

Hilbert transform, spectral filtering and option pricing

Carolyn E. Phelan

Financial Computing and Analytics Group, Department of Computer Science, University College London
c.phelan@cs.ucl.ac.uk

Daniele Marazzina

Dipartimento di Matematica, Politecnico di Milano
daniele.marazzina@polimi.it

Gianluca Fusai

Dipartimento di Studi per l'Economia e l'Impresa (DiSEI), Università del Piemonte Orientale "Amedeo Avogadro", Novara
Department of Finance, Cass Business School, City University London
gianluca.fusai@unipmn.it, gianluca.fusai.1@city.ac.uk

Guido Germano

Financial Computing and Analytics Group, Department of Computer Science, University College London
Systemic Risk Centre, London School of Economics
g.germano@ucl.ac.uk, g.germano@lse.ac.uk

We show how spectral filtering techniques can improve the convergence of numerical schemes which use discrete Hilbert transforms based on a sinc function expansion, and thus ultimately on the fast Fourier transform. This is relevant, for example, for the computation of fluctuation identities, which give the distribution of the maximum or the minimum of a random path, or the joint distribution at maturity with the extrema staying below or above barriers. We use as examples the methods by Feng and Linetsky (2008) and Fusai, Germano and Marazzina (2016) to price discretely monitored barrier options where the underlying asset price is modelled by a Lévy process. Both methods show exponential convergence with respect to the number of grid points in most cases, but are limited to polynomial convergence under certain conditions. We relate these rates of convergence to the Gibbs phenomenon for Fourier transforms and achieve improved results with spectral filtering.

Key words: double-barrier options, discrete monitoring, Lévy processes, Spitzer identity, Wiener-Hopf factorisation, Hilbert transform, Fourier transform, FFT, z -transform, sinc function, Gibbs phenomenon, spectral filtering

1. Introduction

Derivative pricing with Fourier transforms was first investigated by Heston (1993). Carr and Madan (1999) published the first method with both the characteristic function and the payoff in the Fourier domain. Fang and Oosterlee (2008, 2009) devised the COS method based on the Fourier-cosine expansion. The Hilbert transform (King 2009) has also been successfully employed: by Feng and Linetsky (2008) to price barrier options using backward induction in the Fourier space and by Marazzina et al. (2012) and Fusai et al. (2016) to compute the factorisations required by the

Spitzer identities (Spitzer 1956, Kemperman 1963) via the Plemelj-Sokhotsky relations. Feng and Linetsky showed that computing the Hilbert transform with the sinc expansion, as studied by Stenger (1993, 2011), gives errors that reduce exponentially as the number of fast Fourier transform (FFT) grid points increases. Pricing derivatives, especially exotic options, is a challenging problem in the operations research literature. Fusai et al. (2016) provide extensive references for this, as well as for many non-financial applications of the Hilbert transform and the related topics of Wiener-Hopf factorisation and Spitzer identities in insurance, queuing theory, physics, engineering, applied mathematics, etc.

When working in the Fourier domain one has the advantage that there are many pre-existing robust implementations of the discrete Fourier transform (DFT), e.g. the FFTW library by Frigo and Johnson (1998). However, the use of such numerical solutions means that one must manage issues arising from the approximation of an integral over an infinite domain with a finite sum. As long as the truncation limits in the log-price domain are selected judiciously, the main issue to contend with is the so-called Gibbs (or Gibbs-Wilbraham) phenomenon (Wilbraham 1848, Gibbs 1898, 1899). This is commonly observed as oscillations either side of a discontinuity in the original domain and is caused by truncating the function in the Fourier domain. Importantly for the pricing methods considered in this article, it also describes how the shape of the function in the Fourier domain relates to the order of the discontinuities in the original domain.

There have been many papers exploring possible solutions to the Gibbs phenomenon in a general setting, notably by Hewitt and Hewitt (1979), Vandeven (1991), Gottlieb and Shu (1997), Tadmor and Tanner (2005) and Tadmor (2007). More recently Ruijter et al. (2015) explored the use of spectral filtering techniques to solve the problem of slow polynomial error convergence seen when the COS method is used with non-smooth probability distributions.

The application investigated in this paper is the improvement of methods based on the fast Hilbert transform for the pricing of discretely monitored barrier options with Lévy processes. Recent papers have described significant progress in the valuation of these types of financial contracts. Feng and Linetsky (2008) devised a method which gives exponential convergence for many Lévy processes but is limited to polynomial convergence for the variance gamma (VG) processes; this method has a computational time which increases linearly with the number of dates N . Fusai et al. (2016) used the Spitzer identities to devise a method whose computation time is independent of N and which achieves exponential convergence for single-barrier and lookback options, again with the exception of the VG process, but which is limited to polynomial convergence for double-barrier options.

This paper explains the origin of the error performance and presents modified versions of the Fusai, Germano and Marazzina (FGM) and Feng and Linetsky (FL) methods with improved convergence. In order to do this we make use of, and extend, the investigation into the error performance

of the discrete Hilbert transform by Stenger (1993, 2011) and Feng and Linetsky (2008). We show that the error performance is related to both the shape of the characteristic function of the underlying process and to the Gibbs phenomenon. Finally, by making use of the filtering techniques suggested by Gottlieb and Shu (1997) and McKechnan et al. (2010) we are able to achieve improved convergence. The methods we compare are FGM, filtered FGM (FGM-F), Feng and Linetsky (FL) and filtered Feng and Linetsky (FL-F). They are compared for single and double-barrier options for the Kou, normal inverse Gaussian (NIG) and VG processes.

The structure of this paper is as follows. In Section 2 we run briefly through Fourier, Hilbert and z -transforms, give a concise overview of the original pricing schemes and explain our modifications to improve convergence. Section 3 includes a discussion of the performances of the pricing techniques and how they relate to the Gibbs phenomenon and the shape of the characteristic function of the underlying processes. Lastly, Section 4 shows numerical results, comparing the filtered algorithms with the original FGM and FL methods.

2. Background

As this method directly extends the FGM (Fusai et al. 2016) and FL (Feng and Linetsky 2008) pricing methods, we refer to the original papers for a comprehensive introduction. Aspects of the methods which are directly relevant to the error investigation are described here in order to provide a background to the changes that were made to improve convergence.

2.1. Fourier and Hilbert transforms

In this paper we make extensive use of the Fourier transform (see e.g. Polyanin and Manzhirov 1998, Kreyszig 2011), an integral transform with many applications. Historically, it has been widely used in spectroscopy and communications, therefore much of the literature refers to the function in the Fourier domain as its spectrum. According to the usual convention in finance literature, the forward and inverse Fourier transforms are defined as

$$\widehat{f}(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x)] = \int_{-\infty}^{+\infty} f(x) e^{i\xi x} dx, \quad (1)$$

$$f(x) = \mathcal{F}_{\xi \rightarrow x}^{-1} [\widehat{f}(\xi)] = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \widehat{f}(\xi) e^{-i\xi x} d\xi. \quad (2)$$

Let S_t be the price of an underlying asset and $x_t = \log(S_t/S_0)$ its log-price. To find the price $v(x_t, t)$ of an option at time $t = 0$ when the initial price of the underlying is S_0 and thus its log-price is $x_0 = 0$, we need to discount the expected value of the undamped option payoff $\phi(x_T) e^{-\alpha x_T}$ at maturity $t = T$ with respect to an appropriate risk-neutral probability distribution function (PDF)

$p(x, T)$ whose initial condition is $p(x, 0) = \delta(x)$. As shown by Lewis (2001), this can be done using the Plancherel relation,

$$\begin{aligned} v(0, 0) &= e^{-rT} \mathbf{E} [\phi(x_T) e^{-\alpha x_T} | x_0 = 0] = e^{-rT} \int_{-\infty}^{+\infty} \phi(x) e^{-\alpha x} p(x, T) dx \\ &= \frac{e^{-rT}}{2\pi} \int_{-\infty}^{+\infty} \widehat{\phi}(\xi) \widehat{p}^*(\xi + i\alpha, T) d\xi = e^{-rT} \mathcal{F}_{\xi \rightarrow x}^{-1} [\widehat{\phi}(\xi) \widehat{p}^*(\xi + i\alpha, T)] (0). \end{aligned} \quad (3)$$

Here, $\widehat{p}^*(\xi + i\alpha, T)$ is the complex conjugate of the Fourier transform of $e^{-\alpha x} p(x, T)$. To price options using this relation, we need the Fourier transforms of both the damped payoff and the PDF. A double-barrier option has the damped payoff

$$\phi(x) = e^{\alpha x} S_0 (\theta(e^x - e^k))^+ \mathbf{1}_{[l, u]}(x), \quad (4)$$

where α is the damping factor, $\theta = 1$ for a call, $\theta = -1$ for a put, $\mathbf{1}_A(x)$ is the indicator function of the set A , $k = \log(K/S_0)$ is the log-strike, $u = \log(U/S_0)$ is the upper log-barrier, $l = \log(L/S_0)$ is the lower log-barrier, K is the strike price, U is the upper barrier and L is the lower barrier. The Fourier transform of the damped payoff $\phi(x)$ is available analytically,

$$\widehat{\phi}(\xi) = S_0 \left(\frac{e^{(1+i\xi+\alpha)a} - e^{(1+i\xi+\alpha)b}}{1+i\xi+\alpha} - \frac{e^{k+(i\xi+\alpha)a} - e^{k+(i\xi+\alpha)b}}{i\xi+\alpha} \right), \quad (5)$$

where for a call option $a = u$ and $b = \max(k, l)$, while for a put option $a = l$ and $b = \min(k, u)$.

