found by the CONV method with 2^{20} grid points.

It is shown in Tables 3 and 4 that both Discretisation I and II give results of similar accuracy. Discretisation I uses somewhat less CPU time, but Discretisation II shows a regular second order convergence, enabling the use of extrapolation. The computational speed of both discretisations is highly satisfactory.

TABLE 3

CPU time, error and convergence rate pricing a 10-times exercisable Bermudan put under T1-GBM; $K = 110, T = 1$ *and* $V_{ref}(0, S_0) = 11.98745352$,

$(N = 2^n)$	Discretisation I			Discretisation II		
n	time(sec)	error	conv.	time(sec)	error	conv.
7	0.0003	9.09e-3	-	0.006	-3.31e-2	-
8	0.0004	-1.29e-3	7.1	0.007	-8.53e-3	3.9
9	0.0007	1.80e-6	717.7	0.008	-2.13e-3	4.0
10	0.0019	2.71e-5	0.06	0.011	-5.33e-4	4.0
11	0.0041	-9.31e-6	2.9	0.018	-1.33e-4	4.0
12	0.0082	-1.31e-5	0.71	0.029	-3.33e-5	4.0

5.4. American Options under GBM, VG and CGMY. Because Discretisation II yields a regular convergence, we choose it in this section to price American options. We compare the accuracy and CPU time of the two approximation methods mentioned in Section 4.5, i.e. the direct approximation via a Bermudan option, and the repeated Richardson extrapolation technique. For the latter we opted for 2 extrapolations on 3 Bermudan options with 128, 64 and 32 exercise opportuni-

Table 4

CPU time, error and convergence rate pricing a 10-times exercisable Bermudan put under T2-VG; $K = 110, T = 1$ *with reference value* $V_{ref}(0, S_0) = 9.040646119$.

$(N = 2^n)$	Discretisation I			Discretisation II		
n	time(sec)	error	conv.	time(sec)	error	conv.
7	0.0003	8.45e-2	-	0.006	9.53e-2	-
8	0.0005	9.02e-3	9.4	0.007	1.09e-2	8.7
9	0.0010	-1.70e-4	53.1	0.008	2.50e-3	4.4
10	0.0019	-2.04e-4	0.8	0.010	6.51e-4	3.8
11	0.0047	-4.28e-5	4.8	0.016	1.65e-4	4.0
12	0.0094	-1.11e-5	3.9	0.026	4.15e-5	4.0

ties, which gave robust results. In our first test we price an American put under T1-GBM. The reference value was obtained by solving the Black-Scholes PDE on a very fine grid. The performance of both approximation methods is summarised in Table 5, where '$P(N/2)$' denotes that the American option is approximated by an N/2-times exercisable Bermudan option. 'Richardson' denotes the results obtained by the 2-times repeated Richardson extrapolation scheme. It is evident that the extrapolation-based method converges fastest and costs far less CPU time than the direct approximation approach (e.g. to reach an accuracy of 10^{-4}, the extrapolation method is approximately 20 times faster).

In Appendix A the Greeks of the American put from Table 5 are computed.

Table 5

CPU time and errors for an American put under T1-GBM, with: $K = 110, T = 1$, $V_{ref}(0, S(0)) = 12.169417$

$(N = 2^n)$	P(N/2)			Richardson		
n	time(sec)	error	conv.	time(sec)	error	conv.
7	0.025	-6.34e-2	–	0.011	-4.88e-2	–
8	0.055	-2.34-3	2.7	0.020	8.77e-3	5.6
9	0.130	-9.49e-3	2.5	0.038	2.24e-3	3.9
10	0.346	-4.19e-3	2.3	0.078	5.53e-4	4.1
11	1.18	-1.95e-3	2.1	0.181	1.29e-4	4.3
12	3.98	-9.40e-4	2.1	0.436	2.30e-5	5.6

In the remaining tests we demonstrate the ability of the CONV method to price American options accurately under alternative dynamics, using the VG and both CGMY test sets. All reported reference values were generated with the CONV method on a mesh with 2^{20} points and 2-times Richardson extrapolation on 512-, 256- and 128-times exercisable Bermudans. We have included one CGMY test with $Y < 1$, and one with $Y > 1$, as the latter is considered a hard test case when numerically solving the corresponding PIDE. Both CGMY tests stem from the PIDE literature, where reference values for the same American puts were reported as 0.112171 for T3-CGMY [4], and 9.2185249 for T4-CGMY [33]. Our reference result for the latter test differs in the second decimal with the result in [33]. For European options however, the CONV method converges to the exact analytical result, whereas the European reference given in [33] is not exact (the most accurate reported value for a call option with K = 98 and T = 0.25 is 2.2228514, whereas the analytical value of the option is 2.2306557). For this reason we present the CONV reference value. Though the convergence in Table 6 is less stable than for Bermudan options, the results in this section indicate that the CONV method is able to price American options under a variety of Lévy processes. A reasonable accuracy can be obtained quite quickly, so that it might be possible to calibrate a model to the prices of American options [5].

