Fluctuation identities with continuous monitoring and their application to price barrier options

Carolyn E. Phelan

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

Daniele Marazzina

Department of Mathematics, Politecnico di Milano
daniele.marazzina@polimi.it

Gianluca Fusai

Department of Economics and Business Studies (DiSEI), Università del Piemonte
Orientale "Amedeo Avogadro", Novara

Faculty 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

November 30, 2017

Abstract

We present a numerical scheme to calculate fluctuation identities for exponential Lévy processes in the continuous monitoring case. This includes the Spitzer identities for touching a single upper or lower barrier, and the more difficult case of the two-barriers exit problem. These identities are given in the Fourier-Laplace domain and require numerical inverse transforms. Thus we cover a gap in the literature that has mainly studied the discrete monitoring case; indeed, there are no existing numerical methods that deal with the continuous case. As a motivating application we price continuously monitored barrier options with the underlying asset modelled by an exponential Lévy process. We perform a detailed error analysis of the method and develop error bounds to show how the performance is limited by the truncation error of the sinc-based fast Hilbert transform used for the Wiener-Hopf factorisation. By comparing the results for our new technique with those for the discretely monitored case (which is in the Fourier-z domain) as the monitoring time step approaches

1

zero, we show that the error convergence with continuous monitoring represents a limit for the discretely monitored scheme.

Keywords – Finance, Wiener-Hopf factorisation, Hilbert transform, Laplace transform, spectral filter.

1 Introduction

Identities providing the Fourier-z transform of probability distribution functions of the extrema of a random path subject to monitoring at discrete intervals were first published by Spitzer (1956). They were extended to the continuous case by Baxter and Donsker (1957) and to double barriers by Kemperman (1963). The identities for the minimum and maximum of a path, for use with a single upper or lower barrier and for the two-barrier exit problem, are comprehensively described in the discrete monitoring case by Fusai et al. (2016), who proposed numeric methods to compute them for exponential Lévy processes. The discretely and continuously monitored identities are in the Fourier-z and Fourier-Laplace domains respectively. This means that with the application of the inverse z or Laplace transform as appropriate, they can be used within Fourier-transform option pricing methods, which we will use as an example in this paper. The relevance of the Spitzer identity in several fields within operational research is nowadays well recognised. We mention, for example, the application to queuing systems, see the classical contributions by Cohen (1975, 1982) and Prabhu (1974) and more recent work by Bayer and Boxma (1996), Markov chains (Rogers, 1994), insurance (Chi and Lin, 2011), inventory systems (Cohen and Pekelman, 1978; Grassmann and Jain, 1989), and applied probability (Grassman, 1990), as well as in mathematical finance.

Pricing derivatives, especially exotic options, is a challenging problem often covered also in the operations research literature, see e.g. Kou (2008). 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. 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 Fourier space and by Marazzina et al. (2012) and Fusai et al. (2016) to compute the factorisations required by the Spitzer identities 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. However, the Feng and Linetsky method cannot be extended to continuously monitored options because its recursive structure makes it an inherently discrete scheme. In contrast Green et al. (2010) showed that methods based on the Spitzer identities can be extended to continuous monitoring using the Laplace transform in the time domain rather than the z-transform.

In this article we implement a method to numerically calculate the required Wiener-Hopf

2

factors and thence the Spitzer identities in continuous time; we apply this to price continuously monitored options with general exponential Lévy processes. For continuous monitoring, the Wiener-Hopf factorisation can be done analytically if the characteristic exponent is rational, i.e. for the Gaussian and Kou double exponential processes, or in some special cases, e.g. when the jumps are only positive or negative. It is also possible to approximate a non-rational exponent with a rational one that is easily factored (Kuznetsov, 2010). However, an analytical solution for the continuous monitoring case which is usable for any exponential Lévy process and does not require approximation has not been found yet. In the discrete case an analytical Wiener-Hopf factorisation can be done only for a Gaussian process (Fusai et al., 2006), but from a numerical point of view the problem is easier and there are a number of papers dealing with exponential Lévy processes. However, it is well known that the convergence of numerical methods for discrete monitoring to the continuous monitoring limit are very slow, see e.g. Broadie et al. (1997). Therefore this work contributes to the literature, both in probability as well as in applied mathematics and operational research, by providing a method to find exit probabilities in the continuous monitoring case with non-rational characteristic exponents, whereas previous numerical methods have concentrated on the discrete monitoring case.

This method follows the approach suggested by Green et al. (2010) and is based on the Fusai, Germano and Marazzina (FGM) method (Fusai et al., 2016) with spectral filtering (Phelan et al., 2017). While the latter is for discrete monitoring and thus in the Fourier-z domain, here we operate in the Fourier-Laplace domain. Besides the discrete Fourier transform (DFT), or actually the fast Fourier transform (FFT), which is a standard technique, we also require a numerical inverse Laplace transform; for the latter we used a algorithm proposed by Abate and Whitt (1992a, 1995), which is based on a Fourier series and is derived in a similar way to their well established numerical inverse z-transform (Abate and Whitt, 1992b). The error convergence is slightly worse than first-order polynomial; we explain this in detail with reference to the truncation error of the sinc-based discrete Hilbert transform. Our results show that the error convergence is consistent with the error bound and the performance of the discretely monitored technique as the monitoring interval goes to zero.

The structure of this paper is as follows. In Section 2 we briefly run through Fourier, Hilbert, Laplace and z-transforms and explain how they are used for the calculation of the Spitzer identities. We then present a numerical pricing scheme for continuously monitored options and explain its relationship with the FGM pricing scheme with discrete monitoring. Section 3 provides a discussion of the error convergence of the pricing technique with special reference to the truncation error of the sinc-based Hilbert transform. Section 4 shows the results that were achieved, comparing them with the results for the FGM method for discretely monitored options.

2 Fourier transform methods for option pricing

In this paper we make extensive use of the Fourier transform (see e.g. Kreyszig, 2011; Polyanin and Vladimirovich, 1998), an integral transform with many applications. Historically, it has been widely employed 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 financial literature, the forward and inverse Fourier transforms are defined as

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

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

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 $v(x,t)$ of an option at time $t = 0$ when the initial price of the underlying is $S(0) = S_0$, 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,

$$v(0,0) = e^{-rT}\mathrm{E}\left[\phi(x(T))e^{-\alpha x(T)}|x(0) = 0\right] = e^{-rT}\int_{-\infty}^{+\infty}\phi(x)e^{-\alpha x}p(x,T)dx$$

$$= \frac{e^{-rT}}{2\pi}\int_{-\infty}^{+\infty}\widehat{\phi}(\xi)\widehat{p}^*(\xi + i\alpha, T)d\xi = e^{-rT}\mathcal{F}_{\xi\to x}^{-1}\left[\widehat{\phi}(\xi)\widehat{p}^*(\xi + i\alpha, T)\right](0). \tag{3}$$

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

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

where $e^{\alpha x}$ 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,

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

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

$$\Psi(\xi,t) = \mathrm{E}\left[e^{i\xi X(t)}\right] = \int_{-\infty}^{+\infty} e^{i\xi x}p(x,t)dx = \mathcal{F}_{x\to\xi}\left[p(x,t)\right] = \widehat{p}(\xi,t). \tag{6}$$

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

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

The Lévy-Khincine triplet (μ, σ, ν) uniquely defines the Lévy process: μ 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

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

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), normal inverse Gaussian (NIG) (Barndorff-Nielson, 1998), CGMY (Carr et al., 2002), Kou double exponential (Kou, 2002), Merton jump diffusion (Merton, 1976), Lévy alpha stable (Nolan, 2017), Variance Gamma (VG) (Madan and Seneta, 1990) and Meixner (Schoutens, 2003) processes.

Some pricing techniques based on the Fourier transform also use the Hilbert transform, which is an integral transform related to the Fourier transform. Unlike with 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

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

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 $-i\,\text{sgn}\,x$.

Whilst the Fourier and Hilbert transform operate on the state variable (here the log price), the Laplace transform is applied to time. The forward and reverse Laplace transforms are

$$\mathcal{L}_{t \to s}[f(t)] = \widetilde{f}(s) := \int_0^{+\infty} e^{-st} f(t) dt, \quad s \in \mathbb{C} \tag{10}$$

$$\mathcal{L}_{s \to t}^{-1}[\widetilde{f}(s)] = f(t) := \frac{1}{2\pi i} \int_{a-i\infty}^{a+i\infty} e^{st} \widetilde{f}(s) ds, \tag{11}$$

where $a \in \mathbb{R}$ is on the right of all singularities of $\widetilde{f}(s)$ in the complex plane. The Laplace transform is closely related to the z-transform of a discrete function $f(t_n) = f(n)$, $n \in \mathbb{N}_0$,

$$\mathscr{Z}_{n \to q}[f(n)] := \sum_{n=0}^{\infty} q^n f(t_n), \quad q \in \mathbb{C}. \tag{12}$$

Given a continuous function $f_c(t)$, we define the discrete function $f_d(t_n)$ consisting of sampled values of the former, where Δt is the sampling interval and $t_n = n\Delta t$ are the sampling times. Then with a z-transform parameter $q = e^{-s\Delta t}$, the Laplace and z-transforms are related in the

limit $\Delta t \to \infty$:

$$\mathcal{L}_{t \to s}[f_c(t)] = \int_0^\infty e^{-st} f_c(t) dt = \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^\infty e^{-sn\Delta t} f_c(n\Delta t)$$

$$= \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^\infty (e^{-s\Delta t})^n f_d(t_n) = \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^\infty q^n f_d(t_n)$$

$$= \lim_{\Delta t \to 0} \Delta t \mathcal{Z}[f_d(t_n)]. \tag{13}$$

2.1 Spitzer identities for continuous monitoring

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 above. We need the characteristic function of the PDF of the value of a stochastic process $X(t)$ at time $t = T$, conditional on the process remaining inside continuously monitored upper and lower barriers. We use the identities published by Spitzer (1956) which were extended to the continuously monitored case by Baxter and Donsker (1957) and to double-barriers by Kemperman (1963). The Spitzer identities provide the Fourier-z transform of the PDF of a stochastic process $X(t)$ at time $t = T$, conditional on whether $X(t)$ reaches a barrier at discretely monitored times. The Fourier transform is applied to the process values and the z-transform is applied to the discrete monitoring times. Baxter and Donsker (1957) demonstrated that we can obtain the equivalent identities for continuously monitored barriers in the Fourier-Laplace domain. Green et al. (2010) showed that the relationship between the Laplace and z-transforms described in Eq. (13) can be exploited to price continuously monitored options using the Spitzer identities in the Fourier-Laplace domain.