The Fourier transform of the PDF $p(x, t)$ of a stochastic process $X(t)$ is the characteristic function

$$\Psi(\xi, t) = \mathbf{E} [e^{i\xi X(t)}] = \int_{-\infty}^{+\infty} p(x, t) e^{i\xi x} dx = \mathcal{F}_{x \rightarrow \xi} [p(x, t)] = \widehat{p}(\xi, t). \quad (6)$$

For a Lévy process the characteristic function can be written as $\Psi(\xi, t) = e^{\psi(\xi)t}$, where the characteristic exponent $\psi(\xi)$ is given by the Lévy-Khincine formula as

$$\psi(\xi) = ia\xi - \frac{1}{2}\sigma^2\xi^2 + \int_{\mathbb{R}} (e^{i\xi\eta} - 1 - i\xi\eta \mathbf{1}_{[1,1]}(\eta)) \nu(d\eta). \quad (7)$$

The Lévy-Khincine triplet (a, σ, ν) uniquely defines the Lévy process. The value of a defines the linear drift of the process, σ is the volatility of the diffusion part of the process, and the jump part of the process is specified so that $\nu(\eta)$ is the intensity of a Poisson process with jump size η . Under the risk-neutral measure the parameters of the triplet are linked by the equation

$$a = r - q - \frac{1}{2}\sigma^2 - \int_{\mathbb{R}} (e^\eta - 1 - i\eta \mathbf{1}_{[1,1]}(\eta)) \nu(d\eta), \quad (8)$$

where r is the risk-free interest rate and q is the dividend rate. In general the characteristic function of a Lévy process is available in closed form, for example for the Gaussian (Schoutens 2003), NIG (Barndorff-Nielsen 1998), CGMY (Carr et al. 2002), Kou double exponential (Kou 2002), Merton

jump diffusion (Merton 1976), Lévy alpha stable (Nolan 2017), VG (Madan and Seneta 1990) and Meixner (Schoutens 2003) processes.

Some pricing techniques based on the Fourier transform, e.g. FGM and FL, also use the Hilbert transform, which is an integral transform related to the Fourier transform. However, in contrast to the Fourier transform, the function under transformation remains in the same domain, rather than moving between the x and ξ domains. The Hilbert transform of a function in the Fourier domain is defined as

$$\begin{aligned}\mathcal{H}[\widehat{f}(\xi)] &= \text{P.V.} \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{1}{\pi} \left(\int_{\xi-1/\epsilon}^{\xi-\epsilon} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' + \int_{\xi+\epsilon}^{\xi+1/\epsilon} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' \right),\end{aligned}\quad (9)$$

where P.V. denotes the Cauchy principal value. Applying the Hilbert transform in the Fourier domain is equivalent to multiplying the function in the x domain by $-i \operatorname{sgn} x$.

2.2. Applying barriers with Hilbert transforms

The Hilbert transform can be used to obtain the Fourier transform of the part of a function above or below a barrier, without leaving the Fourier domain. For example, with a barrier at 0, the functions $\widehat{f}_+(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x) \mathbf{1}_{\mathbb{R}_+}(x)]$ and $\widehat{f}_-(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x) \mathbf{1}_{\mathbb{R}_-}(x)]$, can be calculated using the Plemelj-Sokhotsky relations,

$$\widehat{f}_+(\xi) = \frac{1}{2} \{ \widehat{f}(\xi) + i\mathcal{H}[\widehat{f}(\xi)] \} \quad (10)$$

$$\widehat{f}_-(\xi) = \frac{1}{2} \{ \widehat{f}(\xi) - i\mathcal{H}[\widehat{f}(\xi)] \}. \quad (11)$$

The shift theorem $\mathcal{F}_{x \rightarrow \xi} [f(x+b)] = \widehat{f}(\xi) e^{-ib\xi}$ allows the Plemelj-Sokhotsky relations to be generalised to an arbitrary barrier b ,

$$\widehat{f}_{b+}(\xi) = \frac{1}{2} \{ \widehat{f}(\xi) + e^{ib\xi} i\mathcal{H}[e^{-ib\xi} \widehat{f}(\xi)] \} \quad (12)$$

$$\widehat{f}_{b-}(\xi) = \frac{1}{2} \{ \widehat{f}(\xi) - e^{ib\xi} i\mathcal{H}[e^{-ib\xi} \widehat{f}(\xi)] \}. \quad (13)$$

Eqs. (12) and (13) can be combined to obtain the Fourier transform the part of a function between two barriers, i.e. $\widehat{f}_{lu}(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x) \mathbf{1}_{[l,u]}(x)]$,

$$\widehat{f}_{lu}(\xi) = \frac{1}{2} \{ e^{il\xi} i\mathcal{H}[e^{-il\xi} \widehat{f}(\xi)] - e^{iu\xi} i\mathcal{H}[e^{-iu\xi} \widehat{f}(\xi)] \}. \quad (14)$$

The Hilbert transform was used by Feng and Linetsky (2008) to price discrete barrier options exploiting the relationship between the price at two successive monitoring dates:

$$v(x, t_{n-1}) = \int_l^u v(x', t_n) p(x - x', \Delta t) dx'. \quad (15)$$

Here $v(x, t_N) = \phi(x)e^{-\alpha x}$, i.e. the payoff of the option, and $p(\cdot, \Delta t)$ denotes the transition density of the underlying process with step size Δt . Using the convolution theorem together with the Hilbert transform, Eqs. (12)–(14) can be employed to express the relationship between the price at two successive dates as

$$\widehat{v}(\xi, t_{n-1}) = \frac{1}{2} \left\{ \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) + e^{i\ell\xi} i\mathcal{H} \left[e^{-i\ell\xi} \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) \right] \right\} \quad (16)$$

for a single-barrier down-and-out option and

$$\widehat{v}(\xi, t_{n-1}) = \frac{1}{2} \left\{ e^{i\ell\xi} i\mathcal{H} \left[e^{-i\ell\xi} \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) \right] - e^{iu\xi} i\mathcal{H} \left[e^{-iu\xi} \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) \right] \right\} \quad (17)$$

for a double-barrier option.

2.3. Spitzer identities

If we wish to use Eq. (3) to price barrier options, the required characteristic functions are more complicated than the closed-form expressions referred to in Section 2.1. For our example we require the characteristic function of the distribution of the value of a process $X(t)$ at time $t = T$, conditional on the process remaining inside upper and lower barriers at discrete monitoring dates t_n , $n = 0, 1, \dots, N$. Fortunately, for a single barrier we can use the identities by Spitzer (1956), and for double barriers their extension by Kemperman (1963). These provide the Fourier- z transform of the required PDF: the Fourier transform is applied to the log-price x and the z -transform is applied to the discrete monitoring times. The z -transform of a discrete function $f(t_n)$ with $n \in \mathbb{N}_0$ is defined as

$$\widetilde{f}(q) = \sum_{n=0}^{\infty} f(t_n) q^n, \quad q \in \mathbb{C}. \quad (18)$$

An important aspect in the calculation of the Spitzer identities is the decomposition of a function into + and - parts, $\widehat{f}_+(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x) \mathbf{1}_{\mathbb{R}_+}(x)]$ and $\widehat{f}_-(\xi) = \mathcal{F}_{x \rightarrow \xi} [f(x) \mathbf{1}_{\mathbb{R}_-}]$. As explained in Section 2.2, this can be done directly in the Fourier domain using the Plemelj-Sokhotsky relations, Eqs. (10) and (11). The calculation of the Spitzer identities also requires the factorisation a function, i.e. obtain $\widehat{g}_+(\xi)$ and $\widehat{g}_-(\xi)$ such that $\widehat{g}(\xi) = \widehat{g}_+(\xi) \widehat{g}_-(\xi)$. This is achieved through a log-decomposition, i.e. decomposing the logarithm $\widehat{h}(\xi) = \log \widehat{g}(\xi)$ and then exponentiating the results to obtain $\widehat{g}_+(\xi) = \exp \widehat{h}_+(\xi)$ and $\widehat{g}_-(\xi) = \exp \widehat{h}_-(\xi)$. Fusai et al. (2016) and Green et al. (2010) go into detail regarding the Spitzer identities and describe methods for single-barrier, double-barrier and lookback options. In this paper we concentrate on the Spitzer identities used for double-barrier and single-barrier down-and-out options. The first step is to factorise $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t) = \Phi_+(\xi, q) \Phi_-(\xi, q)$. In the case of a single-barrier down-and-out option, the z -transform of the required characteristic function is given by

$$\widetilde{p} = \frac{P_+(\xi, q)}{\Phi_+(\xi, q)} \quad (19)$$

where $P_+(\xi, q)$ is obtained from the decomposition of $P(\xi, q) = e^{i\xi} / \Phi_-(\xi, q)$. For a double-barrier option, the z -transform of the required characteristic function is

$$\tilde{p}(\xi, q) = \frac{1}{\Phi(\xi, q)} - e^{i\xi} \frac{J_-(\xi, q)}{\Phi(\xi, q)} - e^{iu\xi} \frac{J_+(\xi, q)}{\Phi(\xi, q)}, \quad (20)$$

where $J_+(\xi, q)$ and $J_-(\xi, q)$ are the solution to the coupled equations

$$\frac{J_-(\xi, q)}{\Phi_-(\xi, q)} = \left[\frac{e^{-i\xi} - e^{i(u-l)\xi} J_+(\xi, q)}{\Phi_-(\xi, q)} \right]_- \quad (21)$$

$$\frac{J_+(\xi, q)}{\Phi_+(\xi, q)} = \left[\frac{e^{-iu\xi} - e^{i(l-u)\xi} J_-(\xi, q)}{\Phi_+(\xi, q)} \right]_+. \quad (22)$$

The Spitzer identity for a single-barrier option can be solved directly. However so far only an iterative solution has been found to the coupled Eqs. (21) and (22) (Fusai et al. 2016).

2.4. Numerical methods

The methods in the previous section are described analytically. However, as they involve some expressions which cannot be solved in closed form, their implementation requires the use of numerical approximation techniques which we discuss in the following.

