University College London

Department of Computer Science

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Computer Science, University College London

Fourier transform methods for the pricing of barrier options and other exotic derivatives

Author
Carolyn E. Phelan

Supervisor
Dr Guido Germano
Department of Computer Science
University College London

Second Supervisor
Dr Fabio Caccioli
Department of Computer Science
University College London

September 10, 2018

This thesis is dedicated to my children: to my son Connor who brings me joy every day and to my daughter Erin who lived fleetingly and is loved always.

Abstract

This thesis focuses on the numerical calculation of fluctuation identities with both discrete and continuous monitoring and the wider application of finding a general numerical solution to the Wiener-Hopf equation on a semi-infinite or finite interval. The motivating application is pricing path-dependent options.

It is demonstrated that, with the use of spectral filters, exponential convergence can be achieved for the pricing of discretely monitored double-barrier options. We thus describe the first exponentially convergent pricing method for this type of option with general Lévy processes and a CPU time which is independent of the number of monitoring dates.

Using a numerical implementation of the inverse Laplace transform, the numerical method to calculate fluctuation identities is extended to continuous monitoring. This provides the first method for calculating continuously monitored fluctuation identities which can be used for general Lévy processes. Furthermore a detailed error bound is obtained, providing additional insight into the pricing methods based on fluctuation identities and the numerical solution of the Wiener-Hopf equation in general.

Pricing algorithms for other exotic options such as α-quantile, perpetual Bermudan and perpetual American options are also devised and a new method to compute the optimal exercise boundary for the latter two types of contract is presented. These methods show excellent error performance with computational time.

Finally, an application of these new numerical methods to the general solution of the Wiener-Hopf equation is presented. The methods are applied to three new test cases which we derived analytically and the results are presented to show that this new method has an error convergence with grid size which has twice the speed of the current state of the art.

Acknowledgments

The support of the Engineering and Physical Sciences Research Council (EPSRC) in funding the UK Centre for Doctoral Training in Financial Computing and Analytics (grant number 1482817) is gratefully acknowledged.

Very many thanks to my supervisor, Dr Guido Germano, for all his invaluable support, help and advice. I am also very grateful to our collaborators Professor Daniele Marazzina and Professor Gianluca Fusai whose comments and suggestions have been so very useful over the course of this work. Sincere thanks must also go to Professor Philip Treleaven who accepted me onto the PhD programme at the UK Centre for Doctoral Training in Financial Computing and Analytics and who has encouraged my development in this new career.

Last but not least I would like to thank my friends and family, most especially my husband Dr Iain Millar, without whose love and support this PhD would not have been possible.

Impact statement

The motivating application for the methods developed in this thesis is the accurate pricing of financial contracts. It is demonstrably true that devising computationally efficient and accurate pricing methods is an important topic. The global financial crisis illustrated, among other things, the risks associated with relying on mathematical models which inaccurately reflect the probability of extreme price movements in financial markets. With this in mind, in addition to producing several publications in academic journals with good impact factors, the work in this thesis has been presented at a number of international conferences attended by industry practitioners.

Notwithstanding the immediate applications within the finance industry, the mathematical insights in this thesis have particular relevance to other fields. The efficient calculation of fluctuation identities has uses in biological, physical and social sciences.

Furthermore, beyond the applications related to probability and fluctuation identities, a computationally efficient numerical solution to the Fredholm equation can be used in many other fields such as electro-magnetics, fluid mechanics and biological sciences.

Decalaration

I, Carolyn Phelan, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signature_____

Date_____

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and overview

This research is primarily motivated by the need to value path-dependent options based on a realistic model for the probability distribution of the price of the underlying asset and the option contract. Classical closed-form solutions to option pricing such as those by Black and Scholes (1973) and Merton (1973) rely on simplifying both the contracts and the underlying processes in order to render the problem analytically tractable. The underlying processes were modelled as log-normal which leads to elegant mathematical solutions but is inaccurate when compared to real-world asset price data whose probability distribution exhibits fat tails, skewness, and whose process includes jumps. In recent years, Lévy processes, which combine both jump and diffusion processes have been considered as an alternative.

In addition, the details of the option contract may affect the pricing requirements and classical methods relied on assumptions such as continuous monitoring of asset price paths when this is not realistic in practice. The importance of this particular aspect for modelling exotic option prices was stressed by Kou (2008) who wrote "Discrete path-dependent options are the options whose payoffs are determined by underlying prices at a finite set of times, whereas the payoff of a continuous path-dependent option depends on the underlying price throughout the life of the option. Due to regulatory and practical issues, most of path-dependent options traded in markets are discrete path-dependent options."

As the work in this field has strived for more realism in modelling processes and contracts, the focus has moved away from closed-form solutions and much of the most recent literature concentrates on the search for efficient numerical methods. This has been accompanied by an enormous increase in computational power since the original closed-form solutions were published in the early 1970s, as per Moore's law (Moore, 1998; Mack, 2011). Therefore numerical techniques which would formerly be dismissed as too

computationally heavy are now perfectly feasible.

Although numerical methods such as Monte Carlo and finite-differences are widely used and understood in the finance industry, they tend to have poor error convergence with computational time: polynomial for finite-differences and sub-polynomial for Monte Carlo. This is particularly problematic for path-dependent options as these techniques also have a computational time which increases linearly with the number of monitoring dates. Indeed this linear relationship with the number of dates is even true for more sophisticated methods such as the one by Feng and Linetsky (2008) who achieved excellent error results using the sinc-based fast Hilbert transform to apply the barrier at each monitoring step.

The focus of this thesis is on developing further methods using the Spitzer identities which are fluctuation identities providing the Fourier-z transform of probability distribution functions of the extrema of a random path, subject to monitoring at discrete intervals. They particularly lend themselves to the pricing of discretely monitored barrier options due to their discrete nature. Moreover, the use of the z-transform means that the time domain is collapsed and therefore the computational time is minimally affected by the number of monitoring dates. This property was exploited by Fusai et al. (2016), who proposed numeric methods to compute the Spitzer identities for exponential Lévy processes and thus used them to price barrier and lookback options. However, whilst the methods by Fusai et al. (2016) and Feng and Linetsky (2008) showed excellent performance in some cases, issues remained with the double barrier case for Fusai et al. (2016) and for the variance gamma process for both methods. Moreover, Fourier-based methods can sometimes be regarded as a "black-box" which, whilst giving highly accurate results, are not particularly well understood by the wider finance community. This, of course, can have a detrimental effect on the uptake of the method, however excellent the performance. Therefore there is also a clear need to provide rigorous error bounds for the methods to aid insight into the error convergence that is obtained.

Thus, the work in this thesis is concerned with two overarching themes. Firstly, to extend the application of, and improve the performance of, existing methods and secondly to perform a detailed error examination of both the existing and new methods to afford a thorough understanding of the performance of these numerical methods.

Although the primary practical focus of the work in this thesis is derivative pricing, this research has the potential to reach far beyond the area of financial mathematics. The calculation of fluctuation identities are relevant, for example, for application to queuing systems, insurance, inventory systems, and applied probability. Moreover, first-passage problems with models based on Markov processes are also ubiquitous in physical, biological, social, actuarial and other sciences. Beyond the applications related to random processes, our numerical methods for calculating fluctuation identities rely on the calculation of the solution to Wiener-Hopf equations which occur in many different applications such as electro-magnetics and biological sciences.

Chapter 2 provides a technical background to the work in this thesis. It is written so that it should contain sufficient information to ensure that this thesis is a self-contained work. However, the breadth of the technical background means that, in a few cases, only information directly related to the work herein has been included. Where this is the case, comprehensive references have also been included for the interested reader.