An important aspect in the calculation of the Spitzer identities is the decomposition of a function $\widehat{f}(\xi)$ into $+$ and $-$ parts, $\widehat{f_+}(\xi) = \mathcal{F}_{x \to \xi}[f(x)\mathbf{1}_{\mathbb{R}_+}(x)]$ and $\widehat{f_-}(\xi) = \mathcal{F}_{x \to \xi}[f(x)\mathbf{1}_{\mathbb{R}_-}(x)]$, such that $\widehat{f}(\xi) = \widehat{f_+}(\xi) + \widehat{f_-}(\xi)$. This can be done directly in the Fourier domain using the Plemelj-Sokhotsky relations (King, 2009; Fusai et al., 2016):

$$\widehat{f_+}(\xi) = \frac{1}{2}\{\widehat{f}(\xi) + i\mathcal{H}[\widehat{f}(\xi)]\} \tag{14}$$

$$\widehat{f_-}(\xi) = \frac{1}{2}\{\widehat{f}(\xi) - i\mathcal{H}[\widehat{f}(\xi)]\}. \tag{15}$$

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

$$\widehat{f_{b+}}(\xi) = \frac{1}{2}\{\widehat{f}(\xi) + e^{ib\xi}i\mathcal{H}[e^{-ib\xi}\widehat{f}(\xi)]\} \tag{16}$$

$$\widehat{f_{b-}}(\xi) = \frac{1}{2}\{\widehat{f}(\xi) - e^{ib\xi}i\mathcal{H}[e^{-ib\xi}\widehat{f}(\xi)]\}. \tag{17}$$

The calculation of the Spitzer identities also requires to factorise 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 by 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 $\widehat{g_-}(\xi) = \exp \widehat{h_-}(\xi)$.

Green et al. (2010) dealt with fluctuation identities that can be used for lookback, single-

barrier and double-barrier options. Here we concentrate on the identities for single-barrier down-and-out and double-barrier options. The first step is always to factorise $\Phi_c(\xi, s) = s - \psi(\xi) = \Phi_{c+}(\xi, s)\Phi_{c-}(\xi, s)$. For a single-barrier down-and-out option, the Laplace transform of the required characteristic function is

$$\widetilde{\widehat{p}}(\xi, s) = \frac{1 - \Phi_{c-}(\xi, s)P_{cl-}(\xi, s)}{\Phi_c(\xi, s)} = \frac{P_{cl+}(\xi, s)}{\Phi_{c+}(\xi, s)}, \tag{18}$$

where $P_c(\xi, s) = 1/\Phi_{c-}(\xi, s)$ is decomposed with respect to the lower log-barrier l using Eqs. (16) and (17). For a double-barrier option, the Laplace transform of the required characteristic function is

$$\widetilde{\widehat{p}}(\xi, s) = \frac{1 - \Phi_{c-}(\xi, s)J_{cl-}(\xi, s) - \Phi_{c+}(\xi, s)J_{cu+}(\xi, s)}{\Phi_c(\xi, s)}, \tag{19}$$

where $J_{cu+}(\xi, s)$ and $J_{cl-}(\xi, s)$ are the solution to the pair of coupled equations

$$J_{cu+}(\xi, s) = \left[\frac{1 - \Phi_{c-}(\xi, s)J_{cl-}(\xi, s)}{\Phi_{c+}(\xi, s)} \right]_{u+} \tag{20}$$

$$J_{cl-}(\xi, s) = \left[\frac{1 - \Phi_{c+}(\xi, s)J_{cu+}(\xi, s)}{\Phi_{c-}(\xi, s)} \right]_{l-}. \tag{21}$$

For $u \to \infty$, $J_{cu+} \to 0$ and $J_{cl-} \to P_{cl-}$, thus recovering the Spitzer identity for the single barrier, Eq. (18). The latter can be calculated directly, while so far only an iterative solution has been found (Fusai et al., 2016; Phelan et al., 2017) to the coupled Eqs. (20) and (21).

2.1.1 Relationship to the Spitzer identities for discrete monitoring

In Section 4 we show numerical results comparing the error convergence obtained using the Spitzer identities for continuous monitoring with the performance of the closely related method using the Spitzer identities for discrete monitoring (Green et al., 2010; Fusai et al., 2016; Phelan et al., 2017).

The relationship between the two methods originates in the connection between the z-transform and the Laplace transform described in Eq. (13). As described in Section 2.1, the first step in pricing continuously monitored barrier options is the calculation of $\Phi_c(\xi, s) = s - \psi(\xi)$ in the Fourier-Laplace domain. The equivalent quantity in the Fourier-z domain for discrete monitoring is $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t)$. We can use the relation in Eq. (13) with $q = e^{-s\Delta t}$ to relate the two:

$$\lim_{\Delta t \to 0} \frac{\Delta t}{\Phi(\xi, q)} = \lim_{\Delta t \to 0} \frac{\Delta t}{1 - q\Psi(\xi, \Delta t)} = \lim_{\Delta t \to 0} \frac{\Delta t}{1 - e^{-s\Delta t}e^{\psi(\xi)\Delta t}}$$

$$= \lim_{\Delta t \to 0} \frac{\Delta t}{1 - e^{(\psi(\xi)-s)\Delta t}} = \frac{1}{s - \psi(\xi)} = \frac{1}{\Phi_c(\xi, s)}. \tag{22}$$

The same factorisation and decomposition steps described in Section 2.1 (Green et al., 2010; Fusai et al., 2016) are applied to both $\Phi(\xi, q)$ and $\Phi_c(\xi, s)$ to price options with respectively discrete or continuous monitoring.

2.2 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 sections.

2.2.1 Discrete Fourier and Hilbert transforms and spectral filters

The forward and inverse Fourier transforms shown in Eq. (1) and Eq. (2) are integrals over an infinite domain and in order to implement them numerically we need to approximate each with a discrete Fourier transform (DFT). We implement this in practice using the built-in Matlab FFT function which is based on the FFTW library by Frigo and Johnson (1998).

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 sign function in between,

$$i\mathcal{H}\big[\widehat{f}(\xi)\big] = \mathcal{F}_{x\to\xi}\big[\mathrm{sgn}(x)\mathcal{F}_{\xi\to x}^{-1}\widehat{f}(\xi)\big]. \tag{23}$$

However, this gives an error convergence 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 studied by Stenger (1993, 2011). Stenger showed that, given a function $\widehat{f}(\xi)$ which is analytic in the whole plane, the function and its Hilbert transform can be expressed as

$$\widehat{f}(\xi) = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi)\frac{\sin(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}, \tag{24}$$

$$\mathcal{H}\big[\widehat{f}(\xi)\big] = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi)\frac{1-\cos(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}, \tag{25}$$

where $\Delta\xi$ is the grid step in the Fourier domain. Stenger (1993) also showed that, when the function $f(\xi)$ is analytic in a strip of the complex plane including the real axis, the expressions in Eqs. (24) and (25) are approximations whose error decays exponentially as $\Delta\xi$ decreases. In addition to discretisation, the infinite sum in Eq. (25) must also be truncated to the grid size M, so that the Hilbert transform approximation becomes

$$\mathcal{H}\big[\widehat{f}(\xi)\big] \approx \sum_{k=-M/2}^{+M/2} \widehat{f}(k\Delta\xi)\frac{1-\cos(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}. \tag{26}$$

Feng and Linetsky (2008, 2009) showed that if $\widehat{f}(\xi)$ decays at least exponentially as $|\xi| \to \infty$, i.e. $\widehat{f}(\xi) \leq \kappa\exp(-c|\xi|^\nu)$, then the error in the Hilbert transform and thus in the Plemelj-Sokhotsky relations caused by truncating the series in Eq. (25) is also exponentially bounded. Furthermore Feng and Linetsky showed that if $\widehat{f}(\xi)$ is polynomially bounded then, although the accuracy of the series in Eq. (25) is retained, the error caused by truncating the sum is no longer exponentially bounded. However, it has subsequently been shown that multiplying the input to the Hilbert transform by a filter can improve the error convergence (Phelan et al.,
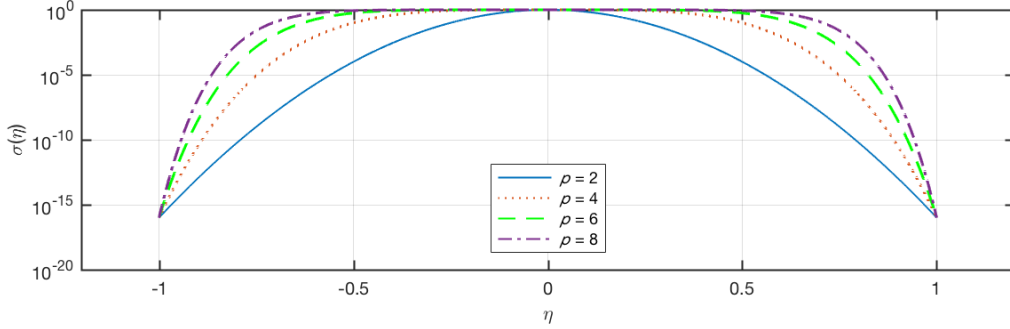
Figure 1: Shape of the exponential filter plotted with different values of p.

2017).

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

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

where p is even. 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 the machine precision, then the filter coefficients are within computational accuracy of the requirements. The exponential filter has the advantages that it has a simple form and that it can be used for any even value of p. Moreover, the order of the filter is a parameter which is directly input to the filter equation. Filter shapes for a range of p values are shown in Figure 1. Many filters other than the exponential can be used, e.g. the Planck taper (Phelan et al., 2017) and the raised cosine (Ruijter et al., 2015).