2.4.1. Discrete Fourier transform and spectral filtering The forward and reverse Fourier transforms in Eqs. (1) and (2) are integrals over an infinite domain and in order to compute them numerically one needs to approximate them with a discrete Fourier transform (DFT). Rather than being defined over an infinite and continuous range of x and ξ values, the DFT is defined on grids of size M in the x and ξ domains. For our scheme both the x and ξ grids are centred around zero and are defined based on the maximum value in the x domain x_{\max} . The step size is $\Delta x = 2x_{\max}/M$ and the x domain grid is defined as

$$x_j = j\Delta x, \quad j = -\frac{M}{2}, -\frac{M}{2} + 1, \dots, \frac{M}{2} - 1. \quad (23)$$

The points in the ξ domain are then calculated according to the Nyquist relation by obtaining the step size $\Delta\xi = \pi/x_{\max}$ and range $\xi_{\max} = \pi/\Delta x$ to give the ξ domain grid as

$$\xi_k = k\Delta\xi, \quad k = -\frac{M}{2}, -\frac{M}{2} + 1, \dots, \frac{M}{2} - 1. \quad (24)$$

The discrete Fourier transform is then

$$\hat{f}_{M, \Delta x}(\xi_k) = \Delta x \sum_{j=-M/2}^{M/2-1} f(x_j) e^{ix_j \xi_k} \quad (25)$$

$$f_{M, \Delta \xi}(x_j) = \frac{\Delta \xi}{2\pi} \sum_{k=-M/2}^{M/2-1} \hat{f}(\xi_k) e^{-ix_j \xi_k}. \quad (26)$$




Figure 1 Illustration of the effect of the Gibbs phenomenon on a rectangular pulse recovered applying the inverse FFT with grid size M to $\text{sinc}(\xi/2\pi)$. The approximated function is shown on the left and the error with respect to an analytical rectangular pulse on the right. On increasing M , the peak error at the discontinuity remains the same, the error away from the discontinuity reduces and the frequency of the oscillations increases.

In practice, we perform this calculation using the built-in MATLAB FFT function based on the FFTW library by Frigo and Johnson (1998).

It can be seen in Eqs. (25) and (26) that the range over which we calculate the Fourier transform is truncated, so we must consider the effect of the Gibbs phenomenon on the error performance. The Gibbs phenomenon describes the way that the shape of the function $f_{M, \Delta \xi}(x)$ approximated by a truncated Fourier series, i.e. the finite sum in Eq. (26), converges to the analytical function $f(x)$ corresponding to an infinite sum. Hewitt and Hewitt (1979) provided a comprehensive guide to this effect which was first observed by Wilbraham (1848) and later described by Gibbs (1898, 1899). An example of this can be seen in Figure 1 which shows how $f_{M, \Delta \xi}(x)$ for a rectangular pulse varies as the value of M increases. The error peaks at the discontinuity $f(x_d)$ and oscillates away from it, with the amplitude decreasing as a function of distance from the discontinuity. The value of the recovered function at the discontinuity $f_{M, \Delta \xi}(x_d)$ will be the mean of the values immediately before and after the discontinuity, i.e. $f_{M, \Delta \xi}(x_d) = \frac{1}{2}[f(x_d^+) + f(x_d^-)]$, and thus stays the same even as the value of M increases. In contrast, it can be observed from Figure 1 that the oscillations increase in frequency and decrease in amplitude as the value of M increases.

An important aspect of the Gibbs phenomenon is that, even for values of x far away from a discontinuity, the speed of convergence of the recovered function is altered by the presence of the discontinuity. If $f(x) \in C^\infty$, $x \in \mathbb{R}$, the discrete Fourier transform converges exponentially, i.e. $\max_j |f(x_j) - f_{M, \Delta \xi}(x_j)| < e^{-\alpha M}$, where $\alpha > 0$ is some constant. However, in the case of a function with a jump we achieve 0th order convergence at the discontinuity and away from the discontinuity we only achieve first order polynomial convergence, i.e. for $x_j \neq x_d$, $|f(x_j) - f_{M, \Delta \xi}(x_j)| \sim O(1/M)$

(Gottlieb and Shu 1997). In general, if the truncation error has k^{th} order convergence, then $|f(x) - f_{M,\Delta\xi}(x)| \sim O(1/M^k)$.

More generally, from the “integration by parts coefficient bound” described by Boyd (2001) (see also Ruijter et al. 2015), if the function is smooth up to and including its $(k-2)^{\text{th}}$ derivative, and its k^{th} derivative is integrable, then the Fourier coefficients decrease as $O(1/\xi^k)$. From Boyd (2001) we also have the “last coefficient error estimate” which states that for polynomial convergence we can approximately bound the error performance of a function with the discontinuity in the $(k-1)^{\text{th}}$ derivative as $O(1/M^{k-1})$. However, these are upper bounds; as observed by Ruijter et al. (2015), it is often the case that an error convergence of $O(1/M^k)$ or even better is seen and that this may be due to the alternating behaviour of the Fourier coefficients.

Investigating and overcoming the Gibbs phenomenon is a mature field with applications in many areas. As a result, there is a large body of literature proposing different solutions to the problem. Some of these are too computationally heavy to be useful for our application, such as adaptive filtering and mollifiers suggested by Tadmor and Tanner (2005) and Tadmor (2007). In this article we adopt the approach of Ruijter et al. (2015) by using simple spectral filtering techniques which are applied by a pointwise multiplication in the Fourier domain and therefore add very little computational load.

In the papers by Vandeven (1991) and Gottlieb and Shu (1997), a filter of order p is defined as a function $\sigma(\eta)$ supported on $\eta \in [-1, 1]$ with the following properties:

$$\begin{aligned} \text{a) } & \sigma(0) = 1, \sigma^{(l)}(0) = 0 \\ \text{b) } & \sigma(\eta) = 0 \text{ for } |\eta| = 1 \\ \text{c) } & \sigma(\eta) \in C^{p-1}. \end{aligned} \tag{27}$$

The scaled variable η is related to ξ in our application as $\eta = \xi/\xi_{\max}$. In this paper we investigate the use of two filters. The exponential filter, described by Gottlieb and Shu (1997) has the form

$$\sigma(\eta) = e^{-\vartheta\eta^p}, \tag{28}$$

where p is even and positive. This does not strictly meet criterion b in Eq. (27) as it does not go exactly to zero when $|\eta| = 1$. However, if we select $\vartheta < \varepsilon \log 10$, where $10^{-\varepsilon}$ is machine precision, then the filter coefficients are within computational accuracy of the requirements. An advantage of the exponential filter is that it has a simple form and the order of the filter is equal to the parameter p which is directly input to the filter equation.




Figure 2 Shape of the exponential filter (left) and Planck taper (right) with different parameter values.

The other filter we study here is the Planck taper (McKechan et al. 2010), which is defined piecewise as

$$\sigma(\eta) = \begin{cases} 0, & \eta \leq \eta_1, & \eta_1 = -1 \\ \frac{1}{e^{z(\eta)+1}}, & z(\eta) = \frac{\eta_2 - \eta_1}{\eta - \eta_1} + \frac{\eta_2 - \eta_1}{\eta - \eta_2}, & \eta_1 < \eta < \eta_2, & \eta_2 = \epsilon - 1 \\ 1, & \eta_2 \leq \eta \leq \eta_3, & \eta_3 = 1 - \epsilon \\ \frac{1}{e^{z(\eta)+1}}, & z(\eta) = \frac{\eta_3 - \eta_4}{\eta - \eta_3} + \frac{\eta_3 - \eta_4}{\eta - \eta_4}, & \eta_3 < \eta < \eta_4, & \eta_4 = 1 \\ 0, & \eta \geq \eta_4. \end{cases} \quad (29)$$

The value of ϵ gives the proportion of the range of η which is used for the slope regions. Outside these regions, it is completely flat with a value of 1. This contrasts with the exponential filter which introduces some, albeit often very minor, distortion for any value of $\eta \neq 0$. In addition the Planck taper has the notable property that for all values of $\epsilon > 0$, $\sigma(\eta, \epsilon) \in C^\infty$ and therefore the order of the Planck taper is ∞ . However, it is clear that different values of ϵ give a different filter shape, so the order of a filter alone cannot be taken as a predictor of performance. Examples of the two filters are shown in Figure 2.

2.4.2. Hilbert transform The calculation of the Hilbert transform of a function $\widehat{f}(\xi)$ can be realised with an inverse/forward Fourier transform pair and multiplication by the signum function,

$$\mathcal{H}[\widehat{f}(\xi)] = -i \mathcal{F}_{x \rightarrow \xi} [\text{sgn}(x) \mathcal{F}_{\xi \rightarrow x}^{-1} \widehat{f}(\xi)]. \quad (30)$$

However, this gives an error performance which is polynomially decreasing with the number of grid points M . In order to obtain exponential error convergence, Feng and Linetsky (2008) and Fusai et al. (2016) have implemented the Hilbert transform using the sinc expansion techniques comprehensively studied by Stenger (1993, 2011). Stenger showed that, given a function $\widehat{f}(\xi)$ which is analytic in the whole plane including the real axis, the function and its Hilbert transform can be expressed as

$$\widehat{f}(\xi) = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi) \frac{\sin(\pi(\xi - k\Delta\xi)/\Delta\xi)}{\pi(\xi - k\Delta\xi)/\Delta\xi}, \quad (31)$$

$$\mathcal{H}[\widehat{f}(\xi)] = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi) \frac{1 - \cos(\pi(\xi - k\Delta\xi)/\Delta\xi)}{\pi(\xi - k\Delta\xi)/\Delta\xi}, \quad (32)$$

where $\Delta\xi$ is the grid step size in the Fourier domain. Stenger (1993) also showed that, when the function $\widehat{f}(\xi)$ is analytic in a strip of the complex plane including the real axis, the expressions in Eqs. (31) and (32) are approximations whose error decays exponentially as $\Delta\xi$ decreases. In addition to discretisation, the infinite sum in Eq. (32) must also be truncated to the grid size M , so that the discrete approximation of the Hilbert transform becomes

$$\mathcal{H}[\widehat{f}(\xi)] \approx \sum_{k=-M/2}^{+M/2} \widehat{f}(k\Delta\xi) \frac{1 - \cos(\pi(\xi - k\Delta\xi)/\Delta\xi)}{\pi(\xi - k\Delta\xi)/\Delta\xi}. \quad (33)$$

Feng and Linetsky (2008, 2009) showed that if $\widehat{f}(\xi)$ decays at least exponentially as $\xi \rightarrow \infty$, i.e. $\widehat{f}(\xi) \leq \kappa \exp(-c|\xi|^\nu)$, then the error in the Hilbert transform and the Plemelj-Sokhotsky relations caused by truncating the infinite sum in Eq. (32) is also exponentially bounded. Furthermore Feng and Linetsky showed that if $\widehat{f}(\xi)$ is polynomially bounded as $\xi \rightarrow \infty$, i.e. $\widehat{f}(\xi) \leq c|\xi|^\nu$, then the error caused by truncating the series is no longer exponentially bounded (Feng and Linetsky 2008, 2009).