In Chapter 3 the use of spectral filtering to improve methods using the sinc-based fast Hilbert transform for the pricing of discretely monitored barrier options with Lévy processes is investigated. The techniques by Fusai et al. (2016) and Feng and Linetsky (2008) are extended and improved and a rigorous error analysis is carried out to explain the results seen by Fusai et al. (2016), aiding insight into the numerical calculation of the Spitzer identities and Fourier-transform based pricing techniques in general. This work has produced the first exponentially convergent pricing technique for discretely monitored double barrier options with general Lévy processes whose CPU time does not depend on the number of monitoring dates.

Chapter 4 describes the extension of the discretely monitored pricing procedure for barrier options to the continuously monitored case. Although, potentially less directly applicable within the finance industry due to the reasons previously explained, this work opens up the use of the technique for the other areas described above. Furthermore the work in this chapter lays the foundations for the work on American options described in Chapter 5 which are one example of an exotic option where continuous path dependence is an accurate reflection of financial contracts. Moreover, the work includes a rigorous derivation of the error bound which aids insight into the techniques in this thesis and elsewhere.

Chapter 5 describes the use of the Spitzer identity based pricing techniques for other exotic options, specifically α-quantile, perpetual Bermudan and perpetual American options. Two techniques have been implemented for the calculation of the early exercise options and a new method for computing the optimal exercise barrier is introduced.

Finally Chapter 6 looks at the possibility of extending the use of the numerical method for the calculation of the Spitzer identities to solving the Wiener-Hopf and Fredholm equations in the general case.

1.2 Publications

The work in this thesis is also described in the following manuscripts:

[1] Chapter 3 has been published as C. E. Phelan, D. Marazzina, G. Fusai, G. Germano, Hilbert transform, spectral filters and option pricing, Annals of Operations Research, DOI 10.1007/s10479-018-2881-4, 2018.

[2] Chapter 4 has been published as C. E. Phelan, D. Marazzina, G. Fusai, G. Germano,

Fluctuation identities with continuous monitoring and their application to the pricing of barrier options, European Journal of Operational Research, DOI 10.1016/j.ejor.2018.04.016, 2018.

[3] Chapter 5 is currently a working paper C. E. Phelan, D. Marazzina, G. Germano. 2018. Spitzer based pricing methods for α-quantile and perpetual American and Bermudan options.

[4] Chapter 6 forms a large part of the working paper G. Germano, C. E. Phelan, D. Marazzina, G. Fusai. Solution of Wiener-Hopf and Fredholm integral equations by fast Hilbert and Fourier transforms.

1.3 Conference communications

The work in this thesis has been presented at the following international conferences:

[1] C. E. Phelan, D. Marazzina, G. Fusai, G. Germano. Progresses on pricing barrier options with the Spitzer identity. Quantitative Finance Workshop 2017, Milan.

[2] C. E. Phelan, D. Marazzina, G. Fusai, G. Germano. Improvement of numerical option pricing methods based on the Hilbert transform using spectral filtering. Econophysics Colloquium 2017, Warsaw.

[3] C. E. Phelan, D. Marazzina, G. Fusai, G. Germano. Numerical pricing of discretely monitored barrier options with Levy jump processes using the Hilbert transform and spectral filtering. International Conference on Computational Finance 2017, Lisbon.

[4] C. E. Phelan, D. Marazzina, G. Germano. Pricing methods for perpetual Bermudan and quantile options based on Spitzer identities. Quantitative Finance Workshop 2018, Milan.

[5] C. E. Phelan, D. Marazzina, G. Germano. Spitzer based pricing methods for alpha-quantile and perpetual early exercise options. 42^{nd} meeting of the Association for Mathematics Applied to Social and Economic Sciences 2018, Naples.

1.4 Test environment

Where numerical tests have been carried out, the software has been implemented in MATLAB2016b and the results were obtained running under OS X Yosemite on a 2015 Retina MacBook Pro with a 2.7GHz Intel Core i5 processor and 8GB of RAM.

1.5 A note on pronouns

This introductory chapter has been written in the third person. However, in order to be consistent with the papers listed in Section 1.2, and also because it is the sincere wish of the author that this thesis will take the reader on a journey through the work described herein, the first person plural has been used in all that follows.

Chapter 2

Technical background and literature review

This chapter provides a technical background to the work in this thesis and a review of the relevant literature. Given the breadth and depth of this subject we have included only sufficient detail to make this thesis self-contained and references have been provided to allow the reader to explore this area in more detail if they wish. We first present the analytic background to the work contained in this thesis. We then describe the types of financial contract which are the subject of the work in later chapters and review the existing literature on pricing methods. Finally we review the original numerical methods which we analyse and modify in later chapters.

2.1 Analytic background

We begin with a description of Fourier and related transforms as they underpin all the techniques that are described in this thesis. We then move on to look at some of the applications of these such as the Plemelj-Sokhotsky relations and the solution to the Wiener-Hopf equations and finally we examine how these methods may be used to find the Spitzer identities based on the formulation by Green et al. (2010).

2.1.1 Fourier transforms

We make extensive use of the Fourier transform, which is an integral transform with many applications. Historically, it has been widely used in spectroscopy and communications, where the transform is moving between a function of time and a function of frequency. Thus much of the literature refers to the function in the Fourier domain as its spectrum. However, in this thesis a more general approach is taken and we move between a function of x in the "state" domain (or space) and a function of ξ in the Fourier domain. For our application of pricing options, x is the log-price; this is a common choice in finance

literature (see Heston, 1993, for example). However, in Chapter 6 we also look at the use of our techniques for the solution of general Wiener-Hopf and Fredholm equations and in these cases x can be any variable.

The Fourier transform and inverse Fourier transforms, as used in the work described herein are

$$\widehat{f}(\xi) = \mathcal{F}_{x\to\xi}[f(x)](\xi) = \int_{-\infty}^{+\infty} f(x)e^{i\xi x}dx, \tag{2.1}$$

$$f(x) = \mathcal{F}^{-1}_{\xi\to x}[\widehat{f}(\xi)](x) = \frac{1}{2\pi}\int_{-\infty}^{+\infty} \widehat{f}(\xi)e^{-ix\xi}d\xi. \tag{2.2}$$

This version of the Fourier transform in Eqs. (2.1) and (2.2) is just one of several conventions that may be used. We select this form as the forward transform in Eq. (2.1) is the same as the formula for a characteristic function of a random variable with probability distribution function $f(x)$. For convenience, we retain this choice for the general solution of the Wiener-Hopf and Fredholm equations. There are many extensive references available for the Fourier transform, see e.g. Polyanin and Manzhirov (1998); Beerends et al. (2003); Kreyszig (2011) for a detailed technical description, with Domínguez (2016) having recently provided a historical guide to its development. However, here we limit our background to the properties which underpin the work in this thesis: shift theorem, convolution theorem, Plancherel's theorem and the Paley-Wiener theorem of the analyticity of half range Fourier transforms.

The shift theorem can operate in either the state or Fourier domain and describes the effect of shifting the functions by a fixed amount, a. For our Fourier transform convention in Eqs. (2.1) and (2.2)

$$\mathcal{F}_{x\to\xi}[f(x)] = \widehat{f}(\xi) \quad \to \quad \mathcal{F}_{x\to\xi}[e^{-iax}f(x)] = \widehat{f}(\xi - a), \tag{2.3}$$

$$\mathcal{F}^{-1}_{\xi\to x}[\widehat{f}(\xi))] = f(x) \quad \to \quad \mathcal{F}^{-1}_{\xi\to x}[e^{ia\xi}\widehat{f}(\xi)] = f(x - a). \tag{2.4}$$

The convolution of two functions $f(x)$ and $g(x)$ is

$$h(x) = f(x) * g(x) = \int_{-\infty}^{+\infty} f(z)g(x - z)dz = \int_{-\infty}^{+\infty} f(x - z)g(z)dz. \tag{2.5}$$

Depending on the form of $f(x)$ and $g(x)$, solving this may involve a complicated integral. However, if both functions are transformed into the Fourier domain to give $\widehat{f}(\xi)$ and $\widehat{g}(\xi)$ then

$$\widehat{h}(\xi) = \widehat{f}(\xi)\widehat{g}(\xi,) \tag{2.6}$$

and $h(x)$ can be obtained by applying the inverse Fourier transform to $\widehat{h}(\xi)$.