2.2.2 Inverse Laplace transform

The Spitzer identities provide the Laplace transform of the characteristic function, so to calculate the option price using Eq. (3) we must apply the inverse Laplace transform. We implement the numerical scheme by Abate and Whitt (1995), which uses the trapezoidal rule to approximate the analytic expression for the inverse Laplace transform shown in Eq. (11) with

$$
f(t) \approx \frac{e^{A/2}}{2t} \mathrm{Re} \widetilde{f}\left(\frac{A}{2t}\right) + \frac{e^{A/2}}{t} \sum_{k=1}^{\infty} (-1)^k \mathrm{Re} \widetilde{f}\left(\frac{A + 2k\pi i}{2t}\right),
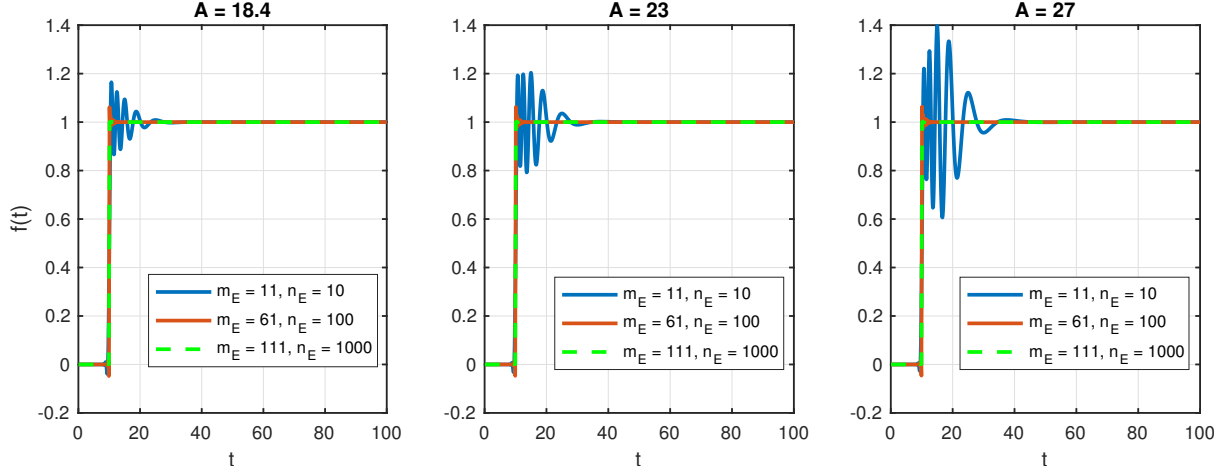\tag{29}
$$

9

Figure 2: Output of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing n_{E} and m_{E} reduces the size of the oscillations, but it is not improved by increasing A.

where $\widetilde{f}\left(\frac{A+2k\pi i}{2t}\right)$ is the Laplace transform $\widetilde{f}(s)$ with $s = \frac{A+2k\pi i}{2t}$. The value of A is selected to control the accuracy of the approximation; for an accuracy of $10^{-\gamma}$ we must select $A = \gamma \log(10)$. We then use the Euler transform to accurately approximate this infinite series. First the partial sums

$$b_k = \frac{e^{A/2}}{2t}\mathrm{Re}\widetilde{f}\left(\frac{A}{2t}\right) + \frac{e^{A/2}}{t}\sum_{j=1}^{k}(-1)^j\mathrm{Re}\widetilde{f}\left(\frac{A+2j\pi i}{2t}\right) \tag{30}$$

are calculated for $k = n_{\mathrm{E}}, \ldots, n_{\mathrm{E}} + m_{\mathrm{E}}$. We then take the binomially weighted average (Euler transform) of these terms, i.e.,

$$f(t) \approx \frac{1}{2^{m_{\mathrm{E}}}}\sum_{k=0}^{m_{\mathrm{E}}}\binom{m_{\mathrm{E}}}{k}b_{n_{\mathrm{E}}+k}. \tag{31}$$

The values of n_{E} and m_{E} are selected large enough to give sufficient accuracy, but low enough to avoid unnecessary computational effort. Numerical tests were carried out inverting the Laplace transform of a delayed unit step function $\widetilde{f}(s) = e^{-\tau s}/s$ where the delay $\tau = 10$. This is an extreme test case as the step function has a jump discontinuity and Abate and Whitt (1992a) state that the performance bound of $10^{-\gamma} = e^{-A}$ does not apply in the presence of jumps. However it is important to consider the performance of the inverse Laplace transform with discontinuities in the time domain as the value of the contracts that we are pricing will abruptly become zero on expiry. The recovered functions for different values of A, n_{E} and m_{E} are shown in Figure 2 and the errors are shown in Figures 3 and 4. The empirical results in Figure 4 show that we can select values for A, m_{E} and n_{E} so that, away from the discontinuity, the performance matches the bound of $10^{-\gamma} = e^{-A}$ specified by Abate and Whitt. Furthermore, we show in Sections 3 and 4 that the error bounds and observed results for the pricing procedure are limited by the performance of the sinc-based Hilbert transform. Therefore, we can use the Abate and Whitt inverse Laplace transform method to price mid- to long-dated options.

We base the selection of the parameters for the inverse Laplace transform on the empirical results. From Figures 2 and 3 we can see that the size of the oscillations due to the disconti-
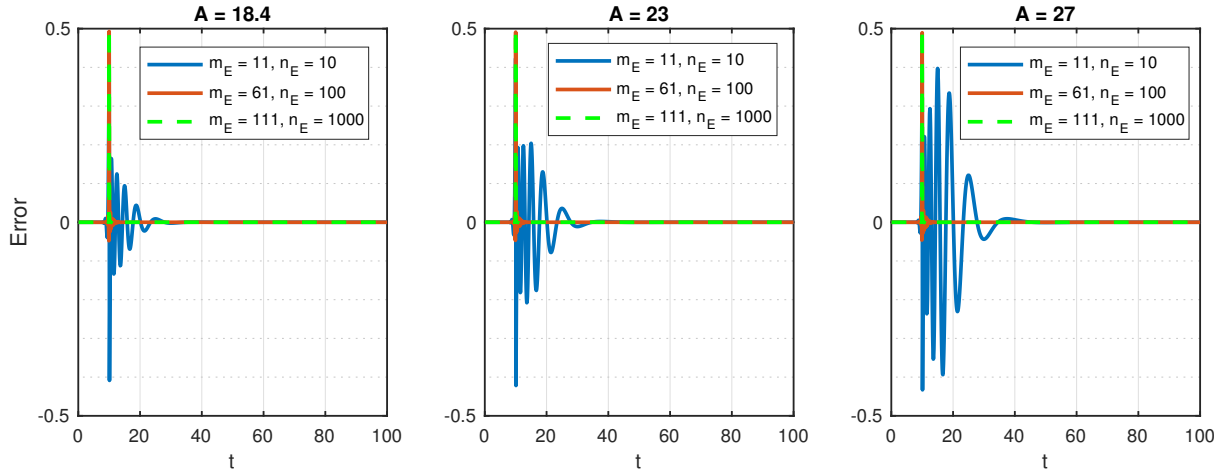
Figure 3: Error of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing n_{E} and m_{E} reduces the size of the errors due to the oscillations, but it is not improved by increasing A.
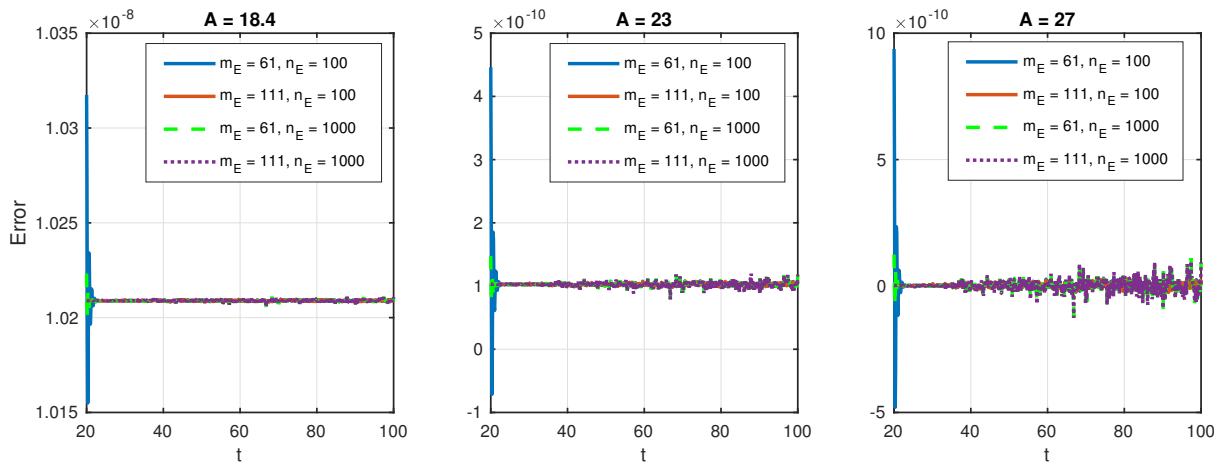

Figure 4: Error of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing A decreases the error floor, while the latter is unaffected by increasing n_{E} and m_{E}. The noise on the error floor is $\approx 10^{-10}$.

nuity are predominantly affected by m_E and n_E. The error floor is controlled by A; the values of 18.4, 23 and 27 in Figures 2–4 correspond to errors of approximately 10^{-8}, 10^{-10} and 10^{-12} respectively. However, Figure 4 shows that the noise around the error floor is $\approx 10^{-10}$ and therefore there is no advantage in selecting values of A larger than 23. For the pricing calculations we use $A = 23$, $m_\mathrm{E} = 61$ and $n_\mathrm{E} = 100$ which give a combination of high accuracy and fast computation time.

2.2.3   Pricing procedure: single-barrier options

We describe the pricing procedure for single-barrier down-and-out options as an example, but the use of the Spitzer identities is equally applicable to other types of barrier options and also to lookback options; the pricing formulae described by Green et al. (2010) include methods for single-barrier up-and-out and knock-in options. The pricing method is adapted from the scheme by Fusai et al. (2016) and Phelan et al. (2017) using the relationship between $\Phi(\xi, q)$ and $\Phi_\mathrm{c}(\xi, s)$ described in Section 2.1.1.

1. Compute the characteristic exponent $\psi(\xi + i\alpha)$, where α is the damping parameter introduced in Section 2, Eq. (4).

2. Use the Plemelj-Sokhotsky relations with the sinc-based Hilbert transform to factorise

$$\Phi_\mathrm{c}(\xi, s) := s - \psi(\xi + i\alpha) = \Phi_{\mathrm{c}+}(\xi, s)\Phi_{\mathrm{c}-}(\xi, s) \tag{32}$$

   for all $s = \frac{A + 2k\pi i}{2t}$ required for the inverse Laplace transform in Eq. (30).

3. Decompose with respect to l

$$P_\mathrm{c}(\xi, s) := \frac{\sigma(\xi/\xi_{\max})}{\Phi_{\mathrm{c}-}(\xi, s)} = P_{cl+}(\xi, s) + P_{cl-}(\xi, s), \tag{33}$$

   and calculate

$$\widetilde{\widehat{p}}(\xi, s) := \frac{P_{cl+}(\xi, s)}{\Phi_{\mathrm{c}+}(\xi, s)}, \tag{34}$$

   where $\sigma(\xi/\xi_{\max})$ is an exponential filter of order p (see Section 3.2).

4. Calculate the option price as

$$v(0, 0) := \mathcal{F}_{\xi \to x}^{-1} \left[\widehat{\phi}^*(\xi) \mathcal{L}_{s \to T}^{-1} \widetilde{\widehat{p}}(\xi, s) \right](0), \tag{35}$$

   where $\widehat{\phi}^*(\xi)$ is the complex conjugate of the Fourier transform of the damped payoff function given in Eq. (5).

2.2.4   Pricing procedure: double-barrier options

The pricing procedure for double-barrier options is very similar to the method for the single-barrier options described in Section 2.2.3, 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.

12

The steps in the pricing procedure are the same as those for single-barrier down-and-out options described in Section 2.2.3 with the exception of Step 3 which is now replaced by the fixed-point algorithm

3 (a) Set $J_{cu+}(\xi, s) = J_{cl-}(\xi, s) = 0$.