2.4.3. Pricing method: single-barrier options with the Spitzer identity Two of the pricing methods that we modify in order to reduce the errors from the discrete Hilbert transform were devised and explained in depth by Fusai et al. (2016). The first method which we examine is the pricing procedure for single-barrier options. Without loss of generality we consider only the down-and-out case; the modifications that we propose are equally applicable to other types of single-barrier options. This method is briefly described here in order to provide a backdrop to the changes that were made to improve convergence.

1. Set the number of dates to $N - 2$ so that the characteristic function acts as a smoothing function for the first and last dates in the scheme.

2. Compute the characteristic function $\Psi(\xi + i\alpha, \Delta t)$, where α is the damping factor used in Section 2.1.

3. Use the Plemelj-Sokhotsky relations with the sinc method to factorise

$$\Phi(\xi, q) := 1 - q\Psi(\xi + i\alpha, \Delta t) = \Phi_+(\xi, q)\Phi_-(\xi, q) \quad (34)$$

with q selected according to the criteria specified by Abate and Whitt (1992b) for the inverse z -transform.

4. Decompose

$$P(\xi, q) := \frac{e^{-i\xi} \Psi(\xi + i\alpha, \Delta t)}{\Phi_-(\xi, q)} = P_+(\xi, q) + P_-(\xi, q) \quad (35)$$

and calculate

$$F(\xi, q) := \widehat{\phi}^*(\xi) \Psi(\xi + i\alpha, \Delta t) e^{i\xi} \frac{P_+(\xi, q)}{\Phi_+(\xi, q)}. \quad (36)$$

5. Calculate the price

$$v(0, N) := e^{-rT} \mathcal{F}_{\xi \rightarrow x=0}^{-1} \mathcal{Z}_{q \rightarrow n=N-2}^{-1} [F(\xi, q)]. \quad (37)$$

The Spitzer identities give the z -transform of the characteristic function, so in order to obtain the price the inverse z -transform must be applied. The method used was devised by Abate and Whitt (1992a,b) and approximates the inverse z -transform by

$$f(t_n) \approx \frac{1}{2n\rho^n} \left[\tilde{f}(\rho) + 2 \sum_{j=1}^{n-1} (-1)^j \operatorname{Re} \tilde{f}\left(\rho e^{\frac{\pi j i}{n}}\right) + (-1)^n \tilde{f}(-\rho) \right]. \quad (38)$$

The number of terms in the summation in Eq. (38) is determined by the number of monitoring dates. However, the Euler acceleration (see e.g. O’Cinneide 1997) allows one to achieve excellent accuracy with a fixed number of terms in the summation which is independent of the number of dates. This is explained in detail in Fusai et al. (2016); the basic idea is to approximate the results by the binomial average, also called the Euler transform, of a smaller number of terms. First the partial sums

$$b_k = \frac{1}{2} \tilde{f}(\rho) + \sum_{j=1}^k (-1)^j \operatorname{Re} \tilde{f}\left(\rho e^{\frac{\pi j i}{n}}\right) \quad (39)$$

are calculated for $k = n_E, \dots, n_E + m_E$ and then the binomial average of the values of b_k is taken, i.e.

$$f(t_n) \approx \frac{1}{2^{m_E} n \rho^n} \sum_{j=0}^{m_E} \binom{m_E}{j} b_{n_E+j}. \quad (40)$$

The parameters n_E and m_E are chosen to be large enough to attain sufficient accuracy and small enough such that $n_E + m_E \ll n$. Tests suggest that a choice of $n_E = 12$ and $m_E = 20$ provides good accuracy. This gives $n_E + m_E = 32$ which is much smaller than the number of dates specified in most option contracts. The parameter ρ controls the accuracy of the inverse z -transform; in order to have an accuracy of $10^{-2\gamma}$, one must set $\rho = 10^{-\gamma/n}$ (Abate and Whitt 1992b). This can result in very small values of ρ and so it has been found in practice that the best achievable performance is of the order of 10^{-12} with $\gamma = 6$. However, this is more than sufficiently low for practical purposes and to show whether exponential convergence is achieved.

Fusai et al. (2016) showed that this method could achieve exponential convergence with a wide range of Lévy processes. However, the performance of the method with the variance gamma process

only achieved polynomial convergence. This is consistent with the error behaviour of the discrete Hilbert transform with the variance gamma process, as explained in Section 2.4.2 above. Section 3 explains in more detail how the error performance is bounded when this process is used.

In order to improve the result, we multiplied the characteristic function by a spectral filter $\sigma(\eta)$ so that the input to both the factorisation and decomposition steps decay exponentially. The expressions in Eqs. (34) and (35) are replaced by

$$\Phi(\xi, q) := 1 - q\Psi(\xi + i\alpha, \Delta t)\sigma\left(\frac{\xi}{\xi_{\max}}\right) = \Phi_+(\xi, q)\Phi_-(\xi, q), \quad (41)$$

$$P(\xi, q) := \frac{e^{-i\xi\eta}\Psi(\xi + i\alpha, \Delta t)\sigma\left(\frac{\xi}{\xi_{\max}}\right)}{\Phi_-(\xi, q)} = P_+(\xi, q) + P_-(\xi, q) \quad (42)$$

Numerical results with the updated method for the double-barrier case are shown in Section 4.

2.4.4. Pricing method: double-barrier options with the Spitzer identity The second method from Fusai et al. (2016) which we examine in this paper is the pricing procedure for double-barrier options. This is very similar to the method for the single-barrier options described in Section 2.4.3, in that it uses Wiener-Hopf factorisation and decomposition to compute the appropriate Spitzer identity. However, the major difference in this case is that the equations cannot be solved directly and so require the use of a fixed-point algorithm. The steps in the pricing procedure for double-barrier options are the same as the procedure described for the single-barrier down-and-out option described in Section 2.4.3 with the exception of Step 4 which is now replaced by the following fixed-point algorithm:

4. (a) Set $J_+(\xi, q) = J_-(\xi, q) = 0$.
- (b) Decompose

$$P(\xi, q) := \frac{e^{-i\xi\eta}\Psi(\xi + i\alpha, \Delta t)}{\Phi_-(\xi, q)} - \frac{e^{i(u-l)\xi}J_+(\xi, q)}{\Phi_-(\xi, q)} = P_+(\xi, q) + P_-(\xi, q) \quad (43)$$

and calculate $J_-(\xi, q) := P_-(\xi, q)\Phi_-(\xi, q)$.

- (c) Decompose

$$Q(\xi, q) := \frac{e^{-i\xi\eta}\Psi(\xi + i\alpha, \Delta t)}{\Phi_+(\xi, q)} - \frac{e^{i(l-u)\xi}J_-(\xi, q)}{\Phi_+(\xi, q)} = Q_+(\xi, q) + Q_-(\xi, q) \quad (44)$$

and calculate $J_+(\xi, q) := Q_+(\xi, q)\Phi_+(\xi, q)$.

- (d) Calculate

$$F(\xi, q) := \hat{\phi}^*(\xi) \frac{\Psi(\xi + i\alpha, \Delta t)}{\Phi(\xi, q)} [\Psi(\xi + i\alpha, q) - e^{i\xi\eta}J_-(\xi, q) - e^{i\xi\eta}J_+(\xi, q)]. \quad (45)$$

(e) If the difference between the new and the old value of $F(\xi, q)$ is less than a predefined tolerance or the number of iterations is greater than a certain value, e.g. 5, then calculate the price using Eq. (37), otherwise return to step (b).

Unlike the direct method for single-barrier options described in Section 2.4.3, this iterative method is limited to polynomial error convergence for all processes. In Section 3 we show that this is due to the Gibbs phenomenon. In order to improve the error performance we placed a filter $\sigma(\eta)$ on the input to each decomposition step in the fixed-point algorithm. The calculations for $P(\xi, q)$ and $Q(\xi, q)$ in Eqs. (43) and (44) are replaced by

$$P(\xi, q) := \sigma\left(\frac{\xi}{\xi_{\max}}\right) \left[\frac{e^{-i\xi} \Psi(\xi + i\alpha, \Delta t)}{\Phi_-(\xi, q)} - \frac{e^{i(u-l)} J_+(\xi, q)}{\Phi_-(\xi, q)} \right] \quad (46)$$

$$Q(\xi, q) := \sigma\left(\frac{\xi}{\xi_{\max}}\right) \left[\frac{e^{-iu\xi} \Psi(\xi + i\alpha, \Delta t)}{\Phi_+(\xi, q)} - \frac{e^{i(l-u)} J_-(\xi, q)}{\Phi_+(\xi, q)} \right]. \quad (47)$$

It must also be noted that this change is only designed to provide significant improvements to the double-barrier method with exponentially decaying characteristic functions. In the case of a polynomially decaying characteristic function such as that of the variance gamma process, this method will also be subject to the same limitations on accuracy as described in Section 3.1 for single-barrier options. Therefore, if we wish to use this scheme with the variance gamma process, we must also apply filtering to the factorisation step as shown in Eq. (41). Numerical results with the updated method for the double-barrier case are shown in Section 4.

2.4.5. Pricing method: Feng and Linetsky The third pricing method that we examine in order to illustrate the improvements obtained by the addition of spectral filtering to the sinc-based Hilbert transform is the recursive one published by Feng and Linetsky (2008) and explained in Section 2.2. In general, the FL method achieves excellent results for both single and double-barrier options (Feng and Linetsky 2008, Fusai et al. 2016); the error converges exponentially with grid size and reaches machine accuracy for fairly small grid sizes. However, with respect to the FGM model, it has the disadvantage that the computational time increases linearly with the number of monitoring dates.