Plancherel's theorem states that if we have two functions $f(x)$ and $g(x)$ with respective

Fourier transforms $\widehat{f}(\xi)$ and $\widehat{g}(\xi)$ then, with our Fourier transform convention,

$$\int_{-\infty}^{+\infty} f(x)g^*(x)dx = \frac{1}{2\pi}\int_{-\infty}^{+\infty} \widehat{f}(\xi)\widehat{g}^*(\xi)d\xi, \tag{2.7}$$

where z^* is the conjugate of z. In the case that $g(x)$ is purely real then we can also use property $\widehat{g}(-\xi) = \widehat{g}^*(\xi)$ to obtain

$$\int_{-\infty}^{+\infty} f(x)g(x)dx = \frac{1}{2\pi}\int_{-\infty}^{+\infty} \widehat{f}(\xi)\widehat{g}(-\xi)d\xi, \tag{2.8}$$

where we use $g^*(x) = g(x)$ as $g(x)$ takes only real values. Originally Fourier transforms were used for real values of ξ only. However, Paley and Wiener (1933) extended ξ to the whole complex plain and showed that the upper and lower half range Fourier transforms of a function are holomorphic in the upper and lower planes of ξ respectively. That is if

$$\widehat{f_+}(\xi) = \mathcal{F}_{x\to\xi}[f(x)\mathbf{1}_{\mathbb{R}^+}(x)](\xi) = \int_0^{+\infty} f(x)e^{i\xi x}dx, \tag{2.9}$$

$$\widehat{f_-}(\xi) = \mathcal{F}_{x\to\xi}[f(x)\mathbf{1}_{\mathbb{R}^-}(x)](\xi) = \int_{-\infty}^{0} f(x)e^{i\xi x}dx, \tag{2.10}$$

where $\mathbf{1}_A(x)$ is the indicator function of set A, then $\widehat{f_+}(\xi)$ is analytic in the upper half plane of ξ (including the real line) and $\widehat{f_-}(\xi)$ is analytic in the lower half plane of ξ (including the real line). Thus the analytic regions of $\widehat{f_+}(\xi)$ and $\widehat{f_-}(\xi)$ overlap in a strip including the real line.

2.1.2 Characteristic functions

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

$$\Psi(\xi,t) = \mathrm{E}\left[e^{i\xi X(t)}\right] = \int_{-\infty}^{+\infty} e^{i\xi x}p(x,t)dx = \mathcal{F}_{x\to\xi}\left[p(x,t)\right] = \widehat{p}(\xi,t). \tag{2.11}$$

The characteristic function uniquely describes a process and can be used to obtain its moments. One advantage that it has over the moment generating function in that it always exists whereas there are some random variables for which the moment generating function does not exist, for example the log-normal process. Furthermore, some probability distributions with extremely complicated PDFs have simple expressions for their characteristic functions. This is the case for many of the Lévy processes which are described in detail in Section 2.2.1. These properties are particularly useful in applications requiring the calculation of an expectation of a function of a random variable. Given a function $\phi(X)$, of a random variable X with PDF $p_X(x)$ the expectation can be expressed in the Fourier

domain using the Plancherel relation which we introduced in Eq. (2.7), i.e:

$$\mathbb{E}[\phi(X)] = \int_{-\infty}^{+\infty} \phi(x) p_X(x) dx = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \widehat{\phi}^*(\xi) \Psi(\xi) d\xi. \tag{2.12}$$

Here, $\widehat{\phi}^*(\xi)$ is the conjugate of the Fourier transform of $\phi(X)$ and $\Psi(\xi)$ is the characteristic function of $p_X(x)$.

2.1.3 Hilbert transforms

Similarly to the Fourier transform, the Hilbert transform is an integral transform, although in this case it does not change the domain of the function to which it is applied. The area of Hilbert transforms is a large topic and was covered extensively by King (2009) in his book on the subject so here we describe only the characteristics which are important for this thesis. The Hilbert transform of a function in the Fourier domain is

$$\begin{aligned} \mathcal{H}\big[\widehat{f}(\xi)\big] &= \text{P.V.} \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' \\ &= \lim_{\epsilon \to 0^+} \frac{1}{\pi} \left(\int_{\xi - 1/\epsilon}^{\xi - \epsilon} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' + \int_{\xi + \epsilon}^{\xi + 1/\epsilon} \frac{\widehat{f}(\xi')}{\xi - \xi'} d\xi' \right), \end{aligned} \tag{2.13}$$

where P.V. denotes the Cauchy principal value. There is no separate inverse transform as the Hilbert transform has the property that $\mathcal{H}[\mathcal{H}[\widehat{f}(\xi)]] = -\widehat{f}(\xi)$. We can relate the Hilbert transform to convolution theorem, desribed in Section 2.1.1 above, by seeing that Eq. (2.13) is a convolution of $f(\xi)$ and $\frac{1}{\pi\xi}$. As $\frac{i}{\pi\xi}$ is the Fourier transform of the signum function, applying the Hilbert transform in the Fourier domain is equivalent to multiplying the function in the state domain by $-i \operatorname{sgn} x$, where

$$\operatorname{sgn} x = \begin{cases} 1 & x > 0, \\ 0 & x = 0, \\ -1 & x < 0. \end{cases} \tag{2.14}$$

2.1.4 Laplace transforms

The Laplace transform is an integral transform similar to the Fourier transform and is

$$\mathcal{L}_{t \to s}[f(t)] = \tilde{f}(s) = \int_0^{\infty} f(t) e^{-st} dt. \tag{2.15}$$

It can be considered more general than the Fourier transform as s can be any complex number unlike $i\xi$ which is purely imaginary. Widder (1945) provided an excellent introduction to the subject and latterly most engineering mathematics books contain detailed chapters on the properties and use of the Laplace transform (see e.g. Kreyszig, 2011).

One significant difference compared to the Fourier transform is that the integral runs from 0 to ∞ rather than from $-\infty$ to $+\infty$. When the integral for the Laplace transform is performed over the range $\pm\infty$ the transform is known as the bilateral Laplace transform (Widder, 1945). However, the use of the single sided Laplace transform here is justified; we use it to transform a price function which is defined in terms of time to expiry ($\tau = T - t$) and therefore set maturity as $\tau = 0$ and so are not concerned with the behaviour at $\tau < 0$.

In many engineering applications, or when solving simple ODEs in general, the inverse Laplace transform is often obtained by comparing the equations in the s-domain with known $\mathcal{L}[f(t)] = \tilde{f}(s)$ pairs. However, for more general applications a formula for the inverse Laplace transform is required. This is known as the Bromwich contour inversion integral:

$$f(t) = \frac{1}{2\pi i} \int_{a-i\infty}^{a+i\infty} \tilde{f}(s) e^{st} ds, \tag{2.16}$$

where a is a real number and lies to the right of all singularities of $\tilde{f}(s)$ in the complex s-plane. The Bromwich integral can be complicated to perform in closed-form and, therefore, a numerical approximation is often required in practice. For the work on continuously monitored barrier options in Chapter 4 we implement and test a numerical method by Abate and Whitt (1995) which is based on a similar philosophy to the well established method for inverse z-transforms described in Section 2.3.3 of this chapter.