(b) Decompose with respect to l

$$P_c(\xi, s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right) \frac{1 - \Phi_{c+}(\xi, s) J_{cu+}(\xi, s)}{\Phi_{c-}(\xi, s)} = P_{cl+}(\xi, s) + P_{cl-}(\xi, s), \qquad (36)$$

and set $J_{cl-}(\xi, s) := P_{cl-}(\xi, s)$.

(c) Decompose with respect to u

$$Q_c(\xi, s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right) \frac{1 - \Phi_{c-}(\xi, s) J_{cl-}(\xi, s)}{\Phi_{c+}(\xi, s)} = Q_{cu+}(\xi, s) + Q_{cu-}(\xi, s), \qquad (37)$$

and set $J_{cu+}(\xi, s) := Q_{cu+}(\xi, s)$.

(d) Calculate

$$\widetilde{\widehat{p}}(\xi, s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right) \frac{1 - \Phi_{c-}(\xi, s) J_{cl-}(\xi, s) - \Phi_{c+}(\xi, s) J_{cu+}(\xi, s)}{\Phi_c(\xi, s)}. \qquad (38)$$

(e) If the difference between the new and the old value of $\widetilde{\widehat{p}}(\xi, s)$ is less than a predefined tolerance or the number of iterations is greater than a certain threshold then continue, otherwise return to step (b). Numerical tests have shown that an iteration threshold of 5 is sufficient, as higher values do not yield improvements.

3 Error convergence of the pricing procedure

We examine the performance of each stage of the pricing procedure and discuss the respective error bounds. In addition, the effect of each step on the shape of the output function in the Fourier domain is investigated, as this influences the error convergence of later steps. Stenger (1993) showed that the discretisation error in Eq. (25) is exponentially convergent when the function $f(\xi)$ is analytic in a strip of the complex plane including the real axis. Therefore the error calculations here concern the truncation error from the approximation in Eq. (26). The truncation error using the sinc-based Hilbert transform depends on the behaviour of the characteristic exponent as $|\xi| \to \infty$: Table 1 shows the characteristic exponents of five Lévy processes. The damping parameter α is omitted to make the notation more concise, which is appropriate as its value becomes insignificant as $|\xi| \to \infty$.

3.1 Factorisation

After calculating the characteristic exponent, the next step in the pricing procedure is the numerical factorisation of $\Phi_c(\xi, s) = s - \psi(\xi)$. In order to understand the error convergence we must consider the way that the function behaves for large values of $|\xi|$. The characteristic

Process	Characteristic exponent $\psi(\xi)$
Normal	$i\xi\mu - \frac{1}{2}\sigma^2\xi^2$
Kou	$i\xi\mu - \frac{1}{2}\sigma^2\xi^2 + \lambda\left(\frac{(1-\rho)\eta_2}{\eta_2+i\xi} + \frac{\rho\eta_1}{\eta_1-i\xi}\right)$
Merton	$i\xi\mu - \frac{1}{2}\sigma^2\xi^2 + \lambda\left(e^{i\alpha\xi-\frac{1}{2}\delta^2\xi^2} - 1\right)$
NIG	$\delta\left(\sqrt{\alpha^2 - (\beta + i\xi)^2} - \sqrt{\alpha^2 - \beta^2}\right)$
VG	$-\frac{1}{\nu}\log\left(1 - i\xi\theta\nu + \frac{1}{2}\nu\sigma^2\xi^2\right)$

Table 1: Characteristic exponent of some Lévy processes.

exponents of the processes listed in Table 1 will take high negative values which will dominate $\Phi_c(\xi, s)$ so that as $|\xi| \to \infty$ we can approximate $s - \psi(\xi) \sim -\psi(\xi)$. The function to be decomposed in the factorisation stage is therefore $\sim \log[-\psi(\xi)]$. This is an increasing function in $|\xi|$, so the bounds for the truncation error of the sinc-based Hilbert transform (Feng and Linetsky, 2008, Theorems 6.4–6.6) cannot be used. Moreover, if we consider the truncation errors from Eq. (26) for positive and negative values of k individually, we obtain two infinite summations that do not converge. However, Table 1 shows that as $|\xi| \to \infty$ the values of $\psi(\xi)$ and $\psi(-\xi)$ will become increasingly similar. We can exploit this similarity to find a bound by combining the positive and negative truncations: the truncation error of $f(\xi) = \mathcal{H}[\log \Phi_c(\xi, s)]$ is bounded as

$$
\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< \Delta\xi \sum_{k<-M/2} \frac{\log \Phi_c(k\Delta\xi, s)}{\pi(\xi - k\Delta\xi)} + \Delta\xi \sum_{k>M/2} \frac{\log \Phi_c(k\Delta\xi, s)}{\pi(\xi - k\Delta\xi)} \\
&= \Delta\xi \sum_{k>M/2} \left(\frac{\log \Phi_c(k\Delta\xi, s)}{\pi(\xi - k\Delta\xi)} + \frac{\log \Phi_c(-k\Delta\xi, s)}{\pi(\xi + k\Delta\xi)}\right) \\
&= \frac{\Delta\xi}{\pi} \sum_{k>M/2} \frac{\xi\left(\log \Phi_c(k\Delta\xi, s) + \log \Phi_c(-k\Delta\xi, s)\right)}{\xi^2 - k^2\Delta\xi^2} \\
&\quad + \frac{\Delta\xi}{\pi} \sum_{k>M/2} \frac{k\Delta\xi\left(\log \Phi_c(k\Delta\xi, s) - \log \Phi_c(-k\Delta\xi, s)\right)}{\xi^2 - k^2\Delta\xi^2},
\end{aligned} \tag{39}
$$

where $f_{\Delta\xi}(\xi)$ is the value of the infinite summation in Eq. (25) and $f_{\Delta\xi,M}(\xi)$ is the result of the truncated summation in Eq. (26).

The next step in bounding the error convergence is to show that the expression in Eq. (39) is dominated by the first sum as $M \to \infty$. As $\psi(k\Delta\xi) \sim \psi(-k\Delta\xi)$ for $k \to \infty$, $\log \Phi_c(k\Delta\xi, s) - \log \Phi_c(-k\Delta\xi, s) \to 0$ as $k \to \infty$. However, $k\Delta\xi$ is also present in the numerator and increases linearly with k. By determining the rate of decrease of $\log \Phi_c(k\Delta\xi, s) - \log \Phi_c(-k\Delta\xi, s)$, we show that the second term is bounded as $O(1/k^2)$ and therefore the first term dominates Eq. (39). We then calculate a bound for the error based on the first summation term in Eq. (39). These steps are carried out in a slightly different way depending on the form of the characteristic exponents shown in Table 1.

3.1.1 Normal, Merton and Kou processes

For the normal, Merton and Kou processes, when $k \to \infty$, $\Phi_c(k\Delta\xi)$ becomes dominated by $\sigma(k\Delta\xi)^2 - i\mu k\Delta\xi$ as shown in Table 1. The parameters μ and σ are specific to the distribution. We can therefore approximate the second expression in the summation by

$$
\begin{aligned}
\frac{k\Delta\xi\big(\log\Phi_c(k\Delta\xi,s) - \log\Phi_c(-k\Delta\xi,s)\big)}{\xi^2 - k^2\Delta\xi^2} &= \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\log\frac{\Phi_c(k\Delta\xi,s)}{\Phi_c(-k\Delta\xi,s)} \\
&\sim \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\log\frac{\sigma^2(k\Delta\xi)^2/2 + i\mu(k\Delta\xi)}{\sigma^2(k\Delta\xi)^2/2 - i\mu(k\Delta\xi)} \\
&= \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\log\frac{1 + 2i\mu/(\sigma^2 k\Delta\xi)}{1 - 2i\mu/(\sigma^2 k\Delta\xi)}.
\end{aligned}
\tag{40}
$$

The logarithm in Eq. (40) is of the form $\log\frac{1+x}{1-x}$ where $x = \frac{2i\mu}{\sigma^2 k\Delta\xi}$. For $x \to 0$, $\log\frac{1+x}{1-x} \sim 2x$, thus

$$
\frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\log\frac{1 + 2i\mu/(\sigma^2 k\Delta\xi)}{1 - 2i\mu/(\sigma^2 k\Delta\xi)} \sim \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\frac{4i\mu}{\sigma^2 k\Delta\xi} = \frac{4i\mu}{\sigma^2(\xi^2 - k^2\Delta\xi^2)}
\tag{41}
$$

gives an approximation for the second term in Eq. (39). Due to the denominator, this is $O(1/k^2)$. Thus, as $\log\Phi_c(k\Delta\xi,s) + \log\Phi_c(-k\Delta\xi,s)$ is increasing in k, the error is indeed dominated by the first term in Eq. (39).

For the normal, Kou and Merton processes, $\Phi_c(k\Delta\xi,s)$ and $\Phi_c(-k\Delta\xi,s) \to 2\log|k\Delta\xi|$ as $k \to \infty$. Therefore, the error bound is

$$
\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< \frac{c\Delta\xi}{\pi}\sum_{k>M/2}\frac{\log\Phi_c(k\Delta\xi,s)}{\xi^2 - k^2\Delta\xi^2}, \\
&< c_1\Delta\xi\sum_{k>M/2}\frac{\log(k^2\Delta\xi^2)}{k^2\Delta\xi^2},
\end{aligned}
\tag{42}
$$

where c and c_1 are some constants. Here, as Eq. (42) gives the error at fixed values of ξ, i.e. the chosen grid points, the ξ can be absorbed into c. However, as M increases, our range of ξ values increases. Therefore, as there is a linear dependence of the error bound on ξ, we should consider the effect of errors at large values of ξ on the error of the final solution. In doing this we can take account of the shape of the output from the factorisation $\Phi_{c\pm}(\xi,s)$ which decays as $|\xi| \to \infty$ and the use of filtering on the input to the next step as described in Section 3.2. These effects combine to mean that the error as a proportion of the value of $\Phi_{c\pm}(\xi,s)$ at high $|\xi|$ is less significant to the error of the overall solution than the relationship between the value of M and the error in $\Phi_{c\pm}(\xi,s)$ for low values of $|\xi|$. Approximating the summation by an integral
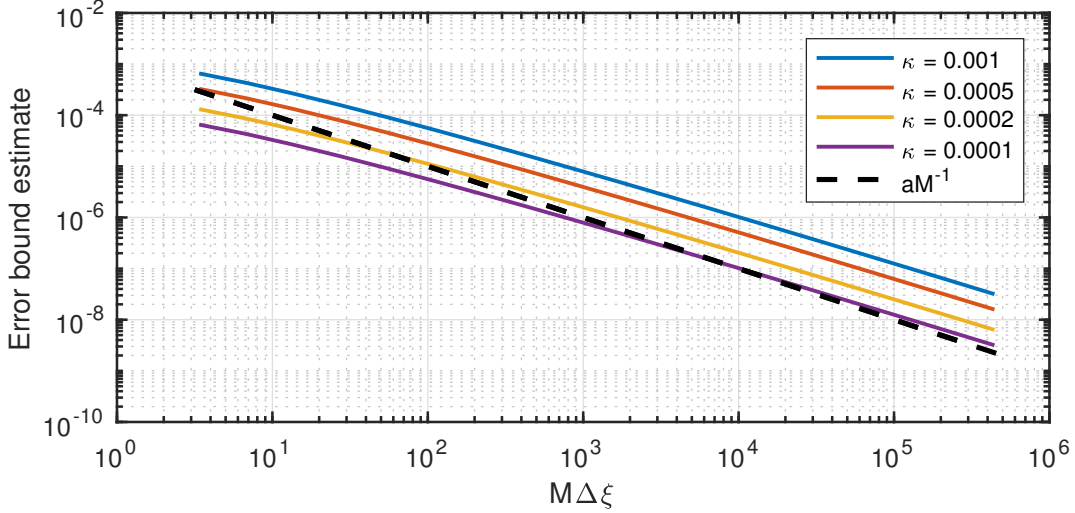
Figure 5: Eq. (44) plotted for different values of κ to show the estimate of the error bound on the sinc-based numerical factorisation of $\Phi_c(\xi, s)$. Notice that the predicted error bound from the factorisation has a decay that increases in slope as $M\Delta\xi$ increases and is slightly shallower than $O(1/M)$ for the values of M which we are using. Sections 3.1.1, 3.1.2 and 3.1.3 show that this bound applies for the normal, NIG, Kou, Merton and VG processes.