Similarly to the FGM method for single-barrier options, exponential error convergence is achieved only for processes where the characteristic function reduces exponentially as $|\xi| \rightarrow \infty$. Therefore, poor error performance is achieved for the variance gamma process which has a characteristic function which only reduces polynomially as $|\xi| \rightarrow \infty$. Feng and Linetsky (2008) explained this in some detail, showing how this is linked to the truncation error of the discrete Hilbert transform. In order to improve the results, we altered the FL method by placing a filter on the input to the Hilbert transform to ensure it decays exponentially. We replaced Eqs. (16) and (17) by

$$\begin{aligned} \hat{v}(\xi, t_{n-1}) = & \frac{1}{2} \left\{ \sigma\left(\frac{\xi}{\xi_{\max}}\right) \Psi(\xi + i\alpha, \Delta t) \hat{v}(\xi, t_n) \right. \\ & \left. + e^{i\xi} i\mathcal{H} \left[e^{-i\xi} \sigma\left(\frac{\xi}{\xi_{\max}}\right) \Psi(\xi + i\alpha, \Delta t) \hat{v}(\xi, t_n) \right] \right\}, \end{aligned} \quad (48)$$

$$\widehat{v}(\xi, t_{n-1}) = \frac{1}{2} \left\{ e^{i\xi} i\mathcal{H} \left[e^{-i\xi} \sigma \left(\frac{\xi}{\xi_{\max}} \right) \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) \right] - e^{iu\xi} i\mathcal{H} \left[e^{-iu\xi} \sigma \left(\frac{\xi}{\xi_{\max}} \right) \Psi(\xi + i\alpha, \Delta t) \widehat{v}(\xi, t_n) \right] \right\}. \quad (49)$$

Numerical results with the updated method are shown in Section 4.

3. Error performance of the pricing procedure

In this section we examine the error performance of the different calculations that make up the original pricing procedures without filtering and show bounds for the individual steps. In doing this, the effect of each step in the procedure on the shape of the output function in the Fourier domain is examined, as this largely determines the error performance of the successive steps. In the FGM and FL pricing methods, the computation of the characteristic function is done directly in the Fourier domain so there are no numerical errors associated with this calculation. All the Lévy processes that we are considering have characteristic functions that decay exponentially as $|\xi| \rightarrow \infty$, with the exception of the variance gamma process where the characteristic function decays polynomially and is bounded as $|\xi|^{-2\Delta t/\nu}$. The damping factor α is omitted from the calculations to make the notation more concise. This is appropriate as the value of $i\alpha$ becomes insignificant as $|\xi| \rightarrow \infty$.

3.1. Pricing single-barrier options with the variance gamma process using the Spitzer identity

Following the calculation of the characteristic function, the next step in the pricing procedure is the factorisation of $\Phi(\xi, q) = [1 - q\Psi(\xi, \Delta t)]$, which means that we need to apply the discrete Hilbert transform to $\log \Phi(\xi, q) = \log[1 - q\Psi(\xi, \Delta t)]$. With the exception of the variance gamma process, as $|\xi| \rightarrow \infty$, $q\Psi(\xi, \Delta t) \sim qe^{-\Delta t\xi^2}$ which quickly becomes very small. Thus we can say that as $|\xi| \rightarrow \infty$, $|\log[1 - q\Psi(\xi, \Delta t)]| < ce^{-\kappa\Delta t\xi^2}$ with c, κ positive constants. Therefore, from the error bounds for the sinc-based Hilbert transform proved by Stenger (1993) and Feng and Linetsky (2008), the output of the decomposition of $\log \Phi(\xi, q)$ has exponential error performance for exponentially decaying characteristic functions.

In the case of the variance gamma process, the characteristic function is

$$\Psi(\xi, t) = \left(1 - i\nu\theta\xi + \frac{1}{2}\nu^2\xi^2\sigma^2 \right)^{-t/\nu}. \quad (50)$$

When the value of ξ is very large then $\Psi(\xi, \Delta t)$ is dominated by $\xi^{-2\Delta t/\nu}$, so when $|\xi| \rightarrow \infty$, $|\log[1 - q\Psi(\xi, \Delta t)]| < c\xi^{-2\Delta t/\nu}$. Therefore, we can bound the truncation error from the decomposition of $\log[1 - q\Psi(\xi, \Delta t)]$. Feng and Linetsky (2008) showed that the truncation error from applying the sinc-based Hilbert transform to a function which decays as $c|\xi|^{-2\Delta t/\nu}$ is bounded by

$\frac{2c\nu}{2\Delta t - \nu}(M\Delta\xi)^{-(2\Delta t/\nu-1)}$, where there is a constraint on the process parameters of $\Delta t > \nu/2$. We show that if we take into account the form of the discrete Hilbert transform and the similarity between the positive and negative tails of the characteristic function, a tighter bound can be defined and the constraints on the parameters can be relaxed. Defining $f_{\Delta\xi}(\xi)$ as the output of the infinite sum from Eq. (32) and $f_{\Delta\xi,M}(\xi)$ as the output of the truncated sum from Eq. (33),

$$\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< c_1\Delta\xi \sum_{k=M/2}^{\infty} \left| \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{\xi - k\Delta\xi} \right| + \sum_{k=-\infty}^{-M/2} \left| \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{\xi - k\Delta\xi} \right| \\
&< c_1\Delta\xi \sum_{k=M/2}^{\infty} \left| \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{\xi - k\Delta\xi} + \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{\xi + k\Delta\xi} \right| \\
&< c_2\Delta\xi \sum_{k=M/2}^{\infty} \left| \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{\xi^2 - (k\Delta\xi)^2} \right| \\
&< c_2\Delta\xi \sum_{k=M/2}^{\infty} \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{(k\Delta\xi)^2} \\
&< c_3 \int_{M\Delta\xi/2}^{+\infty} \xi^{-(\frac{2\Delta t}{\nu}+2)} d\xi < c_4(M\Delta\xi)^{-(\frac{2\Delta t}{\nu}+1)}, \tag{51}
\end{aligned}$$

where c_1 , c_2 , c_3 and c_4 are positive constants. In this case, for the integral to converge we must have $2\Delta t/\nu + 2 > 1$, which is the case for all possible process parameters. When the output of this decomposition is exponentiated to obtain the results of the factorisation, the error will be bounded by

$$\left| \frac{e^{f_{\Delta\xi}(\xi)} - e^{f_{\Delta\xi,M}(\xi)}}{e^{f_{\Delta\xi}(\xi)}} \right| < c_5 \left[1 - e^{c_4(M\Delta\xi)^{-(\frac{2\Delta t}{\nu}+1)}} \right], \tag{52}$$

where c_4 and c_5 are positive constants. For large M this converges as $O(M^{-(2\Delta t/\nu+1)})$, thus the factorisation error convergence is polynomial.

In considering the error performance of the pricing method as a whole, we must also consider the shape of the output of the factorisation in the Fourier domain as this will influence the error performance of the subsequent step. Figure 3 shows that the function flattens out at high values of $|\xi|$ and asymptotically approaches 1. Therefore, if we were to input $\Phi_{\pm}(\xi, \Delta t)$ directly to the Hilbert transform in the decomposition step then we would not be able to bound the truncation error using Feng and Linetsky's error limit for exponentially bounded functions.

However, the last date is taken out of the FGM pricing scheme. This means that we multiply the function to be decomposed by the characteristic function. In the case of exponentially decaying characteristic functions, this restores the exponential decay of the function for high values of ξ which again means that the truncation error of the discrete Hilbert transform is exponentially bounded. However, if the variance gamma process is used then the input to the decomposition is only polynomially decaying and thus we again have polynomial error convergence for this stage.




Figure 3 Input and output functions for the factorisation of $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t)$ with the Kou process for $q = \rho$. The plot shows how the decay of the function is changed by the decomposition.

3.2. Double-barrier options with the unfiltered Spitzer identity

The original pricing procedure for double-barrier options shows polynomial convergence for all processes, even those whose characteristic function decays exponentially. The main difference between the pricing procedure for single and double-barrier options is the presence of the fixed-point algorithm and in this section we show how this causes the polynomial error convergence. As shown in Section 3.1, with an exponentially decaying characteristic functions the factorisation has exponential error convergence. In addition we multiply the input to the fixed-point algorithm by the characteristic function, which means that it is exponentially bounded as $|\xi| \rightarrow \infty$. Provided the input function to the first iteration of the fixed-point algorithm is exponentially bounded, the error on the output of the initial decomposition is exponentially bounded. However, the decomposition operation is equivalent to multiplying the function in the x domain by either $\mathbf{1}_{\mathbb{R}_+}(x)$ or $\mathbf{1}_{\mathbb{R}_-}(x)$, which introduces a jump into the output functions. Due to the Gibbs phenomenon, this means that the output function from the decomposition decays as $O(1/\xi)$ as $\xi \rightarrow \infty$. The effect of this is that the input function to the second iteration of the fixed-point algorithm is no longer exponentially bounded and so, according to Stenger (1993) and Feng and Linetsky (2008), the error from the truncation of the infinite sum in Eq. (32) to give Eq. (33) is no longer exponentially bounded. A bound for this error is

$$\begin{aligned}
 |f_{\Delta\xi} - f_{\Delta\xi, M}| &= \left| \frac{\Delta\xi}{\pi} \left(\sum_{k=M/2}^{\infty} \frac{1}{k\Delta\xi(\xi - k\Delta\xi)} + \sum_{k=-\infty}^{-M/2} \frac{1}{k\Delta\xi(\xi - k\Delta\xi)} \right) \right| \\
 &< c_1 \Delta\xi \sum_{k=M/2}^{\infty} \frac{1}{(k\Delta\xi)^2} \\
 &< c_2 \int_{M\Delta\xi/2}^{+\infty} \frac{1}{\xi^2} d\xi < c_3 \frac{1}{M\Delta\xi}.
 \end{aligned} \tag{53}$$

where c_1, c_2 and c_3 are positive constants. Therefore, using the fixed-point algorithm with more than one iteration means that the error is no longer exponentially bounded. The bound shown in Eq. (53) is $O(1/M)$. However, the error of the pricing procedure actually decays as $O(1/M^2)$; this better performance may be due to the alternating nature of the Fourier coefficients.