2.1.5 Z-TRANSFORMS

The z-transform for functions defined on a discrete domain is similar to the Laplace transform and is

$$\mathcal{Z}_{n\to q}[f(n)] = \tilde{f}(q) = \sum_{n=0}^{\infty} f(n) q^n, \tag{2.17}$$

where q is a complex number and f is a discretely sampled function which takes values $f(n)$, where $n \in \{0, 1, 2, \ldots\}$. It is often defined with a negative signed exponent on q in Eq. (2.17). However, the use of the positive exponent means that it has the same form as the probability generating function and thus matches the convention used in probability related literature. It is also closely related to the Laplace transform, as explored in more detail in Section 2.1.5.1. Similarly to the Laplace transform, the bilateral z-transform is used when the limits on the summation in Equation (2.17) are $\pm\infty$ rather than 0 to $+\infty$. However, the use of the single sided transform can be justified in the same way that the use of the single sided Laplace transform is justified in Section 2.1.4: we use n as number of dates to expiry which means that we are unconcerned with the behaviour for $n < 0$. The z-transform is often used in engineering applications such as the analysis of discrete time signals (see Martin, 1991; Ifeachor and Jervis, 1993, for examples) and, similarly to

the Laplace transform, the inverse is often obtained using known $\mathcal{Z}[f(n)] = \tilde{f}(q)$ pairs. For more general functions the inverse z-transform is

$$f(n) = \frac{1}{2\pi i} \oint_C \frac{\tilde{f}(q)}{q^{n+1}} dq, \tag{2.18}$$

where C is a counterclockwise closed path encircling the origin and entirely in the region of convergence. This is often impossible to solve analytically and so a numerical approximation is required. For the work in this thesis, following its successful use in pricing techniques (Fusai et al., 2016), we implemented the method based on a Fourier series approximation by Abate and Whitt (1992a,b). We describe this in Section 2.3 which details the numerical methods used in the later chapters of this thesis.

2.1.5.1 Relationship between the z-transform and the Laplace transform

We have seen in the previous sections that the Laplace transform is a single sided integral of a continuous function and the z-transform is a single sided summation of a discrete time function. We can explore this similarity further to show the mathematical equivalence between the two transforms.

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

$$\begin{aligned}
\mathcal{L}_{t\to s}[f_{\mathrm{c}}(t)] &= \int_0^\infty e^{-st} f_{\mathrm{c}}(t) dt \\
&= \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^{\infty} e^{-sn\Delta t} f_{\mathrm{c}}(n\Delta t) \\
&= \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^{\infty} (e^{-s\Delta t})^n f_{\mathrm{d}}(n) \\
&= \lim_{\Delta t \to 0} \Delta t \sum_{n=0}^{\infty} q^n f_{\mathrm{d}}(n) \\
&= \lim_{\Delta t \to 0} \Delta t \mathcal{Z}\left[f_{\mathrm{d}}(n)\right].
\end{aligned} \tag{2.19}$$

Thus the Laplace transform is equal to the limit of Δt multiplied by the z-transform as $\Delta t \to 0$. In addition to the relationship described above, the connection between the s-plane and the z-plane can be considered (Wang, 2014). This is a conformal mapping, preserving the angles between contours on the two planes and is illustrated in Figure 2.1. The vertical lines on the s-plane map to circles of radius $e^{-\Re s}$ on the z-plane and the horizontal lines map to lines emanating from the origin at an angle of $-\Im s$ to the positive real axis on the s-plane. As the angle between a line emanating from the origin and the circumference of a circle centred at the origin is $\frac{\pi}{2}$ then this is a conformal mapping. It

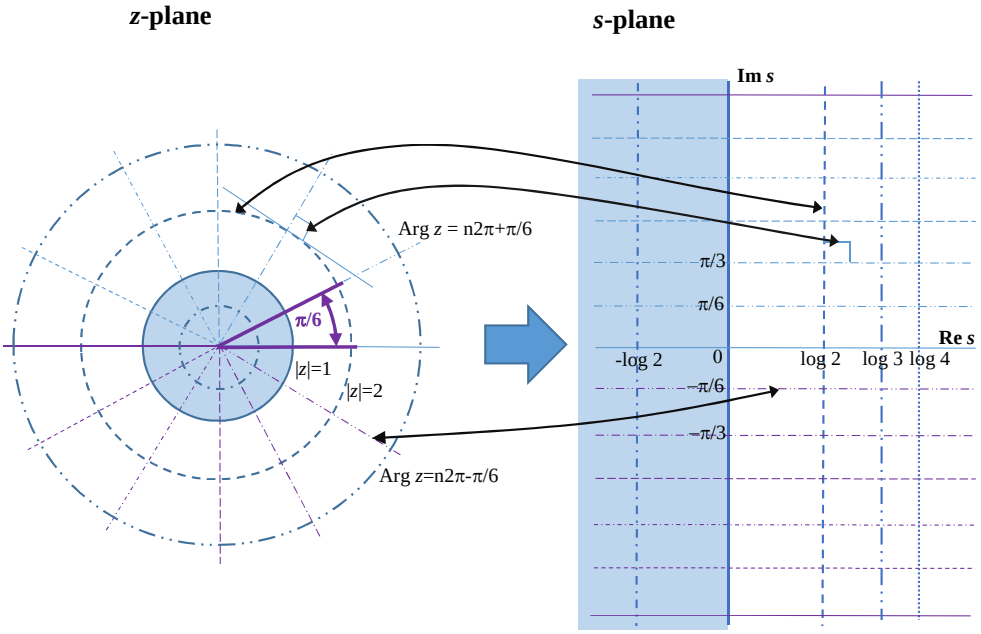

Figure 2.1: Conformal mapping between the z-plane and the s-plane.

is also important to note that the mapping from the s-plane to the z-plane is not unique. The point $s = \sigma + i\omega$ will map to the same point on the z-plane as $s = \sigma + i(k2\pi + \omega)$, where k is an integer.

The relationship between the inverse Laplace and inverse z-transforms can also be considered. The inverse Laplace transform in Eq. (2.16) can be written in terms of the z-transform

$$f(t) = \frac{1}{2\pi i}\int_{a-i\infty}^{a+i\infty} \lim_{\Delta t\to 0} \Delta t \mathcal{Z}[f(n)]e^{st}ds. \tag{2.20}$$

As $q = e^{-s\Delta t}$, then $ds = -dq/(\Delta t q)$, in addition the path of q now traces a contour of radius e^{-at} around the origin on the q plane rather than a line on the s plane as shown in Figure 2.1. Then Eq. (2.20) is equivalent to

$$\begin{aligned} f(t) &= \frac{1}{2\pi i}\oint_C \lim_{\Delta t\to 0} \Delta t \mathcal{Z}[f(n)] \frac{-1}{\Delta t q}\frac{1}{q^n}dq \\ &= \frac{1}{2\pi i}\oint_C \mathcal{Z}[f(n)]\frac{1}{q^{n+1}}dq. \end{aligned} \tag{2.21}$$

This is the same as the expression for the inverse z-transform in Eq. (2.18).