Table 6
CPU time and errors for American puts under VG and CGMY

	T2-VG		T3-CGMY		T4-CGMY	
	$K = 110, T = 1$		$K = 1, T = 1$		$K = 98, T = 0.25$	
$(N = 2^n)$	$V_{ref}(0, S(0)) = 10.0000$		$V_{ref}(0, S(0) = 0.112152$		$V_{ref}(0, S(0) = 9.225439$	
n	time(sec)	error	time(sec)	error	time(sec)	error
7	0.073	-3.49e-1	0.062	-6.35e-3	0.074	-1.89e-1
8	0.096	4.13e-2	0.097	1.38e-4	0.093	2.93e-2
9	0.115	1.37e-2	0.116	1.16e-4	0.118	-1.30e-3
10	0.157	-6.17e-3	0.160	1.10e-5	0.162	-3.97e-4
11	0.270	6.03e-3	0.275	1.18e-5	0.278	2.89e-4
12	0.466	1.31e-3	0.482	-2.35e-6	0.483	9.59e-5

5.5. 4D Basket Options under GBM. The CONV method can easily be generalised to higher dimensions. The only assumption that the multi-dimensional model is required to satisfy is the independent increments assumption in (11). We do not state the multi-dimensional version of Algorithm 1 here as it is a trivial generalisation of the univariate case. Its ability to price options of a moderate dimension is demonstrated by considering a 4-asset basket put option. Upon exercise at time t_i, the payoff is:

$$V(t_i, \mathbf{S}(\mathbf{t_i})) = \max(\frac{1}{4}\sum_{p=1}^{4} S_p(t_i) - K, 0). \tag{55}$$

The results of pricing a European and a 10-times exercisable Bermudan put under T5-GBM are summarised in Table 7. The CPU times on the tensor-product grids are very satisfactory, especially as the results on the coarse grids obtained in only a few seconds seem to have converged to within a practical tolerance level. In order to be able to price higher-dimensional problems our future research will aim to combine the multi-dimensional CONV method with sparse grids.

Table 7
CPU time and prices for multi-asset European and 10-times exercisable Bermudan basket put options under T5-GBM, $K = 40, T = 1$

European			10-times exerc. Bermudan	
N	result	time (sec)	result	time (sec)
16^4	1.6428	0.02	1.7721	0.18
32^4	1.6537	0.51	1.7390	3.40
64^4	1.6539	9.5	1.7394	65.7
128^4	1.6538	202.4	1.7393	1526.3

6. Conclusions. In this paper we have presented a novel FFT-based method for pricing options with early-exercise features, the CONV method. Like other FFT-based methods, it is flexible with respect to the choice of asset price process and the type of option contract, which has been demonstrated in numerical examples for European, Bermudan and American options. Path-dependent exotics can in principle also be valued by a forward propagation in time, though this has not

[5]The majority of exchange-traded options are American.

been demonstrated here. The crucial assumption of the method is that the underlying assets are driven by processes with independent increments, whose characteristic function is readily available. Though we have mainly focused on univariate exponential Lévy models, the techniques presented here certainly also extend to multivariate models, as Section 5.5 has shown. By using the FFT to calculate convolutions we achieve a complexity of $O(MNlogN)$, where N is the number of grid points and M is the number of exercise opportunities of the option contract. In comparison, the QUAD method of [1] is $O(MN^2)$. The speed of the method may make it possible to calibrate models to the prices of American options, as exchange-traded options are mainly of the American type. Future research will focus on the usage of more advanced quadrature rules, combined with speeding up the method for high-dimensional problems.

Acknowledgments: The authors would like to thank Coen Leentvaar for his help in producing numbers for Table 7.

Large parts of research for this paper were performed when the first author was employed by the Modelling and Research Department at Rabobank International and the Tinbergen Institute at the Erasmus University of Rotterdam. The authors are grateful to seminar participants at Rabobank International.

REFERENCES

[1] A.D. Andricopoulos, M. Widdicks, P.W. Duck and D.P. Newton, *Universal Option Valuation Using Quadrature*, J. Financial Economics, 67,3: 447-471, 2003,

[2] J. Abate and W. Whitt, The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems*, 10: 5–88, 1992.