with $M' = M/2$, we obtain

$$
\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M'}(\xi)| &< c_2 \int_{M'\Delta\xi}^{+\infty} \frac{\log\xi'}{\xi'^2}d\xi' \\
&= c_2 \left[\frac{\log\xi'}{\xi'} + \frac{1}{\xi'}\right]_{+\infty}^{M'\Delta\xi} \\
&= c_2 \left[\frac{\log M'\Delta\xi}{M'\Delta\xi} + \frac{1}{M'\Delta\xi}\right]
\end{aligned}
\tag{43}
$$

where c_2 is some constant. Having applied the sinc-based discrete Hilbert transform we can calculate the positive and negative functions using the Plemelj-Sokhotsky relations and then exponentiate the results to obtain the Wiener-Hopf factors. Therefore, using the expression in Eq. (43), we can bound the truncation error of the Wiener-Hopf factors, and by extension the total error as the truncation error dominates, as

$$
\left|\frac{\Phi_{\Delta\xi,c\pm}(\xi) - \Phi_{\Delta\xi,M',c\pm}(\xi)}{\Phi_{\Delta\xi,c\pm}(\xi)}\right| < \left|1 - (eM'\Delta\xi)^{\frac{\kappa}{M'\Delta\xi}}\right|,
\tag{44}
$$

where κ is some constant. Here, $\Phi_{\Delta\xi,c\pm}(\xi)$ denotes the (theoretical) Wiener-Hopf factors generated using the series in Eq. (25) and $\Phi_{\Delta\xi,M',c\pm}(\xi)$ denotes the Wiener-Hopf factors calculated using the truncated summation in Eq. (26).

Figure 5 shows Eq. (44) plotted against $M\Delta\xi$ for different values of κ. This demonstrates that the predicted error bound from the factorisation has a decay that increases in slope as $M\Delta\xi$ increases and is slightly shallower than $O(1/M)$ for the values of M which we are using.

16

3.1.2 Normal inverse Gaussian process

In the case of the NIG process the characteristic exponent is

$$\psi(\xi) = \delta\left(\sqrt{\alpha^2 - (\beta + i\xi)^2} - \sqrt{\alpha^2 - \beta^2}\right). \tag{45}$$

The presence of a square root around the $i\xi$ and ξ^2 terms means that as $|k| \to \infty$, the equivalent expression to the logarithm in Eq. (40) is $\frac{1}{2}\log\frac{1+2i\beta/(k\Delta\xi)}{1-2i\beta/(k\Delta\xi)}$. Furthermore, $\Phi_c(k\Delta\xi, s)$ and $\Phi_c(-k\Delta\xi, s)$ become dominated by $\log|k\Delta\xi|$ as $|k| \to \infty$. Therefore the only difference between the truncation error bound for the NIG process and the result in Eq. (43) is the size of the constants.

3.1.3 Variance gamma process

The characteristic function of the VG process is

$$\psi(\xi) = -\frac{1}{\nu}\log\left(1 - i\xi\theta\nu + \frac{1}{2}\nu\sigma^2\xi^2\right). \tag{46}$$

This is significantly different from the other characteristic exponents that we have considered, being the log of a polynomial. Similarly to the previous methods, we show that as $k \to \infty$ the decrease rate of $\log\frac{\Phi_c(k\Delta\xi, s)}{\Phi_c(-k\Delta\xi, s)}$ is at least $O(1/k)$:

$$\begin{aligned}
\log\frac{\Phi_c(k\Delta\xi, s)}{\Phi_c(-k\Delta\xi, s)} &\sim \log\frac{\log(-ik\Delta\xi\theta\nu + \nu\sigma^2(k\Delta\xi)^2/2)}{\log(ik\Delta\xi\theta\nu + \nu\sigma^2(k\Delta\xi)^2/2)} \\
&= \log\frac{\log\left(1 - \frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2}\right) + \log\frac{\nu\sigma^2(k\Delta\xi)^2}{2}}{\log\left(1 + \frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2}\right) + \log\frac{\nu\sigma^2(k\Delta\xi)^2}{2}} \\
&\sim \log\frac{-\frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2} + \log\frac{\nu\sigma^2(k\Delta\xi)^2}{2}}{\frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2} + \log\frac{\nu\sigma^2(k\Delta\xi)^2}{2}} \\
&= \log\frac{1 - \frac{2i\theta}{\sigma^2 k\Delta\xi\log(\nu\sigma^2(k\Delta\xi)^2/2)}}{1 + \frac{2i\theta}{\sigma^2 k\Delta\xi\log(\nu\sigma^2(k\Delta\xi)^2/2)}} \\
&\sim \frac{-4i\theta}{\sigma^2 k\Delta\xi\log\left(\nu\sigma^2(k\Delta\xi)^2/2\right)}. \tag{47}
\end{aligned}$$

This decreases quicker than $O(1/k)$ and thus Eq. (39) is dominated by the first term. The equivalent expression to Eq. (42) for the VG process is

$$|f_{\Delta\xi} - f_{\Delta\xi,M}| < \frac{c\Delta\xi}{\pi}\sum_{k>M/2}\frac{\log\Phi_c(k\Delta\xi, s)}{\xi^2 - k^2\Delta\xi^2}$$

$$< c_1\Delta\xi\sum_{k>M/2}\frac{\log\log(k^2\Delta\xi^2)}{k^2\Delta\xi^2}. \tag{48}$$

As $|\log(\log x)|$ is bounded by $|\log x|$ as $x \to \infty$, the factorisation error of the method with the VG process is also bounded by Eq. (44).
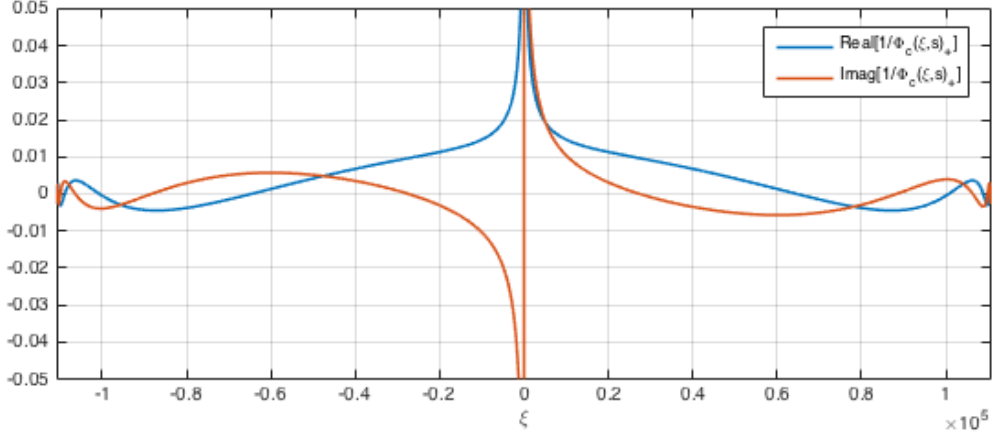
Figure 6: Example plot of the real and imaginary parts of $\Phi_{c+}(\xi, s)$ plotted against ξ with $s = A/(2t)$, as specified for the Abate and Whitt inverse Laplace transform. Notice that although the value of $|\Phi_{c+}(\xi, s)|$ is bounded by a constant as $|\xi| \to \infty$, the rate of decay is very slow and we have not been able to determine a decreasing bound.

3.2 Decomposition error

The output of the factorisation is shown in Figure 6. The next step in the calculation is to find the positive part with respect to l of $P_c(\xi, s) = \frac{1}{\Phi_{c-}(\xi, s)}$. We can attempt to bound the truncation error of this calculation by combining the errors from the positive and negative truncations as before:

$$
\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &= \frac{\Delta\xi}{\pi} \left| \sum_{k>M/2} \frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi} + \sum_{k<-M/2} \frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi} \right| \\
&= \frac{\Delta\xi}{\pi} \left| \sum_{k>M/2} \frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi} + \frac{P_c(-k\Delta\xi)}{\xi + k\Delta\xi} \right| \\
&= \frac{\Delta\xi}{\pi} \left| \sum_{k>M/2} \frac{\xi[P_c(k\Delta\xi) + P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2} + \frac{k\Delta\xi[P_c(k\Delta\xi) - P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2} \right| .
\end{aligned}
\tag{49}
$$

Figure 6 shows that for high $|\xi|$, $|P_c(\xi) - P_c(-\xi)|$ is bounded from above by a constant. However, we do not have a decreasing bound for $|P_c(\xi) - P_c(-\xi)|$. Therefore we can only bound the second term in Eq. (49) as

$$
\sum_{k>M/2} \frac{k\Delta\xi[P_c(k\Delta\xi) - P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2} < c \sum_{k>M/2} \frac{k\Delta\xi}{\xi^2 - (k\Delta\xi)^2}
\tag{50}
$$

where c is some positive constant; this does not converge. We can compare it with the discretely monitored version from Fusai et al. (2016) where the first date is taken out of the scheme, meaning that the function undergoing decomposition is multiplied with the characteristic function. For processes other than VG, this restores the exponential slope of the function for high values

18

of ξ which again means that the truncation error of the sinc-based discrete Hilbert transform is exponentially bounded. To improve the error of the decomposition in the continuously monitored case we can improve the slope of the function on the input to the Hilbert transform by using a spectral filter. We use an exponential filter which has previously achieved good results in option pricing applications (Ruijter et al., 2015; Phelan et al., 2017). The filter is described by Eq. (28) and its shape is shown in Figure 1. Numerical tests have shown that the use of this filter improves the error of the decomposition step so that it no longer limits the error convergence of the pricing scheme. However, the overall error of the pricing procedure will be continue to be limited by the error from the initial factorisation step as described in Eq. (44) and shown in Figure 5.