2.1.6 Plemelj-Sokhotsky relations

One of the most important concepts for the work in this thesis is that of he "+" and "−" parts of a function. These are defined as $f_+(x) = f(x)\mathbf{1}_{\mathbb{R}_+}(x)$ and $f_-(x) = f(x)\mathbf{1}_{\mathbb{R}_-}(x)$ and many of the techniques described herein require their Fourier transforms, i.e. $\widehat{f_+}(\xi) = \mathcal{F}_{x\to\xi}\left[f(x)\mathbf{1}_{\mathbb{R}_+}(x)\right]$ and $\widehat{f_-}(\xi) = \mathcal{F}_{x\to\xi}\left[f(x)\mathbf{1}_{\mathbb{R}_-}(x)\right]$ which are equivalent to the half range Fourier transforms described in Eqs. (2.9) and (2.10). To obtain these directly in the Fourier domain, the Hilbert transform described in Section 2.1.3 can be used to calculate the Plemelj-Sokhotsky relations (see Germano et al., 2018, for example). These relations use the property described in Section 2.1.3 that applying the Hilbert transform in the Fourier domain is equivalent to multiplying the function by $-i\operatorname{sgn} x$ in the state domain. Therefore as $f(x)\operatorname{sgn} x = f_+(x) - f_-(x)$ and the Hilbert transform is a linear operation due to it being an integral transform we obtain

$$-i\mathcal{H}[\widehat{f}(\xi)] = \widehat{f_+}(\xi) - \widehat{f_-}(\xi). \tag{2.22}$$

The Fourier transform is also a linear operation and as $f(x) = f_+(x) + f_-(x)$ then

$$\widehat{f}(\xi) = \widehat{f_+}(\xi) + \widehat{f_-}(\xi). \tag{2.23}$$

Eqs. (2.22) and (2.23) can then be combined to give the Plemelj-Sokhotsky relations which split a function into its positive and negative parts, a processes also known as decomposition,

$$\widehat{f_+}(\xi) = \frac{1}{2}\left[\widehat{f}(\xi) + i\mathcal{H}\widehat{f}(\xi)\right], \tag{2.24}$$

$$\widehat{f_-}(\xi) = \frac{1}{2}\left[\widehat{f}(\xi) - i\mathcal{H}\widehat{f}(\xi)\right]. \tag{2.25}$$

We can also use the shift theorem described in Section 2.1.1 to generalise the Plemelj-Sokhotsky relations to decompose a function above and below an arbitrary barrier b,

$$\widehat{f_{b+}}(\xi) = \frac{1}{2}\left\{\widehat{f}(\xi) + e^{ib\xi}i\mathcal{H}\left[e^{-ib\xi}\widehat{f}(\xi)\right]\right\}, \tag{2.26}$$

$$\widehat{f_{b-}}(\xi) = \frac{1}{2}\left\{\widehat{f}(\xi) - e^{ib\xi}i\mathcal{H}\left[e^{-ib\xi}\widehat{f}(\xi)\right]\right\}. \tag{2.27}$$

Eqs. (2.26) and (2.27) can be combined to obtain the Fourier transform the part of a function between two barriers, i.e. $\widehat{f_{lu}}(\xi) = \mathcal{F}_{x\to\xi}\left[f(x)\mathbf{1}_{[l,u]}(x)\right]$,

$$\widehat{f_{lu}}(\xi) = \frac{1}{2}\left\{e^{il\xi}i\mathcal{H}\left[e^{-il\xi}\widehat{f}(\xi)\right] - e^{iu\xi}i\mathcal{H}\left[e^{-iu\xi}\widehat{f}(\xi)\right]\right\}. \tag{2.28}$$

The numerical methods described in this thesis also require the factorisation of a function, i.e. the calculation of $\widehat{f_\oplus}(\xi)$ and $\widehat{f_\ominus}(\xi)$ such that $\widehat{f}(\xi) = \widehat{f_\oplus}(\xi)\widehat{f_\ominus}(\xi)$. Note the use of

$\cdot_\oplus$ and $\cdot_\ominus$ to donate factorisation instead of decomposition. Factorisation is achieved by arithmetically decomposing the logarithm $\widehat{h}(\xi) = \log \widehat{f}(\xi)$ i.e.

$$\widehat{h_+}(\xi) = \frac{1}{2}\big[\widehat{h}(\xi) + i\mathcal{H}\widehat{h}(\xi)\big] = \frac{1}{2}\big[\log \widehat{f}(\xi) + i\mathcal{H}\log \widehat{f}(\xi)\big], \tag{2.29}$$

$$\widehat{h_-}(\xi) = \frac{1}{2}\big[\widehat{h}(\xi) - i\mathcal{H}\widehat{h}(\xi)\big] = \frac{1}{2}\big[\log \widehat{f}(\xi) - i\mathcal{H}\log \widehat{f}(\xi)\big], \tag{2.30}$$

and then exponentiating the results to obtain $\widehat{f_\oplus}(\xi) = \exp \widehat{h_+}(\xi)$ and $\widehat{f_\ominus}(\xi) = \exp \widehat{h_-}(\xi)$.

For the factorisation of a function multiplied by a constant e.g. $\widehat{g}(\xi) = T\,\widehat{f}(\xi)$, we apply the logarithm for $\widehat{h}(\xi) = \log \widehat{g}(\xi) = \log T + \log \widehat{f}(\xi)$. The Plemelj-Sokotsky relations then give

$$\widehat{h_+}(\xi) = \frac{1}{2}\big[\widehat{h}(\xi) + i\mathcal{H}\widehat{h}(\xi)\big] = \frac{1}{2}\big[\log T + \log \widehat{f}(\xi) + i\mathcal{H}\log \widehat{f}(\xi)\big], \tag{2.31}$$

$$\widehat{h_-}(\xi) = \frac{1}{2}\big[\widehat{h}(\xi) - i\mathcal{H}\widehat{h}(\xi)\big] = \frac{1}{2}\big[\log T + \log \widehat{f}(\xi) - i\mathcal{H}\log \widehat{f}(\xi)\big], \tag{2.32}$$

where we use the Hilbert transform properties of linearity and zero value for a constant input. These differ from Eqs. (2.29) and (2.30) only by the term $0.5 \log T$, which gives the following relation between the factors of $\widehat{f}(\xi)$ and $\widehat{g}(\xi)$

$$\widehat{g_\oplus}(\xi) = \sqrt{T}\,\widehat{f_\oplus}(\xi), \tag{2.33}$$

$$\widehat{g_\ominus}(\xi) = \sqrt{T}\,\widehat{f_\ominus}(\xi). \tag{2.34}$$

Thus the effect of the constant T is equally split between the positive and negative factors so that they are both multiplied by $\sqrt{T}$.

2.1.7 Wiener-Hopf technique

An important subject for the work in this thesis, both as a stand-alone topic and for the application of option pricing, is the solution of the Wiener-Hopf (or Fredholm) equations which are of the form

$$\lambda f(x) - \int_a^b k(x - x')f(x')dx' = g(x) \quad x \in (a, b) \tag{2.35}$$

(see e.g. Polyanin and Manzhirov, 1998), where we wish to solve for $f(x)$. The functions $k(x)$ and $g(x)$ are known and are respectively referred to as the kernel and forcing function. If either $a = -\infty$ or $b = +\infty$ then the expression in Eq. (2.35) is a Wiener-Hopf equation (Wiener and Hopf, 1931; Noble, 1958; Lawrie and Abrahams, 2007) and if a and b are both finite then it is referred to as a Fredholm equation (Fredholm, 1903). When $\lambda = 0$ then it is an equation of the first kind and when $\lambda \neq 0$ it is of the second kind.

From its application to microwave propagation, as in Daniele and Lombardi (2006), the former is often described as a classical Wiener-Hopf equation (CWHE) and the latter

as a longitudinally modified Wiener-Hopf equation (LMWHE) (Daniele and Zich, 2014). In addition to radio propagation, the overview of the Wiener-Hopf method by Lawrie and Abrahams (2007) described further applications such as crystal growth, fracture mechanics, flow mechanics (see also Choi et al., 2005), geophysics and diffusion. It is also connected to problems in probability such as the Spitzer identities which are discussed in detail in Section 2.1.8 below. The Wiener-Hopf technique has been used several times in the finance literature to price discretely monitored path-dependent options. This has involved both the calculation the price directly such as in Fusai et al. (2006) and also finding probability distributions of underlying processes subject to discrete monitoring such as in Green et al. (2010), Fusai et al. (2016) and Marazzina et al. (2012).