[3] A. Almendral and C.W. Oosterlee, Highly Accurate Evaluation of European and American Options Under the Variance Gamma Process. *J. Comp. Finance* 10(1): 21-42, 2006.

[4] A. Almendral and C.W. Oosterlee, Accurate Evaluation of European and American Options Under the CGMY Process., *SIAM J. Sci. Comput.* 29: 93-117, 2007.

[5] S. I. Boyarchenko and S. Z. Levendorskiĭ, *Non-Gaussian Merton-Black-Scholes theory*, vol. 9 of Advanced Series on Statistical Science & Appl. Probability, World Scientific Publishing Co. Inc., River Edge, NJ, 2002.

[6] M. Broadie and Y. Yamamoto, *A double-exponential Fast Gauss transform algorithm for pricing discrete path-dependent options.* Operations Research 53(5): 764–779, 2005

[7] P. P. Carr, H. Geman, D. B. Madan, and M. Yor, *The fine structure of asset returns: An empirical investigation*, J. of Business, 75, 305–332, 2002.

[8] P. P. Carr and D. B. Madan, *Option valuation using the Fast Fourier Transform*, J. Comp. Finance, 2: 61–73, 1999.

[9] P. P. Carr, D. B. Madan, and E. C. Chang, *The Variance Gamma process and option pricing*, European Finance Review, 2: 79–105, 1998.

[10] C-C Chang, S-L Chung and R.C. Stapleton, *Richardson extrapolation technique for pricing American-style options* Proc. of 2001 Taiwanese Financial Association, Tamkang University Taipei, June 2001. Available at `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=313962`.

[11] R. Cont and P. Tankov, *Financial modelling with jump processes*, Chapman & Hall, Boca Raton, FL, 2004.

[12] H Dubner and J. Abate, *Numerical inversion of Laplace transforms by relating them to the finite Fourier cosine transform.* Journal of the ACM 15(1): 115–123, 1998.

[13] D. Duffie, J. Pan and K. Singleton, *Transform analysis and asset pricing for affine jump-diffusions.* Econometrica 68: 1343–1376, 2000.

[14] G. Fan and G.H. Liu, *Fast Fourier Transform for discontinuous functions,* IEEE Trans. Antennas and Propagation 52(2): 461-465, 2004.

[15] R. Geske, H. Johnson, *The American put valued analytically* J. of Finance 39: 1511-1542, 1984.

[16] J. Gil-Pelaez, *Note on the inverse theorem.* Biometrika 37: 481-482, 1951.

[17] J. Gurland, *Inversion formulae for the distribution of ratios.* Ann. of Math. Statistics 19: 228-237, 1948.

[18] S. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, Rev. Financ. Stud., 6: 327–343, 1993.

[19] D. J. Higham, *An Introduction to Financial Option Valuation*, Cambridge University Press, Cambridge, UK, 2004.

[20] A. Hirsa and D. B. Madan, *Pricing American Options Under Variance Gamma*, J. Comp. Finance, 7, 2004.
[21] S. Howison, *A matched asymptotic expansions approach to continuity corrections for discretely sampled options. Part 2: Bermudan options* Working Paper, Oxford Univ., 2005.
[22] Z. Hu, J. Kerkhof, P. McCloud, and J. Wackertapp, *Cutting edges using domain integration,* Risk, 19(11): 95-99, 2006.
[23] P. Hunt, J. Kennedy and A.A.J. Pelsser, *Markov-functional interest rate models.* Finance and Stochastics 4(4): 391-408, 2000.
[24] P. den Iseger, *Numerical transform inversion using Gaussian quadrature* Probab. in the Eng. and Inform. Sciences 20(1): 1-44, 2006.
[25] R. Lee, Option Pricing by Transform Methods: Extensions, Unification, and Error Control. *J. Computational Finance,* 7(3): 51-86, 2004.
[26] A. Lewis *A simple option formula for general jump-diffusion and other exponential Lévy processes.* SSRN working paper, 2001. Available at: `http//ssrn.com/abstract=282110`.
[27] R. Lord and C. Kahl, *Optimal Fourier inversion in semi-analytical option pricing.* SSRN working paper, 2006. Available at: `http//ssrn.com/abstract=921336`.
[28] A. M. Matache, P. A. Nitsche, and C. Schwab, *Wavelet Galerkin pricing of American options on Lévy driven assets*, working paper, ETH, Zürich, 2003.
[29] S. Raible *Lévy Processes in Finance: Theory, Numerics and Emperical Facts* PhD Thesis, Inst. für Math. Stochastik, Albert-Ludwigs-Univ. Freiburg, 2000.
[30] E. Reiner, *Convolution Methods for Path-Dependent Options,* Financial Math. workshop, IPAM UCLA, Jan. 2001. Available through `http://www.ipam.ucla.edu/publications/fm2001/fm2001_4272.pdf`.
[31] K-I Sato *Basic Results on Lévy Processes, In:* Lévy Processes, 3–37, Birkhäuser Boston, Boston MA, 2001.
[32] C. O'Sullivan, *Path Dependent Option Pricing under Levy Processes* EFA 2005 Moscow Meetings Paper, Available at SSRN: http://ssrn.com/abstract=673424, Febr. 2005.
[33] I. Wang, J.W. Wan and P. Forsyth, *Robust numerical valuation of European and American options under the CGMY process.* Techn. Report U. Waterloo, Canada, 2006.
[34] P. Wilmott, J. Dewynne, and S. Howison, *Option pricing*, Oxford: Financial Press, 1993.