3.3. Feng and Linetsky pricing method with the variance gamma process

The FL method is described in Eqs. (16) and (17), which show how the Hilbert transform is applied for each monitoring date. As explained in Section 3.2, the application of the Hilbert transform introduces a discontinuity into the function in the log-price domain, therefore the Fourier coefficients on the output of the Hilbert transform will decay as $O(1/\xi)$ as $\xi \rightarrow \infty$. However, before the Hilbert transform is applied for the next monitoring date, the Fourier domain function is multiplied by the characteristic function of the underlying process. Therefore, as explained by Feng and Linetsky (2008), if the characteristic function is exponentially decaying, this will result in an exponentially convergent error. However, with polynomially decaying characteristic functions, such as that of the variance gamma process, then a polynomially convergent error will be achieved.

3.4. Error performance with filtering on the sinc-based Hilbert transform

The multiplication by a filter with exponentially decaying coefficients as $|\xi| \rightarrow \infty$ gives an exponentially convergent truncation error for the sinc-based discrete Hilbert transform compared with the non-truncated version. However, filtering distorts the function somewhat. The numerical results with the updated method are shown in Section 4 and the prices calculated with the filtered version have been compared with the price calculated using the unfiltered FL method with the maximum grid size to confirm that any distortion error is less significant than the improvement in error convergence. Due to the error being influenced by these two opposing effects, we have not attempted to devise a tight error bound which closely matches the improvement in performance achieved in practice. It is often seen in the literature on the Gibbs phenomenon that the empirical results outstrip the calculated error bounds. For example, Ruijter et al. (2015) suggest that the faster convergence they see may be due to the alternating nature of the Fourier coefficients.

4. Numerical results

We performed numerical tests using the pricing schemes updated to include filtering, as described in Section 2. The results for the FGM method for double-barrier options with exponentially decaying characteristic functions are presented in Section 4.1. Section 4.2 contains results for all methods with the variance gamma process. Details of the contract and the model parameters are included in Table 4 in the Appendix. The numerical results were obtained using MATLAB R2016b running under OS X Yosemite on a 2015 Retina MacBook Pro with a 2.7GHz Intel Core i5 processor and 8GB of RAM.

4.1. Results with exponentially decaying characteristic functions

We present results for the FGM method for double-barrier options with filtering included in the fixed-point algorithm as described in Section 2.4.4. We examined the performance for both the Kou and NIG processes with $N = 4, 52$ and 252 . The values of 52 and 252 represent weekly and daily monitoring over 1 year. Results with $N = 4$ are presented in order to show the performance of the method with very few monitoring dates. Figure 4 shows results for the Kou process and Figure 5 shows results for the NIG process. The original FL and FGM methods are labelled “FL” and “FGM”. The FGM method with filtering is labelled “FGM-E, p =order” for results with the exponential filter and “FGM-P, ϵ =parameter” with the Planck taper. Comparing the results for all methods, we see that the FL method gives the best error convergence versus grid size. This is due to the error of the FGM method being limited by the performance of the inverse z -transform. Comparing the filtered FGM methods, the exponential filter gives better results but the Planck taper is less sensitive to variations in the filter shape. The best results were achieved with an exponential filter of order $p = 12$.

Tables 1 and 2 present the number of iterations and the computational time for a range of dates. The results demonstrate that as the number of dates increases, the number of iterations and computational time either does not increase, or minimally increases, and thus confirm that the computational time is independent of the number of monitoring dates. Figures 4 and 5 show how the convergence of the numerical techniques changes with the grid size and Figures 6 and 7 show how the convergence behaviour corresponds to computational time with an exponential filter of order 12.

The inclusion of a filter in the FGM method produces a large improvement compared to the unfiltered method. Despite this improvement, for low numbers of monitoring dates the FL method shows the best performance. However, for 252 monitoring dates, the filtered FGM method performs around the same as the FL method for errors $> 10^{-10}$ and for higher number of dates, the filtered FGM method shows the best performance for errors $> 10^{-10}$. Including the filter in the FL method produces a result with very slightly worse absolute error performance but which still retains exponential convergence. We can relate this to the error discussion in Section 3.4: the filter causes a slight distortion which degrades the absolute error performance, but there is no improvement to be gained in the rate of convergence as the unfiltered method already achieves exponential convergence.




Figure 4 Error vs. grid size M for the Kou process and varying number of monitoring dates N . The filter improves the performance of the FGM method from polynomial to exponential. The best results are obtained with an exponential filter of order $p = 12$.

Dates	Tolerance	M	Average iterations	Price	Error	CPU time
4	E-8	1024	2.000	0.00721968941	4.12E-14	5.63E-03
52	E-8	1024	2.000	0.00518403635	3.07E-13	3.81E-02
104	E-8	1024	2.000	0.00490517113	5.54E-13	3.99E-02
252	E-8	1024	2.000	0.00465711572	4.29E-12	3.72E-02
504	E-8	1024	2.000	0.00452396360	4.31E-09	3.80E-02
4	E-10	1024	2.000	0.00721968941	4.12E-14	1.82E-02
52	E-10	1024	2.000	0.00518403635	3.07E-13	3.50E-02
104	E-10	1024	2.091	0.00490517113	5.62E-13	3.88E-02
252	E-10	1024	2.121	0.00465711572	4.31E-12	3.71E-02
504	E-10	1024	2.152	0.00452396360	4.31E-09	3.90E-02

Table 1 Results for the Kou process with the fixed-point algorithm tolerance set to 10^{-8} and 10^{-10} .

Dates	Tolerance	M	Average iterations	Price	Error	CPU time
4	E-8	1024	2.000	0.00545479385	2.38E-13	1.50E-02
52	E-8	1024	2.000	0.00359559460	5.07E-13	8.57E-02
104	E-8	1024	2.000	0.00341651334	5.92E-10	8.58E-02
252	E-8	1024	2.091	0.00328484367	3.15E-07	9.63E-02
504	E-8	1024	2.182	0.00322814330	6.84E-07	9.34E-02
4	E-10	4096	2.000	0.00545479385	7.17E-14	1.45E-02
52	E-10	4096	2.242	0.00359559460	6.70E-13	2.20E-01
104	E-10	4096	2.303	0.00341651275	3.80E-13	2.15E-01
252	E-10	4096	2.364	0.00328453104	2.33E-09	2.08E-01
504	E-10	4096	2.485	0.00322753427	7.53E-08	2.21E-01

Table 2 Results for the NIG process with the fixed-point algorithm tolerance set to 10^{-8} and 10^{-10} .

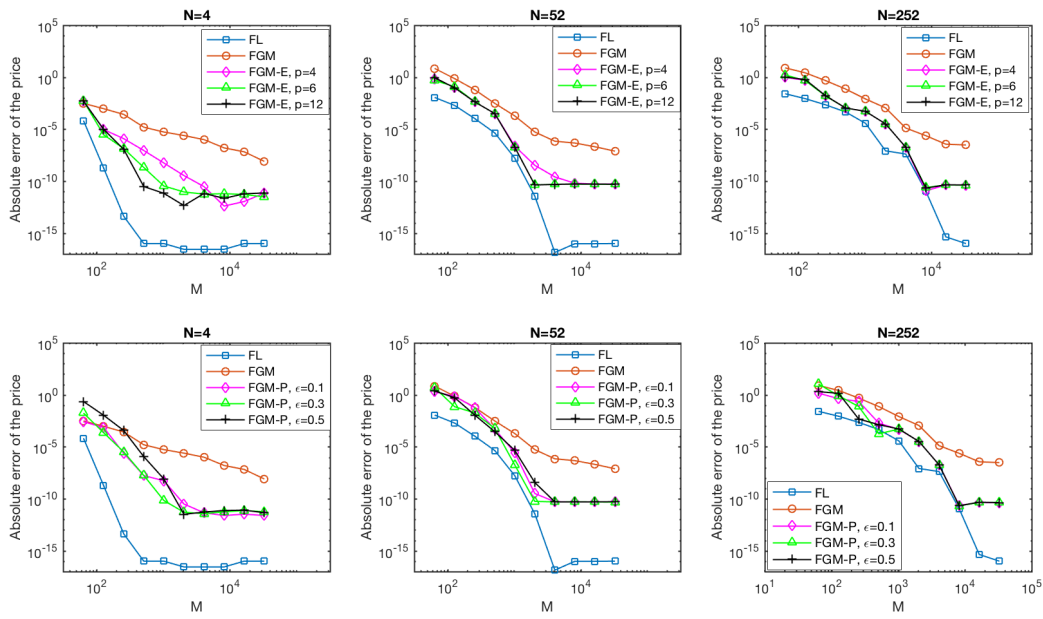


Figure 5 Error vs. grid size M with the NIG process and varying number of monitoring dates N . The filter improves the performance of the FGM method from polynomial to exponential. The best results are obtained with an exponential filter of order $p = 12$.




Figure 6 Error vs. CPU time for a double-barrier option with the Kou process and varying numbers of monitoring dates N . The filter improves the FGM method for all N ; FGM-F is the fastest method for an error of 10^{-8} with $N > 252$.