The Wiener-Hopf method was described in great detail in the book by Noble (1958) with Lawrie and Abrahams (2007) providing both a historical perspective and references for the more recent developments in the technique. The latter is especially recommended to the interested reader wishing for an introduction to the subject. Here we present a brief description of the solution to the Wiener-Hopf and Fredholm equations of the second kind in order to provide a framework for the numerical methods developed in the later chapters of this thesis.

We first describe the Wiener-Hopf method for Eq. (2.35) over a semi-infinite range and therefore set $b = \infty$. The lower integration limit can also be shifted to $a = 0$ without loss of generality (Germano et al., 2018) and so Eq. (2.35) becomes

$$\lambda f(x) - \int_0^\infty k(x - x')f(x')dx' = g(x) \quad x \in (0, \infty). \tag{2.36}$$

The similarity of the integral to the convolution in Eq. (2.5) is immediately apparent but as the range of the integration is not infinite the convolution theorem cannot be applied directly. In order to use the convolution theorem we must first extend Eq. (2.36) over the whole range of $x \in \mathbb{R}$. We write

$$\lambda f(x)\mathbf{1}_{\mathbb{R}^+}(x) - \int_{-\infty}^{+\infty} k(x - x')f(x')\mathbf{1}_{\mathbb{R}^+}(x')dx' + f_{\mathrm{s}}(x)\mathbf{1}_{\mathbb{R}^-}(x) = g(x)\mathbf{1}_{\mathbb{R}^+}(x) \quad x \in \mathbb{R} \tag{2.37}$$

where, as before, $\mathbf{1}_A(x)$ is the indicator function of the set A. The use of the indicator function within the integral allows the extension of x' over the full range of $\mathbb{R}$. The supplementary function $f_s(x)$ is introduced as

$$f_s(x) = \int_{-\infty}^{\infty} k(x - x')f(x')\mathbf{1}_{\mathbb{R}^+}(x')dx', \tag{2.38}$$

and allows the function to be defined over all values of $x \in \mathbb{R}$. For ease and clarity of

notation the new functions are defined

$$f_+(x) = \mathbf{1}_{\mathbb{R}^+}(x) f(x), \tag{2.39}$$

$$g_+(x) = \mathbf{1}_{\mathbb{R}^+}(x) g(x), \tag{2.40}$$

$$f_-(x) = \mathbf{1}_{\mathbb{R}^-}(x) f_s(x). \tag{2.41}$$

Notice how the $\cdot_+$ subscript corresponds to a "+" function and the $\cdot_-$ subscript to a "−" function. We can thus rewrite Eq. (2.37) more concisely as

$$\lambda f_+(x) - \int_{-\infty}^{+\infty} k(x-x') f_+(x') dx' + f_-(x) = g_+(x), \quad x \in \mathbb{R}. \tag{2.42}$$

Having extended the Wiener-Hopf equation to $x \in \mathbb{R}$, we can now apply the Fourier transform to Eq. (2.42) and use the convolution theorem to obtain

$$\lambda \widehat{f_+}(\xi) - \widehat{k}(\xi)\widehat{f_+}(\xi) + \widehat{f_-}(\xi) = \widehat{g_+}(\xi), \tag{2.43}$$

and by defining $\widehat{l}(\xi) = \lambda - \widehat{k}(\xi)$ we can write Eq. (2.43) more compactly as

$$\widehat{l}(\xi)\widehat{f_+}(\xi) + \widehat{f_-}(\xi) = \widehat{g_+}(\xi). \tag{2.44}$$

The first step in the Wiener-Hopf method is to factorise $\widehat{l}(\xi)$ so that $l(\xi) = \widehat{l_\oplus}(\xi)\widehat{l_\ominus}(\xi)$. We have the constraint that if $\widehat{l_\ominus}(\xi)$ is a "–" function then $\widehat{l_\ominus}^{-1}(\xi)$ must also be a "–" function (Noble, 1958), but we can show that if we use the factorisation method described in Section 2.1.6 then, subject to conditions on $\log \widehat{l}(\xi)$[1], this requirement is met. Decomposing $\log \widehat{l}(\xi) = [\log \widehat{l}(\xi)]_+ + [\log \widehat{l}(\xi)]_-$ gives a function $[\log \widehat{l}(\xi)]_+$ $\left([\log \widehat{l}(\xi)]_-\right)$ which, by the Paley-Wiener theorem, is analytic in the upper (lower) ξ plane. Using the same notation as Section 2.1.6 we obtain

$$\widehat{l_\oplus}(\xi) = e^{[\log \widehat{l}(\xi)]_+}, \tag{2.45}$$

$$\widehat{l_\ominus}(\xi) = e^{[\log \widehat{l}(\xi)]_-}. \tag{2.46}$$

As $\widehat{l_\oplus}(\xi)$ $\left(\widehat{l_\ominus}(\xi)\right)$ is calculated via the exponentiation of an analytic function we know that it is both analytic and non-zero in the upper (lower) ξ plane and therefore $\widehat{l_\oplus}^{-1}(\xi)$ $\left(\widehat{l_\ominus}^{-1}(\xi)\right)$ is also analytic and non-zero in the upper (lower) ξ plane. Having factorised $l(\xi)$, the next step is to divide Eq. (2.43) through by $\widehat{l_\ominus}(\xi)$ for

$$\widehat{l_\oplus}(\xi)\widehat{f_+}(\xi) + \widehat{l_\ominus}^{-1}(\xi)\widehat{f_-}(\xi) = \widehat{l_\ominus}^{-1}(\xi)\widehat{g_+}(\xi). \tag{2.47}$$

Then as $\widehat{l_\oplus}(\xi)\widehat{f_+}(\xi)$ is purely a "+" function and $\widehat{l_\ominus}^{-1}(\xi)\widehat{f_-}(\xi)$ is purely a "–" function,

[1] We use the Paley-Wiener theorem which requires that $f(x) \in L_2(\mathbb{R}_+)$ for the Fourier transform over ≥ 0 and $f(x) \in L_2(\mathbb{R}_+)$ for the Fourier transform over ≤ 0

we can decompose $\widehat{l_{\ominus}}^{-1}(\xi)\widehat{g_{+}}(\xi)$ and

$$\widehat{l_{\oplus}}(\xi)\widehat{f_{+}}(\xi) = \big[\widehat{l_{\ominus}}^{-1}(\xi)\widehat{g_{+}}(\xi)\big]_{+}, \tag{2.48}$$

$$\widehat{l_{\ominus}}^{-1}(\xi)\widehat{f_{-}}(\xi) = \big[\widehat{l_{\ominus}}^{-1}(\xi)\widehat{g_{+}}(\xi)\big]_{-}. \tag{2.49}$$

The Fourier transform of the unknown function $\widehat{f_{+}}(\xi)$ can then be calculated by rearranging Eq. (2.48)

$$\widehat{f_{+}}(\xi) = \widehat{l_{\oplus}}^{-1}(\xi)\big[\widehat{l_{\ominus}}^{-1}(\xi)\widehat{g_{+}}(\xi)\big]_{+}. \tag{2.50}$$

If we extend the method to the Fredholm equation then a and b are both finite. Therefore defining Eq. (2.35) over the entire range of x gives

$$\begin{aligned}\lambda f(x)\mathbf{1}_{(a,b)}(x) - \int_{-\infty}^{+\infty} k(x-x')f(x')\mathbf{1}_{(a,b)}(x')dx' \\ + f_{\mathrm{s}}(x)\mathbf{1}_{(-\infty,a)}(x) + f_{\mathrm{s}}(x)\mathbf{1}_{(b,\infty)}(x) = g(x)\mathbf{1}_{(a,b)}(x) \quad x \in \mathbb{R}.\end{aligned} \tag{2.51}$$