4 Results

We present results for the Spitzer-Laplace pricing procedure for continuously monitored single and double-barrier options for the NIG, Kou and VG processes. We also show that the error convergence represents a limiting case of the performance of the FGM method for discretely monitored options as $N \to \infty$ and $\Delta t \to 0$, where N is the number of monitoring dates and Δt is the time interval between them.

4.1 Results for Spitzer-Laplace method for continuously monitored options

The error convergence for single-barrier down-and-out options is shown in Figure 7. Figure 8 shows the results for double-barrier options. The computed prices are given in Tables 2 and 3 for single and double-barrier options respectively. Although the absolute error is worse for double-barrier options, the speed of error convergence is very similar for all cases and is slightly worse than $O(1/M)$, which concurs with the simulated results for the factorisation error shown in Figure 5. The details of the contract and underlying processes are shown in Table 6 in Appendix A.
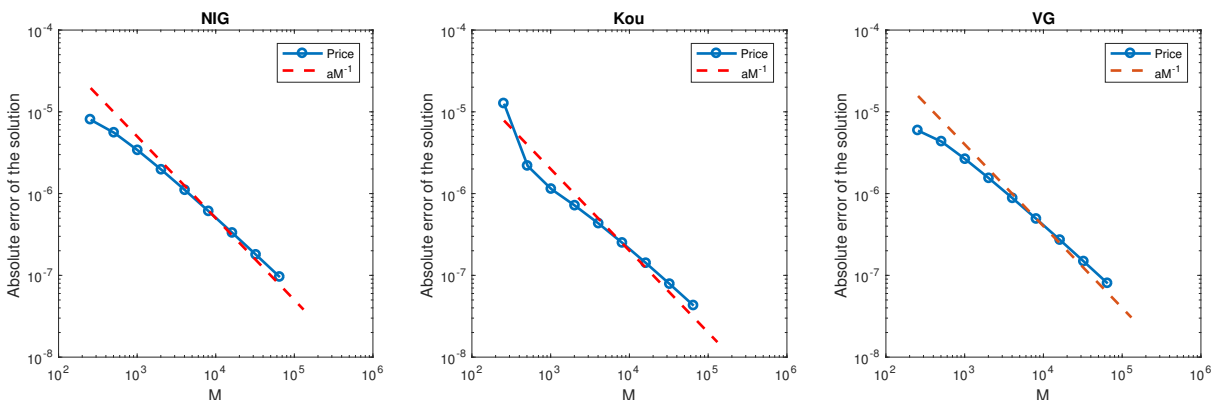

Figure 7: Error convergence for a continuously monitored single-barrier option.The error convergence conforms to the calculated error bound (only the first point of the Kou process deviates slightly from the overall behaviour) and shows the typical sub-polynomial error convergence for higher values of M.
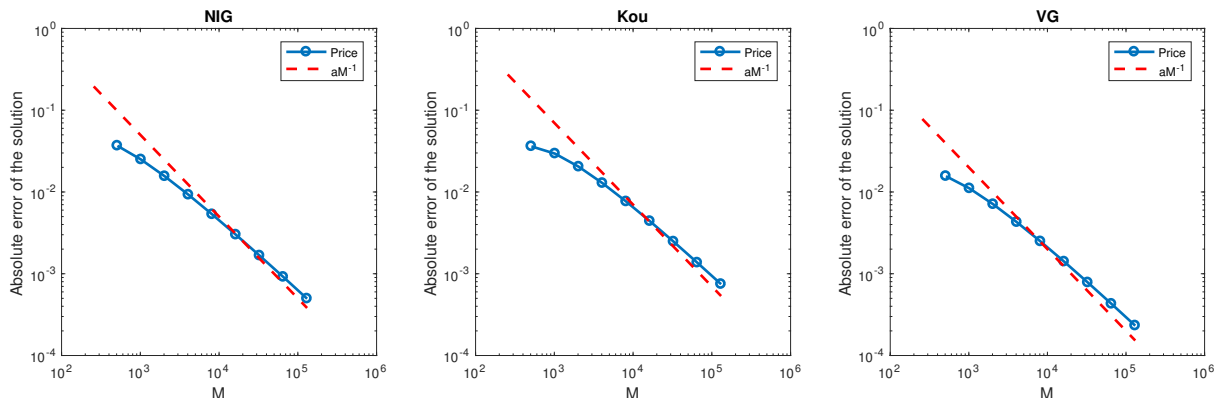
Figure 8: Error convergence for the continuously monitored double-barrier option. The absolute error is worse than that for the single-barrier option but the error convergence conforms to the calculated error bound.

Process	Price
Normal inverse Gaussian	4.77403523401E-2
Kou	4.32042632202E-2
Variance gamma	4.70627023105E-2

Table 2: Prices calculated for single-barrier options with the contract details and process parameters described in Table 6 in Appendix A and $M = 2^{17}$.

Process	Price
Normal inverse Gaussian	2.78787488E-2
Kou	3.30368034E-2
Variance gamma	2.82666693E-2

Table 3: Prices calculated for double-barrier options with the contract details and process parameters described in Table 6 in Appendix A and $M = 2^{17}$.

4.2 Comparison with the error convergence of Spitzer-z pricing method for discretely monitored options

In Section 2.1.1 we described the relationship between the Fourier-Laplace based method for continuously monitored options and the FGM method, based in the Fourier-z domain, for discretely monitored options. The latter method, as measured for a single barrier in Fusai et al. (2016) and double barriers in Phelan et al. (2017), with the number of monitoring dates up to $N \approx 10^3$, is exponentially convergent with the number of grid points for the NIG and Kou processes and better than second order polynomially convergent for the VG process. Therefore we investigate the performance of the discretely monitored method with a very high number of dates (i.e. $\Delta t \to 0$), to better understand the difference in performance between the two pricing schemes.

In Green et al. (2010) the error between the discretely and continuously monitored prices was shown to be bounded as $O(1/\sqrt{N})$, where N is the number of monitoring dates. We therefore also consider whether lower errors might be achieved by approximating the price for a continuously monitored option with the price for a discretely monitored option with a very high number of monitoring dates.

We use the same implementation as the one described in Fusai et al. (2016) for single-barrier options and Phelan et al. (2017) for double-barrier options, although the maximum number of monitoring dates is far higher than would ever be used for discretely monitored options in practice. Due to the $O(1/\sqrt{N})$ error bound between the prices for continuously and discretely monitored options, we must select a very large number of monitoring dates in order for this effect to be less significant than the error from the continuously monitored pricing method. The error convergence of the discrete pricing method as $N \to \infty$ (or $\Delta t \to 0$) is shown in Figures 10 and 11. The results show that as $\Delta t \to 0$, the error convergence for discrete monitoring degrades until it approaches that of continuously monitored options. Moreover, it demonstrates that, rather than being an anomaly, the error convergence of the continuously monitored method is consistent with that of the discretely monitored method as $\Delta t \to 0$. This can be understood by considering how $\Psi(\xi, \Delta t)$ changes with Δt for the discrete example. As $\Delta t \to 0$, $\Psi(\xi, \Delta t) = e^{\psi(\xi)\Delta t}$ decays to 0 more and more slowly as $|\xi| \to \infty$. Therefore the error convergence of the pricing technique for continuously monitored barrier options is a limit of the error convergence for the discrete case as $\Delta t \to 0$.

Computed prices for continuously and discretely monitored options are plotted against M in Figures 12 and 13. In addition, computation times of the pricing methods for the discrete and continuously monitored methods are shown in Tables 4 and 5. Figures 12 and 13 show that, as expected, the larger the number of monitoring dates the closer the price is to the continuously monitored price. However, they also show that the direction of convergence depends on the type of option and the process being used. Therefore, in order to obtain the CPU times in Tables 4 and 5 we take the lowest time where the convergence error for the discretely monitored method is significantly lower (i.e. ten times) than the error compared to the price for the continuously monitored case with maximum M. This shows that for relatively high errors, $\approx 10^{-4}$ for a single barrier and $\approx 10^{-2}$ for double barriers, the discretely monitored method is slightly quicker. However, the discretely monitored method is unable to achieve the lower

errors, $\approx 10^{-6}$ for a single barrier and $\approx 10^{-4}$ for double barriers, which are attained by the continuously monitored method and therefore is not a sufficiently accurate approximation.

Can we can achieve a better approximation of the continuous method by increasing the number of monitoring dates further? Previous literature, e.g. Green et al. (2010), has shown that the convergence of the discrete method to the continuous method with increasing monitoring dates is $O(1/\sqrt{N})$. From Figure 9 we can observe that, although the discrete method with the Kou process does indeed have this rate of convergence, it achieves approximately $O(1/N)$ with the NIG and VG processes. Therefore, if we wished to decrease the error of the discrete approximation so that it is significantly (i.e. ten times) less than the continuous case then we would have to increase the maximum number of monitoring dates in Tables 4 and 5 by 100 times for the NIG and VG processes and by 200^2 times for the Kou process. We can see from Figures 10 and 11 that at $M = 2^{17}$ the discrete methods with these numbers of dates have an error which is worse than the required accuracy of ten times better than the continuous method. Thus the only possibility would be to also increase M, and by extension the CPU time of the discrete monitoring method, causing its computational cost to be greater than that for continuous monitoring.