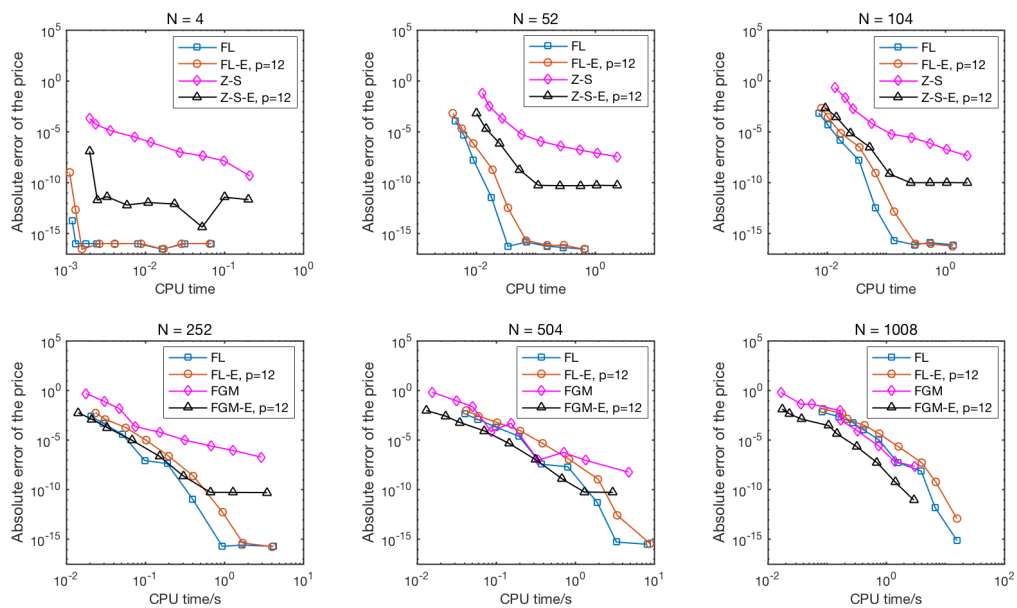


Figure 7 Error vs. CPU time for a double-barrier option with the NIG process and varying numbers of monitoring dates N . The filter improves the FGM method for all N ; FGM-F is the fastest method for an error of 10^{-8} with $N \geq 504$.

4.2. Polynomially decaying characteristic functions

We present results for the FL and FGM methods for a process with a polynomially decaying characteristic function, i.e. the variance gamma process. Figures 8 and 9 show the results of tests for single and double-barrier options where we have applied exponential filtering as described in Section 2.4.

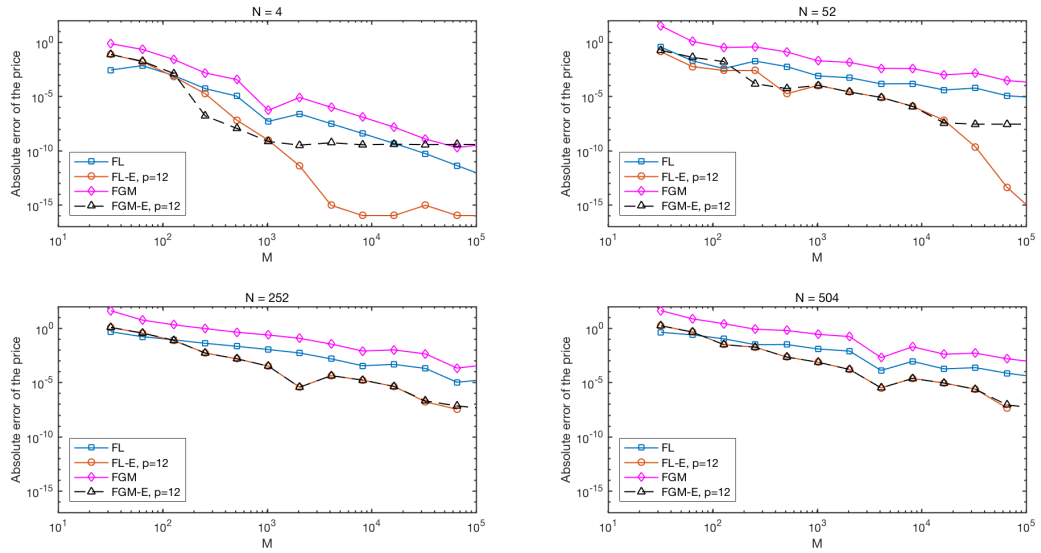


Figure 8 Error vs. grid size M for a single-barrier down-and-out option with the variance gamma process and varying numbers of monitoring dates N . The filter improves both the FGM and FL methods, with the FL-F method performing best at low numbers of dates.

The performance for a low number of dates shows a good improvement with the addition of filtering for both the FGM and FL methods. This demonstrates that the performance of the sinc-based discrete Hilbert transform of polynomially decaying functions can be improved even when the polynomial decay is a true representation of the function shape and not simply an artefact of the fixed-point algorithm as was the case in Section 4.1. For a higher number of dates, the error convergence vs grid size for the FGM method is improved so that it is the same as the FL method with or without filtering. This is a significant improvement as the FGM method has the advantage over the FL method that its computation time beyond a small threshold is independent of the number of dates, unlike the linear increase of the FL method. This is demonstrated by the results shown in Figures 10 and 11. The filtered methods show the best performance for all dates; filtered FL is the best performing method for low numbers of monitoring dates and filtered FGM is the best performing method for higher numbers of dates.

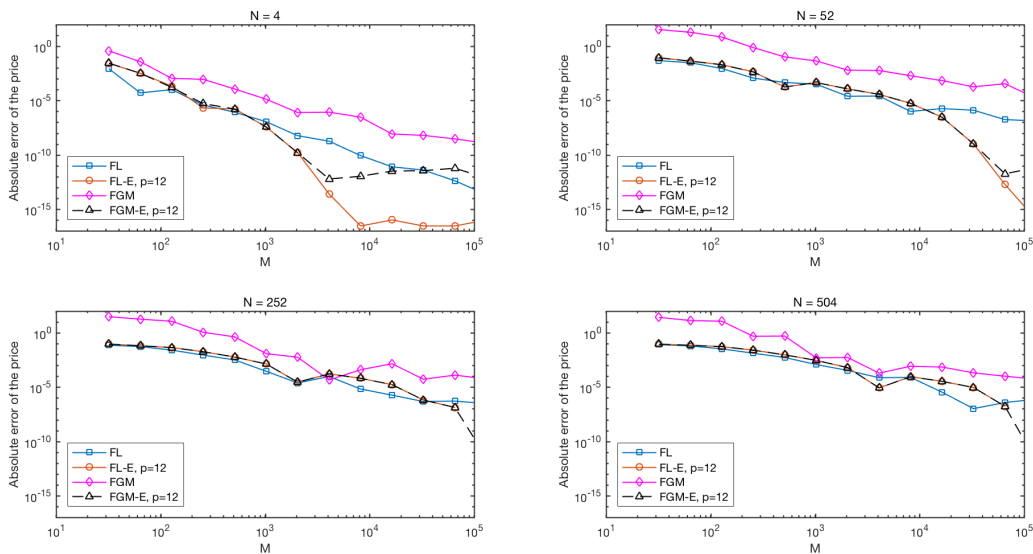


Figure 9 Error vs. grid size M for a double-barrier option with the variance gamma process and varying numbers of monitoring dates N . The filter improves both the FGM methods for all numbers of monitoring dates and improved the FL method for low numbers of dates.




Figure 10 Error vs. CPU time for a single-barrier option with the variance gamma process and varying numbers of monitoring dates N . The best performance of the new filtered methods, FL-F and FGM-F, either equals or exceeds the performance of the existing methods over all numbers of dates.

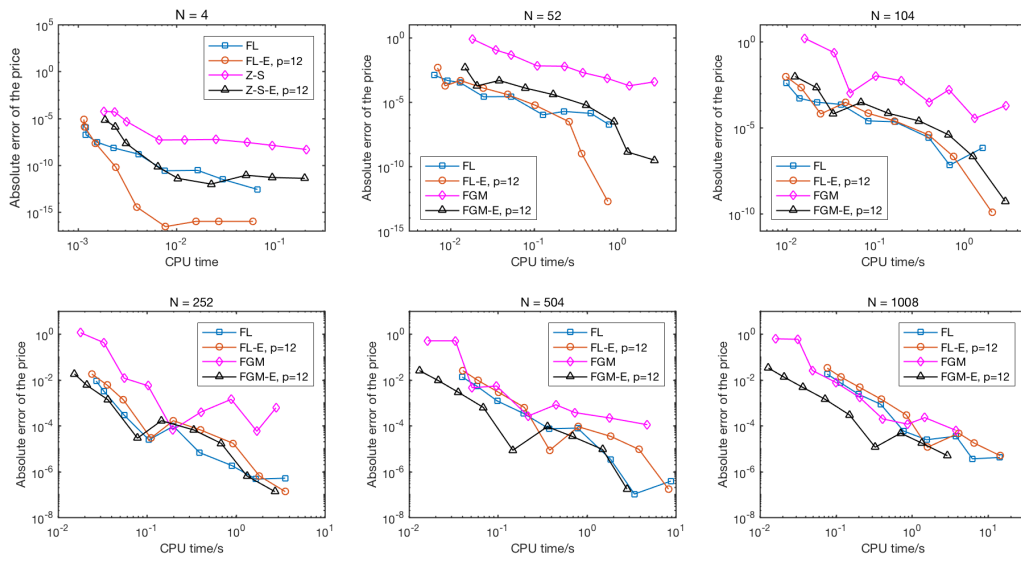


Figure 11 Error vs. CPU time for a double-barrier option with the variance gamma process and varying numbers of monitoring dates N . The best performance of the new filtered methods, FL-F and FGM-F, either equals or exceeds the performance of the existing methods over all numbers of dates.

4.3. Summary of results

Table 3 shows a summary of the best performing methods in terms of CPU time for different processes and types of options.

Dates	Single barrier	Double barrier		
	VG	Kou	NIG	VG
4	FL-E	FL	FL	FL-E
52	FL-E	FL	FL	FL-E
104	FL-E, FGM-E	FL	FL	FL-E
252	FL-E, FGM-E	FGM-E, FL	FGM-E, FL	FGM-E, FL-E, FL
504	FGM-E*	FGM-E	FGM-E	FGM-E*
1008	FGM-E*	FGM-E	FGM-E	FGM-E*

Table 3 Quickest method for an error of 10^{-8} . Due to the slower convergence of all methods with the variance gamma process, entries marked with an asterisk show the quickest method for an error of 10^{-5} . Green: a filtered method provides the best performance. Blue: the performance of the filtered methods equals, but does not exceed, the best performance of an existing method. Red: the few cases where an existing method performs best.