As before, $f_s(x)$ is a supplementary function which allows the extension of the Fredholm equation to the entire range of $x \in \mathbb{R}$ and is

$$f_s(x) = \int_{-\infty}^{\infty} k(x-x')f(x')\mathbf{1}_{(a,b)}(x')dx', \tag{2.52}$$

Again, for brevity and clarity we define

$$f_0(x) = \mathbf{1}_{(a,b)}(x)f(x), \tag{2.53}$$

$$g_0(x) = \mathbf{1}_{(a,b)}(x)g(x), \tag{2.54}$$

$$f_+(x) = \mathbf{1}_{(b,\infty)}(x)f_s(x), \tag{2.55}$$

$$f_-(x) = \mathbf{1}_{(-\infty,a)}(x)f_s(x), \tag{2.56}$$

$$k_0(x) = \mathbf{1}_{(a-b,b-a)}(x)k(x), \tag{2.57}$$

so Eq. (2.51) is now

$$\lambda f_0(x)_0 - \int_{-\infty}^{+\infty} k_0(x-x')f_0(x')dx' + f_-(x) + f_+(x) = g_0(x) \quad x \in \mathbb{R}. \tag{2.58}$$

We again apply the Fourier transform to obtain

$$\widehat{l}(\xi)\widehat{f_0}(\xi) + \widehat{f_-}(x) + \widehat{f_+}(x) = \widehat{g_0}(x), \tag{2.59}$$

where $\widehat{l}(\xi) = \lambda - \widehat{k_0}(\xi)$. However, we now have two supplementary functions, $\widehat{f_-}(x)$ and $\widehat{f_+}(x)$, so unlike the solution to the Wiener-Hopf solution over a semi-infinite range of x

described in Eqs. (2.42)–(2.50) we cannot solve this directly. An iterative solution was proposed by Henery (1977), but this is a theoretical solution and has not been successfully implemented. The iterative solution successfully realised by Germano et al. (2018) is based on producing two versions of Eq. (2.59), one shifted so that b is the origin and one shifted so that a is the origin:

$$e^{-ib\xi}\widehat{l}(\xi)\widehat{f}_0(\xi) + e^{-ib\xi}\widehat{f_-}(x) + e^{-ib\xi}\widehat{f_+}(x) = e^{-ib\xi}\widehat{g}_0(x), \tag{2.60}$$

$$e^{-ia\xi}\widehat{l}(\xi)\widehat{f}_0(\xi) + e^{-ia\xi}\widehat{f_-}(x) + e^{-ia\xi}\widehat{f_+}(x) = e^{-ia\xi}\widehat{g}_0(x). \tag{2.61}$$

We then factorise $\widehat{l}(\xi)$ and divide Eqs. (2.60) and (2.61) by $\widehat{l_\oplus}(\xi)$ and $\widehat{l_\ominus}(\xi)$ respectively to obtain,

$$e^{-ib\xi}\widehat{l_\ominus}(\xi)\widehat{f}_0(\xi) + e^{-ib\xi}\widehat{l_\oplus}^{-1}(\xi)\widehat{f_-}(x) + e^{-ib\xi}\widehat{l_\oplus}^{-1}(\xi)\widehat{f_+}(x) = e^{-ib\xi}\widehat{l_\oplus}^{-1}(\xi)\widehat{g}_0(x), \tag{2.62}$$

$$e^{-ia\xi}\widehat{l_\oplus}(\xi)\widehat{f}_0(\xi) + e^{-ia\xi}\widehat{l_\ominus}^{-1}(\xi)\widehat{f_-}(x) + e^{-ia\xi}\widehat{l_\ominus}^{-1}(\xi)\widehat{f_+}(x) = e^{-ia\xi}\widehat{l_\ominus}^{-1}(\xi)\widehat{g}_0(x). \tag{2.63}$$

Here, the terms $e^{-ib\xi}\widehat{l_\ominus}(\xi)\widehat{f}_0(\xi)$ and $e^{-ia\xi}\widehat{l_\ominus}^{-1}(\xi)\widehat{f_-}(x)$ are "$-$" and the terms $e^{-ia\xi}\widehat{l_\oplus}(\xi)\widehat{f}_0(\xi)$ and $e^{-ib\xi}\widehat{l_\oplus}^{-1}(\xi)\widehat{f_+}(x)$ are "+". Therefore, by decomposing the remaining terms in Eqs. (2.62) and (2.63) and rearranging, we can obtain expressions for $\widehat{f_+}(\xi)$ and $\widehat{f_-}(\xi)$ which can be used to solve Eq. (2.59)

$$\widehat{f_+}(x) = \widehat{l_\oplus}(\xi)e^{ib\xi}\left[e^{-ib\xi}\left(\frac{\widehat{g_0}(x) - \widehat{f_-}(x)}{\widehat{l_\oplus}(\xi)}\right)\right]_+ = \widehat{l_\oplus}\left[\frac{\widehat{g_0}(x) - \widehat{f_-}(x)}{\widehat{l_\oplus}(\xi)}\right]_{b+}, \tag{2.64}$$

$$\widehat{f_-}(x) = \widehat{l_\ominus}(\xi)e^{ia\xi}\left[e^{-ia\xi}\left(\frac{\widehat{g_0}(x) - \widehat{f_+}(x)}{\widehat{l_\ominus}(\xi)}\right)\right]_- = \widehat{l_\ominus}\left[\frac{\widehat{g_0}(x) - \widehat{f_+}(x)}{\widehat{l_\ominus}(\xi)}\right]_{a-}, \tag{2.65}$$

where the notation $[\cdot]_{a-}$ and $[\cdot]_{b+}$ denotes decomposition around a barrier as described in the generalized Plemelj-Sokhotsky relations in Eqs. (2.26) and (2.27). The expressions in Eqs. (2.64) and (2.65) cannot be solved directly; they form a system of equations which can be solved using an iterative scheme such as the fixed-point algorithm employed by Fusai et al. (2016) and Germano et al. (2018). The value of $\widehat{f_+}(\xi)$ is initialised as 0 and then Eq. (2.65) is solved to give an estimate for $\widehat{f_-}(x)$ which is then input to Eq. (2.64) to provide an estimate for $\widehat{f_+}(x)$, and so on, until a required tolerance or maximum number of iterations is reached. The computed values for $\widehat{f_+}(x)$ and $\widehat{f_-}(x)$ are then used to solve Eq. (2.59) for $\widehat{f}_0(\xi)$:

$$\widehat{f}_0(\xi) = \frac{\widehat{g_0}(x) - \widehat{f_-}(x) - \widehat{f_+}(x)}{\widehat{l}(\xi)}. \tag{2.66}$$

Work by Fusai et al. (2016) using this iterative method found that convergence to the final error level is achieved within about 2–3 iterations, beyond which little improvement is possible.

The work in this thesis is based on the iterative method described above, however for completeness we provide here a summary of the latest developments finding a direct solution to the Fredhom equation; for more details the reader is directed to Germano et al. (2018).