Process	Dates	Error	Time	M
and	cont	3.21E-04	0.07	1024
	1008	1.60E-04	0.07	4096
	cont	1.86E-04	0.13	2048
	cont	5.74E-05	0.50	8192
	cont	1.69E-05	2.03	32768
	252	2.31E-04	0.04	2048
	cont	2.57E-04	0.06	1024
	504	1.13E-04	0.11	4096
VG	cont	1.49E-04	0.17	2048
	1008	5.29E-05	0.14	8192
	cont	4.74E-05	0.49	8192
	cont	1.43E-05	2.01	32768
	cont	4.19E-06	11.52	131072
	252	3.02E-04	0.01	512
	cont	1.57E-04	0.02	256
	504	2.03E-04	0.01	512
Kou	cont	1.57E-04	0.02	256
	1008	1.47E-04	0.02	1024
	cont	1.57E-04	0.02	256
	cont	6.69E-05	0.12	2048
	cont	2.35E-05	0.49	8192
	cont	7.36E-06	2.07	32768

Table 4: CPU times and errors for the continuously monitored method and the discretely monitored method as an approximation to continuous monitoring for the single-barrier case. The CPU times for the discretely monitored price are chosen for the grid size M which gives the lowest CPU time where the convergence error is significantly (about ten times) lower than the error compared to the continuously monitored price.
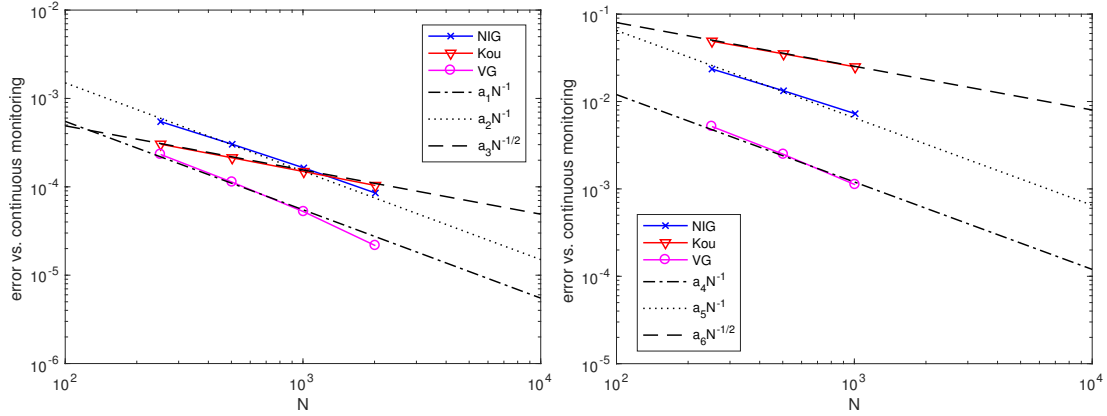
Figure 9: Error for discretely monitored barrier options used as an approximation of the price for the continuously monitored case, plotted as a function of the number of monitoring dates. The error is calculated as the difference between the prices for discrete and continuous monitoring at the maximum grid size of 2^{17} Results for single and double-barrier options are displayed on the left and right hand plots respectively. Notice that the error for the Kou process converges as $O(1/\sqrt{N})$, whereas the error for the NIG and VG processes converges a rate of approximately $O(1/N)$.
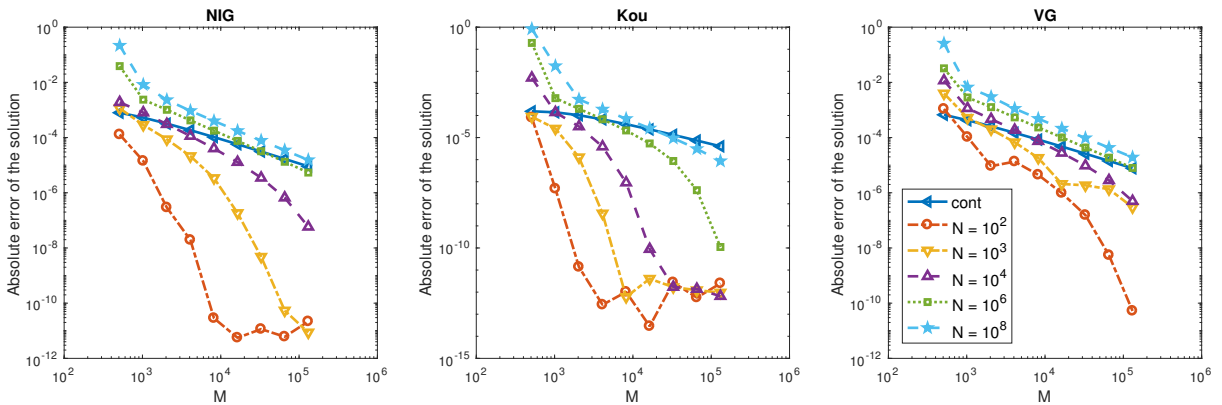

Figure 10: Error as a function of the grid size M for continuously monitored single-barrier options compared to discretely monitored options as the number of monitoring dates N increases. The error for each number of dates (including continuous monitoring) is calculated against the price for the same number of dates with 2^{18} grid points. For all processes, as $\Delta t \to 0$ the slope of the error convergence of the discretely monitored scheme approaches that of the continuously monitored scheme.
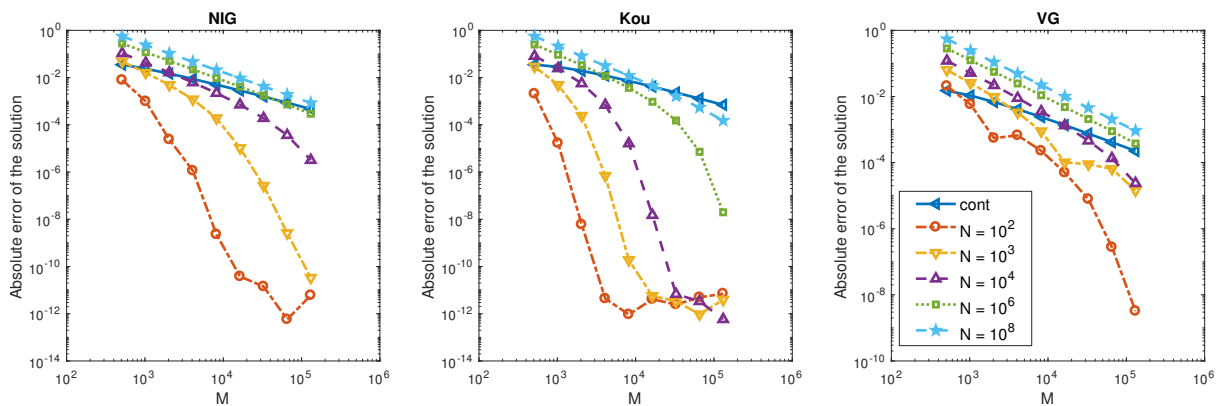
Figure 11: Error as a function of the grid size M for continuously monitored double-barrier options compared to discretely monitored options as the number of monitoring dates N increases. The error for each number of dates (including continuous monitoring) is calculated against the price for the same number of dates with 2^{18} grid points. For all processes, as $\Delta t \to 0$ the slope of the error convergence of the discretely monitored scheme approaches that of the continuously monitored scheme.
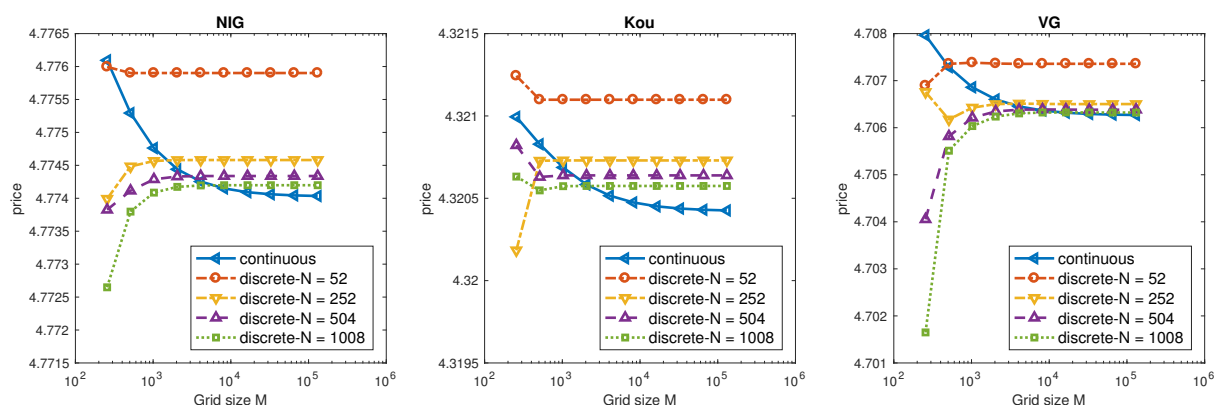

Figure 12: Price plotted against the grid size M for continuously monitored single-barrier options compared to discretely monitored options as the number of monitoring dates N increases. Note that the larger the value of N, the closer the price is to the continuously monitored option price.

Process	Dates	Error	Time	M
	252	2.35E-02	0.16	4096
	cont	2.38E-02	0.11	512
	504	1.32E-02	0.14	4096
NIG	cont	1.48E-02	0.24	1024
	1008	7.24E-03	0.29	8192
	cont	8.78E-03	0.54	2048
	cont	1.58E-03	5.03	16384
	cont	4.69E-04	20.94	65536
	252	5.10E-03	0.08	2048
	cont	6.86E-03	0.24	1024
	504	2.51E-03	0.13	4096
VG	cont	2.40E-03	1.15	4096
	1008	1.15E-03	0.29	8192
	cont	1.36E-03	2.44	8192
	cont	7.56E-04	5.16	16384
	cont	2.25E-04	21.21	65536
	252	4.90E-02	0.03	1024
	cont	3.54E-02	0.07	256
	504	3.51E-02	0.07	2048
Kou	cont	3.54E-02	0.07	256
	1008	2.47E-02	0.04	1024
	cont	2.84E-02	0.14	512
	cont	7.23E-03	1.19	4096
	cont	2.33E-03	5.30	16384
	cont	7.03E-04	21.05	65536

Table 5: CPU times and errors for the continuously monitored method and the discretely monitored method as an approximation to continuous monitoring for the double-barrier case. The CPU times for the discretely monitored price are chosen for the grid size M which gives the lowest CPU time where the convergence error is significantly (about ten times) lower than the error compared to the continuously monitored price.
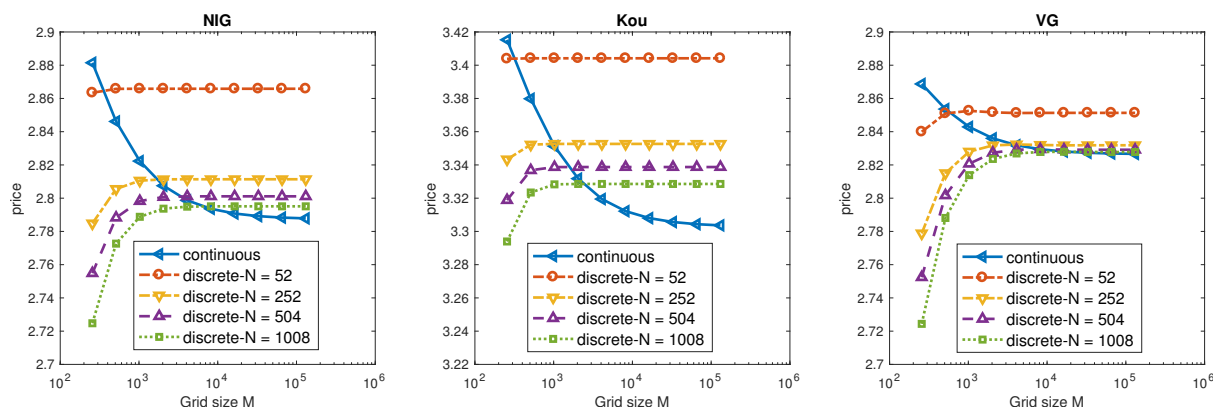

Figure 13: Price plotted against grid size M for continuously monitored double-barrier options compared to discretely monitored options as the number of monitoring dates N increases. Note that the larger the value of N, the closer the price is to the continuously monitored option price.