Appendix A. The Hedge Parameters. Here, we present the CONV formulae for two important hedge parameters Δ and Γ, defined as,

$$\Delta = \frac{\partial V}{\partial S} \ = \frac{1}{S}\frac{\partial V}{\partial x}, \ \Gamma = \frac{\partial^2 V}{\partial S^2} = \frac{1}{S^2}\left(-\frac{\partial V}{\partial x} + \frac{\partial^2 V}{\partial x^2}\right). \tag{56}$$

As it is relatively easy to derive the corresponding CONV formulae, we merely present them here. For notational convenience we define:

$$\mathcal{F}\{e^{\alpha x}V(t_0,x)\} = e^{-r\Delta t}A(u), \tag{57}$$

where $A(u) = \mathcal{F}\{e^{\alpha y}V(t_1,y)\}\cdot\phi(-u+i\alpha)$, and we assume $t_1 > 0$. We now obtain the CONV formula for Δ, as

$$\Delta = \frac{e^{-\alpha x}e^{-r\Delta t}}{S}\Big[\mathcal{F}^{-1}\{-iuA(u)\} - \alpha\mathcal{F}^{-1}\{A(u)\}\Big], \tag{58}$$

and for Γ:

$$\Gamma = \frac{e^{-\alpha x}e^{-r\Delta t}}{S^2}\Big[\mathcal{F}^{-1}\{(-iu)^2A(u)\} - (1+2\alpha)\mathcal{F}^{-1}\{-iuA(u)\} + \alpha(\alpha+1)\mathcal{F}^{-1}\{A(u)\}\Big]. \tag{59}$$

Note that the only additional calculations occur at the final step of the CONV algorithm, where we calculate the value of the option given the continuation and exercise values at time t_1. Since differentiation is exact in Fourier space the rate of convergence of the Greeks will be the same as that of the value. To demonstrate this we evaluate the delta and gamma under T1-GBM of the European call from Table 2 and the American put from Table 5. For both tests we choose Discretisation II. Tables 8 and 9 present the results. The reference values for the European call option are analytic solutions, for the American call these were found by numerically solving the Black-Scholes PDE on a very fine grid. Note that the delta and gamma of the American put converge to a slightly different value - this is due to our approximation of the American option via 2 Richardson extrapolations on 128-, 64- and 32-times exercisable Bermudans. If we would increase the number of exercise opportunities of the Bermudan options the delta and gamma would, at the cost of a longer computation time, converge to their true values.

Table 8
Accuracy of hedge parameters for a European call under T1-GBM; $K = 110, T = 0.1$

$(N = 2^d)$	European call			
	$\Delta_{ref} = 0.933029$		$\Gamma_{ref} = 0.01641389$	
d	Δ error	conv.	Γ error	conv.
7	-3.75e-4	–	3.79e-5	–
8	-9.37e-5	4.0	9.43e-6	4.0
9	-2.34e-5	4.0	2.35e-6	4.0
10	-5.86e-6	4.0	5.88e-7	4.0
11	-1.46e-6	4.0	1.47e-7	4.0
12	-3.66e-7	4.0	3.68e-8	4.0

Table 9
Values of hedge parameters for an American put under T1-GBM; $K = 110, T = 0.1$

$(N = 2^d)$	American put:	
d	$\Delta_{ref} = -0.62052$	$\Gamma_{ref} = 0.0284400$
7	-0.62170	0.028498
8	-0.62035	0.028687
9	-0.62050	0.028464
10	-0.62053	0.028463
11	-0.62054	0.028463
12	-0.62055	0.028463

Appendix B. Error Analysis of the Trapezoidal Rule.