5. Conclusions

In this article we showed that numerical methods for pricing derivatives based on the Hilbert transform computed with a sinc function expansion can be modified with the addition of spectral filters to improve their convergence. Furthermore, we expanded on the work by Stenger and Feng and Linetsky which showed how the shape of the function on the input to the Hilbert transform relates to the resultant error on the output of the Hilbert transform. We showed that due to the Gibbs phenomenon, an algorithm using successive Hilbert transforms will achieve polynomial performance unless additional filtering is applied after the first Hilbert transform. Moreover, we demonstrated that simple spectral filters such as the exponential filter or the Planck taper are sufficient to improve performance so that exponential convergence can be achieved. In addition we showed that the pricing schemes by Feng and Linetsky and Fusai et al., which have relatively poor performance with the variance gamma process, even for single-barrier options, can also be improved by spectral filters. This article directly concerns the pricing of barrier option pricing but the findings are relevant for any application which is related to jump-diffusion in the presence of barriers and requires the solution of the Wiener-Hopf or Fredholm equation.

Appendix. Parameters

Table 4 contains all the parameters used for the numerical experiments which produced the results presented in Section 4.

Description	Symbol	Value
Option parameters		
Maturity	T	1 year
Initial spot price	S_0	1
Strike	K	1.1
Upper barrier (down-and-out)	U	$+\infty$
Upper barrier (double-barrier)	U	1.15
Lower barrier	L	0.85
Risk-free rate	r	0.05
Dividend rate	q	0.02
NIG process parameters, $\Psi(\xi, t) = e^{-t(\sqrt{\alpha^2 - (\beta + i\xi)^2} + \sqrt{\alpha^2 - \beta^2})}$		
	α	15
	β	-5
	δ	0.5
Kou process parameters, $\Psi(\xi, t) = e^{-t\left(\frac{\sigma^2 \xi^2}{2} - \lambda\left(\frac{(1-p)\eta_2}{\eta_2 + i\xi} + \frac{p\eta_1}{\eta_1 - i\xi} - 1\right)\right)}$		
	p	0.3
	λ	3
	σ	0.1
	η_1	40
	η_2	12
Variance gamma parameters, $\Psi(\xi, t) = (1 - i\nu\xi\theta + \nu\sigma^2\xi^2/2)^{-t/\nu}$		
	θ	$\frac{1}{9}$
	σ	$\frac{1}{3\sqrt{3}}$
	ν	0.25

Table 4 Parameters for the numerical tests and processes used; $\Psi(\xi, t)$ is the characteristic function of the process that models the underlying asset.

Acknowledgements

The support of the Economic and Social Research Council (ESRC) in funding the Systemic Risk Centre is gratefully acknowledged (Grant number ES/K002309/1).

References

- Abate J, Whitt W (1992a) The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems* 10(1-2):5-88, URL <http://dx.doi.org/10.1007/BF01158520>.
- Abate J, Whitt W (1992b) Numerical inversion of probability generating functions. *Operations Research Letters* 12(4):245-251, URL [http://dx.doi.org/10.1016/0167-6377\(92\)90050-D](http://dx.doi.org/10.1016/0167-6377(92)90050-D).
- Barndorff-Nielsen OE (1998) Processes of normal inverse Gaussian type. *Finance and Stochastics* 2(1):41-68, URL <http://dx.doi.org/10.1007/s007800050032>.
- Boyd JP (2001) *Chebyshev and Fourier Spectral Methods* (Heidelberg: Springer), URL <http://dx.doi.org/10.1002/zamm.19910710715>.

- Carr P, Geman H, Madan DB, Yor M (2002) The fine structure of asset returns: An empirical investigation. *Journal of Business* 75(2):305–332, URL <http://dx.doi.org/10.1086/338705>.
- Carr P, Madan D (1999) Option valuation using the fast Fourier transform. *Journal of Computational Finance* 2(4):61–73.
- Fang F, Oosterlee CW (2008) A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing* 31(2):826–848, URL <http://dx.doi.org/10.1137/080718061>.
- Fang F, Oosterlee CW (2009) Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik* 114(1):27–62, URL <http://dx.doi.org/10.1007/s00211-009-0252-4>.
- Feng L, Linetsky V (2008) Pricing discretely monitored barrier options and defaultable bonds in Lévy process models: a Hilbert transform approach. *Mathematical Finance* 18(3):337–384, URL <http://dx.doi.org/10.1111/j.1467-9965.2008.00338.x>.
- Feng L, Linetsky V (2009) Computing exponential moments of the discrete maximum of a Lévy process and lookback options. *Finance and Stochastics* 13(4):501–529, URL <http://dx.doi.org/10.1007/s00780-009-0096-x>.
- Frigo M, Johnson SG (1998) FFTW: An adaptive software architecture for the FFT. *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 3, 1381–1384 (Piscataway: IEEE), URL <http://dx.doi.org/10.1109/ICASSP.1998.681704>.
- Fusai G, Germano G, Marazzina D (2016) Spitzer identity, Wiener-Hopf factorisation and pricing of discretely monitored exotic options. *European Journal of Operational Research* 251(4):124–134, URL <http://dx.doi.org/10.1016/j.ejor.2015.11.027>.
- Gibbs JW (1898) Fourier’s series. *Nature* 59:200.
- Gibbs JW (1899) Fourier’s series. *Nature* 59:606.
- Gottlieb D, Shu C (1997) On the Gibbs phenomenon and its resolution. *SIAM Review* 39(4):644–668, URL <http://dx.doi.org/10.1137/S0036144596301390>.
- Green R, Fusai G, Abrahams ID (2010) The Wiener-Hopf technique and discretely monitored path-dependent option pricing. *Mathematical Finance* 20(2):259–288, URL <http://dx.doi.org/10.1111/j.1467-9965.2010.00397.x>.
- Heston SL (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6(2):327–343, URL <http://dx.doi.org/10.1093/rfs/6.2.327>.
- Hewitt E, Hewitt RE (1979) The Gibbs-Wilbraham phenomenon: an episode in Fourier analysis. *Archive for History of Exact Sciences* 21(2):129–160, URL <http://dx.doi.org/10.1007/BF00330404>.

- Kemperman JHB (1963) A Wiener-Hopf type method for a general random walk with a two-sided boundary. *Annals of Mathematical Statistics* 34(4):1168–1193, URL <http://www.jstor.org/stable/2238330>.
- King FW (2009) *Hilbert Transforms* (Cambridge: Cambridge University Press).
- Kou S (2002) A jump-diffusion model for option pricing. *Management Science* 48(8):1086–1101, URL <http://dx.doi.org/10.1287/mnsc.48.8.1086.166>.
- Kreyszig E (2011) *Advanced Engineering Mathematics* (New York: Wiley), 10 edition.
- Lewis A (2001) A simple option formula for general jump-diffusion and other exponential Lévy processes. SSRN 282110.
- Madan DB, Seneta E (1990) The variance gamma (V.G.) model for share market returns. *Journal of Business* 63(4):511–524, URL <http://www.jstor.org/stable/2353303>.
- Marazzina D, Fusai G, Germano G (2012) Pricing credit derivatives in a Wiener-Hopf framework. Cummins M, Murphy F, Miller JJH, eds., *Topics in Numerical Methods for Finance*, volume 19 of *Springer Proceedings in Mathematics & Statistics*, 139–154 (New York: Springer).
- McKechan DJA, Robinson C, Sathyaprakash BS (2010) A tapering window for time-domain templates and simulated signals in the detection of gravitational waves from coalescing compact binaries. *Classical and Quantum Gravity* 27(8):084020, URL <http://stacks.iop.org/0264-9381/27/i=8/a=084020>.
- Merton RC (1976) Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3(1):125–144, URL [http://dx.doi.org/10.1016/0304-405X\(76\)90022-2](http://dx.doi.org/10.1016/0304-405X(76)90022-2).
- Nolan JP (2017) *Stable Distributions — Models for Heavy Tailed Data* (Boston: Birkhäuser), in progress, Chapter 1 online at <http://academic2.american.edu/~jpnolan>.
- O’Cinneide CA (1997) Euler summation for Fourier series and Laplace transform inversion. *Stochastic Models* 13(2):315–337.
- Polyanin AD, Manzhirov AV (1998) *Handbook of Integral Equations* (Boca Raton: CRC Press).
- Ruijter MJ, Versteegh M, Oosterlee CW (2015) On the application of spectral filters in a Fourier option pricing technique. *Journal of Computational Finance* 19(1):75–106, URL <http://oai.cwi.nl/oai/asset/23070/23070D.pdf>.
- Schoutens W (2003) *Lévy Processes in Finance* (New York: Wiley).
- Spitzer F (1956) A combinatorial lemma and its application to probability theory. *Transactions of the American Mathematical Society* 82(2):323–339, URL <http://dx.doi.org/10.1090/S0002-9947-1956-0079851-X>.
- Stenger F (1993) *Numerical Methods Based on Sinc and Analytic Functions* (Berlin: Springer).
- Stenger F (2011) *Handbook of Sinc Numerical Methods* (Boca Raton: CRC Press).
- Tadmor E (2007) Filters, mollifiers and the computation of the Gibbs phenomenon. *Acta Numerica* 16:305–378, URL <http://dx.doi.org/10.1017/S0962492906320016>.

Tadmor E, Tanner J (2005) Adaptive filters for piecewise smooth spectral data. *IMA Journal of Numerical Analysis* 25(4):635–647, URL <http://dx.doi.org/10.1093/imanum/dri026>.

Vandeven H (1991) Family of spectral filters for discontinuous problems. *Journal of Scientific Computing* 6(2):159–192, URL <http://dx.doi.org/10.1007/BF01062118>.

Wilbraham H (1848) On a certain periodic function. *Cambridge and Dublin Mathematical Journal* 3:198–201.