25

5 Conclusions

We showed that the numerical method for calculating the discretely monitored Spitzer identities described by Fusai et al. (2016) can be modified for continuous monitoring by using the Fourier-Laplace domain instead of the Fourier-z domain. We implemented this with the inverse Laplace transform by Abate and Whitt (1992a, 1995) which achieves an accuracy of approximately 10^{-11}, sufficient for our chosen application of pricing barrier options. We presented results showing that the conversion from discrete to continuous monitoring means that exponential convergence is no longer achieved, but instead the error convergence becomes sub-polynomial due to the performance of the Wiener-Hopf factorisation. By examining the effect of truncating the sinc-based discrete Hilbert transform, we were able to provide an error bound which is well matched to the observed accuracy of the pricing procedure for continuously monitored options.

It is notable that previous papers have shown that the discretely monitored case achieves exponential convergence (Fusai et al., 2016; Phelan et al., 2017) but the continuous case described here does not. However, we showed that, as the number of monitoring dates increases and $\Delta t \to 0$, the error convergence for the discretely monitored case degrades and approaches that of the continuously monitored case. Thus, the performance of the pricing technique for continuously monitored barrier options is consistent with previous results, being a limit of the error convergence for the discretely monitored case.

Furthermore we have compared the error vs. computational time of the continuously monitored scheme with that of an approximate solution generated by the discretely monitored scheme with a high number of monitoring dates. We show that, for higher errors, the discrete scheme may produce a rapidly calculated approximation to the continuously monitored scheme, but when lower errors are required the continuously monitored scheme must be used.

Acknowledgments

References

Abate, J., W. Whitt. 1992a. The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems* **10**(1–2) 5–88. doi:10.1007/BF01158520.

Abate, J., W. Whitt. 1992b. Numerical inversion of probability generating functions. *Operations Research Letters* **12**(4) 245–251. doi:10.1016/0167-6377(92)90050-D.

Abate, J., W. Whitt. 1995. Numerical inversion of Laplace transforms of probability distributions. *ORSA Journal on Computing* **7**(1) 36–43. doi:10.1287/ijoc.7.1.36.

Barndorff-Nielson, O. E. 1998. Processes of normal inverse Gaussian type. *Finance and Statistics* **2**(1) 41–68. doi:10.1007/s007800050032.

Baxter, G., M. D. Donsker. 1957. On the distribution of the supremum functional for processes with stationary independent increments. *Transactions of the American Mathematical Society* **85**(1) 73–87.

Bayer, N., O. J. Boxma. 1996. Wiener-Hopf analysis of an M/G/1 queue with negative customers and of a related class of random walks. *Queueing Systems* **23**(1) 301–316. doi: 10.1007/BF01206563.

Broadie, M., P. Glasserman, S. Kou. 1997. A continuity correction for discrete barrier options. *Mathematical Finance* **7**(4) 325–349. doi:10.1111/1467-9965.00035.

Carr, P., H. Geman, D. B. Madan, M. Yor. 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business* **75**(2) 305–332. doi:10.1086/338705.

Carr, P., D. Madan. 1999. Option valuation using the fast Fourier transform. *Journal of Computational Finance* **2**(4) 61–73.

Chi, Y., S. X. Lin. 2011. On the threshold dividend strategy for a generalized jump-diffusion risk model. *Insurance: Mathematics and Economics* **48**(3) 326–337. doi: 10.1016/j.insmatheco.2010.11.006.

Cohen, J. W. 1975. The Wiener-Hopf technique in applied probability. J. Gani, ed., *Perspectives in Probability and Statistics: Papers in Honour of M. S. Bartlett on the Occasion of His Sixty-Fifth Birthday*. Academic Press, London, 145–156.

Cohen, J. W. 1982. *The Single Server Queue*. North Holland, Amsterdam.

Cohen, M. A., D. Pekelman. 1978. LIFO inventory systems. *Management Science* **24**(11) 1150–1162. doi:10.1287/mnsc.24.11.1150.

Fang, F., C. W. Oosterlee. 2008. A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing* **31**(2) 826–848. doi:10.1137/080718061.

Fang, F., C. W. Oosterlee. 2009. Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik* **114**(1) 27–62. doi:10.1007/s00211-009-0252-4.

Feng, L., V. Linetsky. 2008. Pricing discretely monitored barrier options and defaultable bonds in Lévy process models: a Hilbert transform approach. *Mathematical Finance* **18**(3) 337–384. doi:10.1111/j.1467-9965.2008.00338.x.

Feng, L., V. Linetsky. 2009. Computing exponential moments of the discrete maximum of a Lévy process and lookback options. *Finance and Stochastics* **13**(4) 501–529. doi:10.1007/s00780-009-0096-x.

Frigo, M., S. G. Johnson. 1998. FFTW: An adaptive software architecture for the FFT. *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 3. IEEE, 1381–1384. doi:10.1109/ICASSP.1998.681704.

Fusai, G., I. D. Abrahams, C. Sgarra. 2006. An exact analytical solution for discrete barrier options. *Finance and Stochastics* **10**(1) 1–26.

Fusai, G., G. Germano, D. Marazzina. 2016. Spitzer identity, Wiener-Hopf factorisation and pricing of discretely monitored exotic options. *European Journal of Operational Research* **251**(4) 124–134. doi:10.1109/ICASSP.1998.681704.

Gottlieb, D., C. Shu. 1997. On the Gibbs phenomenon and its resolution. *SIAM Review* **39**(4) 644–668. doi:10.1137/S0036144596301390.

Grassman, W. K. 1990. Computational methods in probability theory. D. P. Heyman, M. J. Sobel, eds., *Handbooks in Operations Research and Management Science*, vol. 2. Elsevier, Amsterdam, 199–254.

Grassmann, W. K., J. L. Jain. 1989. Numerical solutions of the waiting time distribution and idle time distribution of the arithmetic GI/G/1 queue. *Operations Research* **37**(1) 141–150. doi:10.1287/opre.37.1.141.

Green, R., G. Fusai, I. D. Abrahams. 2010. The Wiener-Hopf technique and discretely monitored path-dependent option pricing. *Mathematical Finance* **20**(2) 259–288. doi:10.1111/j.1467-9965.2010.00397.x.

Heston, S. L. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* **6**(2) 327–343. doi:10.1093/rfs/6.2.327.

Kemperman, J. H. B. 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, F. W. 2009. *Hilbert Transforms*. Cambridge University Press, Cambridge.

Kou, S. G. 2002. A jump-diffusion model for option pricing. *Management Science* **48**(8) 1086–1101. doi:10.1287/mnsc.48.8.1086.166.

Kou, S. G. 2008. Discrete barrier and lookback options. J. Birge, V. Linetsky, eds., *Handbooks in Operations Research and Management Science*, vol. 15. Elsevier, Amsterdam, 343–373.

Kreyszig, E. 2011. *Advanced Engineering Mathematics*. 10th ed. Wiley, New York.

Kuznetsov, A. 2010. Wiener-Hopf factorization and distribution of extrema for a family of Lévy processes. *Annals of Applied Probability* **20**(5) 1801–1830. doi:10.1214/09-AAP673.

Lewis, A. 2001. A simple option formula for general jump-diffusion and other exponential Lévy processes. *SSRN 282110* .

Madan, D. B., E. Seneta. 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., G. Fusai, G. Germano. 2012. Pricing credit derivatives in a Wiener-Hopf framework. M. Cummins, F. Murphy, J. J. H. Miller, eds., *Topics in Numerical Methods for Finance*. Springer, New York, 139–154. Springer Proceedings in Mathematics & Statistics, vol. 19.

Merton, R. C. 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* **3**(1) 125–144. doi:10.1016/0304-405X(76)90022-2.

Nolan, J. P. 2017. *Stable Distributions — Models for Heavy Tailed Data*. Birkhäuser, Boston. In progress, Chapter 1 online at academic2.american.edu/~jpnolan.

Phelan, C., G. Fusai, D. Marazzina, G. Germano. 2017. Hilbert transform, spectral filters and option pricing. Submitted.

Polyanin, A. D., A. Vladimirovich. 1998. *Handbook of Integral Equations*. CRC Press, Boca Raton.

Prabhu, N. U. 1974. Wiener-Hopf techniques in queueing theory. A. B. Clarke, ed., *Mathematical Methods in Queueing Theory*. Springer, Berlin, 81–90. Series on Lecture Notes in Economics and Mathematical Systems, vol. 98.

Rogers, L. C. G. 1994. Fluid models in queueing theory and Wiener-Hopf factorization of Markov chains. *Annals of Applied Probability* **4**(2) 390–413. URL http://www.jstor.org/stable/2245162.

Ruijter, M. J., M M. Versteegh, C. W. Oosterlee. 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*. Wiley, New York.

Spitzer, F. 1956. A combinatorial lemma and its application to probability theory. *Transactions of the American Mathematical Society* **82**(2) 323–339. doi:10.1090/S0002-9947-1956-0079851-X.

Stenger, F. 1993. *Numerical Methods Based on Sinc and Analytic Functions*. Springer, Berlin.

Stenger, F. 2011. *Handbook of Sinc Numerical Methods*. CRC Press, Boca Raton.

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

Appendix A

Table 6 contains all the parameters used for the numerical experiments which produced the results presented in Section 4.

	Description	Symbol	Value
	Maturity	T	1 year
	Initial spot price	S_0	1
	Strike	K	1.1
	Upper barrier (double-barrier)	U	1.40
Option parameters	Lower barrier (double-barrier)	U	0.60
	Upper barrier (down-and-out)	U	$+\infty$
	Lower barrier (down-and-out)	L	0.80
	Risk-free rate	r	0.05
	Dividend rate	q	0.02
Model	$\Psi(\xi, t)$	Symbol	Value
		α	15
NIG	$e^{-t\left(\sqrt{\alpha^2-(\beta+i\xi)^2}+\sqrt{\alpha^2-\beta^2}\right)}$	β	-5
		δ	0.5
		p	0.3
		λ	3
Kou	$e^{-t\left(\frac{\sigma^2\xi^2}{2}-\lambda\left(\frac{(1-p)\eta_2}{\eta_2+i\xi}+\frac{p\eta_1}{\eta_1-i\xi}-1\right)\right)}$	σ	0.1
		η_1	40
		η_2	12
		θ	$\frac{1}{9}$
VG	$(1-i\nu\xi\theta+\nu\sigma^2\xi^2/2)^{-t/\nu}$	σ	$\frac{1}{3\sqrt{3}}$
		ν	0.25

Table 6: Parameters for the numerical tests and processes used; $\Psi(\xi, t)$ is the characteristic function of the process that models the log return of the underlying asset.