Suppose we are integrating $f \in C^\infty$ over an interval $[a, b]$. The discretisation error induced by approximating this integral with the trapezoidal rule follows from the Euler-Maclaurin summation formula:

$$\int_a^b f(x)dx - T(a, b, f, \Delta x) = \sum_{j=1}^{\infty} (\Delta x)^{2j} \frac{B_{2j}}{(2j)!} \left(f^{(2j-1)}(b) - f^{(2j-1)}(a) \right), \tag{60}$$

where B_j is the j-th Bernoulli number and $T(a, b, f, \Delta x)$ is the trapezoidal sum:

$$T(a, b, f, \Delta x) = \Delta x \{ \sum_{j=1}^{N-1} f(x_j) - \frac{1}{2}(f(a) + f(b)) \}, \tag{61}$$

with $\Delta x = (b - a)/(N - 1)$ and $x_j = a + j\Delta x$. From (60) it is clear that if the value of the first derivative is not the same in a and b, the trapezoidal rule is of order $1/N^2$.

The trapezoidal rule can obviously also be applied to functions that are piecewise continuously differentiable. The convergence may however be less stable if we do not know the exact location of the discontinuities. To see this, suppose that f can be written as:

$$f(x) = \begin{cases} g(x) & x \le z \\ h(x) & x > z \end{cases}. \tag{62}$$

Further, we define the following points:

$$\ell = \max\{j | x_j \leq z,\ j = 0, \ldots, N-1\}, \tag{63}$$

so that the interval $[x_\ell, x_{\ell+1}]$ contains z. Placing the discontinuity on the grid would result in the same order of convergence as the trapezoidal rule itself:

$$\int_a^b f(x)dx \approx T(a, x_\ell, g, \Delta x) + T(x_{\ell+1}, b, h, \Delta x) + \frac{1}{2}(z - x_\ell)(g(x_\ell) + g(z)) + \frac{1}{2}(x_{\ell+1} - z)(h(z) + h(x_{\ell+1})). \tag{64}$$

A straightforward application of the trapezoidal rule would lead to $T(a, b, f, \Delta x)$. The difference with (64) is:

$$\frac{1}{2}\Delta x g(x_\ell) + \frac{1}{2}\Delta x h(x_{\ell+1}) - \frac{1}{2}(z - x_\ell)(g(x_\ell) + g(z)) - \frac{1}{2}(x_{\ell+1} - z)(h(z) + h(x_{\ell+1})).$$

Expanding both g and h around the point of discontinuity z yields:

$$\frac{1}{2}(x_{\ell+1} + x_\ell - 2z)(g(z) - h(z)) + \frac{1}{2}(x_{\ell+1} - z)(z - x_\ell)(g^{(1)}(z) - h^{(1)}(z)) + \frac{1}{2}(x_{\ell+1} - z)\sum_{j=1}^{\infty}\frac{1}{j!}g^{(j)}(z) + \frac{1}{2}(z - x_\ell)^j h^{(j)}(z).$$

If f is continuous, but the first derivatives of g and h do not match at z, the order of convergence is still $1/N^2$ since $(x_{\ell+1} - z)(z - x_\ell) \leq (\Delta x)^2$. It is clear that as N changes, the ratio of $(x_{\ell+1} - z)(z - x_\ell)$ to $(\Delta x)^2$ may vary strongly, leading to non-smooth convergence. If f is discontinuous, i.e., if the values of g and h in z disagree, the order of convergence is of $O(1/N)$.

Now suppose that we have computed g and h at grid points $x_j,\ j = 0, \ldots, N-1$. We know that $g(z) = h(z)$, though we do not know the exact location of z. All we know is that it is contained in $[x_\ell, x_{\ell+1}]$. This is a situation we encounter in the pricing of Bermudan options, as outlined in Section 4.4. If we proceed to integrate f on this grid, we will not obtain smooth convergence. A simple approximation of the discontinuity can however be found by assuming a linear relationship between x and $g(x) - h(x)$. This leads to

$$z \approx \frac{x_{\ell+1}(g(x_\ell - h(x_\ell)) - x_\ell(g(x_{\ell+1}) - h(x_{\ell+1}))}{(g(x_\ell - h(x_\ell)) - (g(x_{\ell+1}) - h(x_{\ell+1}))} + O\left(\Delta x^2\right), \tag{65}$$

where the error estimate follows from linear interpolation. Now suppose that we recalculate g and h such that either X_ℓ or $x_{\ell+1}$ coincide with this approximation of z, and redo the numerical integration. It is easy to see that smooth convergence will be restored, as the combination of the error term in(65) to the error term in (64) will be of $O\left(\Delta x^3\right)$. Note that if we use higher-order Newton-Côtes rules, a higher order interpolation step will be required.