The approach by Noble (1958) puts Eqs. (2.62) and (2.63) in matrix form

$$\begin{pmatrix} \widehat{l}(\xi) & e^{i(b-a)\xi} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \widehat{f_0}(\xi)e^{-ia\xi} \\ \widehat{f_+}(\xi)e^{-ib\xi} \end{pmatrix} + \begin{pmatrix} 0 & 1 \\ \widehat{l}(\xi) & e^{i(a-b)\xi} \end{pmatrix} \begin{pmatrix} \widehat{f_0}(\xi)e^{-ia\xi} \\ \widehat{f_-}(\xi)e^{-ib\xi} \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ \widehat{l}(\xi) & e^{i(a-b)\xi} \end{pmatrix} \begin{pmatrix} 0 \\ \widehat{g_0}(\xi)e^{-ia\xi} \end{pmatrix}. \tag{2.67}$$

Multiplying by $\begin{pmatrix} 0 & 1 \\ \widehat{l}(\xi) & e^{i(a-b)\xi} \end{pmatrix}^{-1}$ from the left gives

$$\begin{pmatrix} -e^{i(b-a)\xi} & 0 \\ \widehat{l}(\xi) & e^{i(b-a)\xi} \end{pmatrix} \begin{pmatrix} \widehat{f_0}(\xi)e^{-ia\xi} \\ \widehat{f_+}(\xi)e^{-ib\xi} \end{pmatrix} + \begin{pmatrix} \widehat{f_0}(\xi)e^{-ia\xi} \\ \widehat{f_-}(\xi)e^{-ib\xi} \end{pmatrix} = \begin{pmatrix} 0 \\ \widehat{g_0}(\xi)e^{-ia\xi} \end{pmatrix}. \tag{2.68}$$

When this is written in short form as

$$\widehat{\mathbf{l}}(\xi)\widehat{\mathbf{f}_+}(\xi) + \widehat{\mathbf{f}_-}(\xi) = \widehat{\mathbf{g}_+}(\xi), \tag{2.69}$$

the similarity with Eq. (2.44) is immediately apparent. Thus, if we could factorise $\widehat{\mathbf{l}}(\xi)$ we could solve the Fredholm equation directly without the use of iterative methods. Significant advances were made by Jones (1984, 1991) on the factorisation of 2×2 matrices. However the conditions that must be met to apply this method are not fulfilled due to the presence of elements such as $e^{i(a-b)\xi}$. Similarly to the conditions required by the Wiener-Hopf method that if $\widehat{l_\ominus}(\xi)$ is "$-$" then $\widehat{l_\ominus}^{-1}(\xi)$ must also be "$-$" (Noble, 1958), we also require that $\widehat{\mathbf{l}_\ominus}^{-1}(\xi)$ is "$-$". Unfortunately the inversion of a "$-$" element such as $e^{i(a-b)\xi}$ gives a "$+$" element $e^{i(b-a)\xi}$ [2] (Daniele and Zich, 2014).

An alternate factorisation method was proposed by Voronin (2004). This requires the conversion of a known matrix factorisation such that $\widehat{\mathbf{m}}(\xi) = \widehat{\mathbf{m}_\ominus}(\xi)\widehat{\mathbf{m}_\oplus}(\xi)$ into one where $\widehat{\mathbf{m}}(\xi) = \widehat{\mathbf{m}_\oplus}(\xi)\widehat{\mathbf{m}_\ominus}(\xi)$ and again whilst methods exist for this operation, the form of our matrices is not such that we can apply them (Jones, 1984). Since these major developments in the 1980s, there have been a number of advances in the area of Wiener-Hopf matrix factorization, both for 2×2 matrices such as by Abrahams (1997); Kisil (2015); Mishuris and Rogosin (2016) and for the more general case of the factorisation of an $n \times n$ matrix (see Veitch and Abrahams, 2007). However, it remains the case that no general solution for the factorisation of an $n \times n$ Wiener-Hopf matrix has been discovered. Therefore we are

[2] We can easily see this by recognising that $e^{ic\xi}$ is the Fourier transform of a shifted delta function $\delta(x - c)$ and if c changes sign the delta function will change from being a "+" to a "-" function or vice versa

required to use an iterative scheme such as the one by Fusai et al. (2016) described above, or to use other solutions to integral equations such as quadrature method, as described in Press et al. (2007) Eq. (4.1.12).

2.1.8 Spitzer identities and the Green et al. formulation

Spitzer (1956) devised his eponymous fluctuation identities using combinatorial arguments; these were extended to the continuous case by Baxter and Donsker (1957) and to the double-barrier case by Kemperman (1963). The Spitzer identities provide the Fourier-z transform of the probability distributions of the extrema of a process subject to discrete monitoring and of the final value of a process conditional on the process crossing a discretely monitored barrier. The link to Wiener-Hopf factorisation was first discussed by Baxter (1961); this has been explored in much greater detail by Green et al. (2007); Green (2009); Green et al. (2010). This work used a modified Wiener-Hopf technique (Jones' method) to recast the Spitzer identities in a form that is both numerically tractable and useful for the pricing of path-dependent options.

The relevance of the Spitzer identity in several fields 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. The reader is directed to the work by Green (2009) for a comprehensive guide to the formulation of the Spitzer identities used for the work in this thesis. However, we include here a brief summary of this technique in order to provide sufficient background information for the work herein.

For a discretely monitored process X_n, $(n = 0, 1, 2, ...)$ with the transition density between monitoring dates of $k(x - x')$ for $X_n = x$ and $X_{n-1} = x'$, we wish to know the probability density of $X_n = x$ subject to the process not having crossed upper or lower barriers at previous monitoring dates. We adapt the slightly unusual designation of $k(\cdot|\Delta t)$ for the transition density, where Δt is the time between monitoring dates, to clearly illustrate the link to the Wiener-Hopf method described in Section 2.1.7.

We dub the probability density which we wish to compute $p(x, n)$ and the idea behind Green's formulation of the Spitzer identities is that the relationship between $p(x, n)$ at subsequent monitoring dates can be expressed recursively as

$$p(x, n) = \int_d^u k(x - x'|\Delta t)p(x', n - 1)dx', \tag{2.70}$$

where $p(x_0) = \delta(x_0)$, the Dirac delta function. Applying the z-transform defined in

Eq. (2.17) to Eq. (2.70) gives

$$\widetilde{p}(x,q) - \delta(x) = q\int_d^u k(x-x'|\Delta t)\widetilde{p}(x',q)dx', \tag{2.71}$$

where $\widetilde{p}(x,q) = \mathcal{Z}[p(x,n)] = \sum_{j=0}^{\infty} q^n p(x,n)$ and q is the z-transform parameter. The expression in Eq. (2.71) can be rewritten as a Wiener-Hopf equation similar to Eq. (2.35), where the z-transformed probability distribution $\widetilde{p}(x',q)$ is used instead of $f(x)$ for the unknown function, the kernel $k(x)$ is replaced by the transition density $k(x|\Delta t)$and $\delta(x)$ is the forcing function $g(x)$. The techniques used to solve the Wiener-Hopf and Fredholm equations in Section 2.1.7 can then be used to provide an expression for the required probability distribution in the Fourier-z domain.

For the work on barrier options in Chapters 3 and 4 and early exercise options in Chapter 5, the identities we require are

$$p_l(x,n)dx = \mathrm{P}\left[X_n \in (x, x+dx) \cap \max_{j<n} X_j > l\right], \tag{2.72}$$

$$p_{lu}(x,n)dx = \mathrm{P}\left[X_n \in (x, x+dx) \cap \max_{j<n} X_j < u \cap \min_{j<n} X_j > d\right]. \tag{2.73}$$

It should be noted that these are defective probability distributions in x because, defining the cumulative distribution function (CDF) as $F(x) = \int_{-\infty}^{x} p(x,n)dx$, then $F(+\infty) - F(-\infty) < 1$. Eq. (2.72) is calculated in the Fourier-z domain as

$$p_l(x,n) = \begin{cases} \mathcal{Z}^{-1}_{q\to n}\left[\mathcal{F}^{-1}_{\xi\to x}\left[\frac{P_{l+}(\xi,q)}{\Phi_{\oplus}(\xi,q)}\right]\right] & x \geq l, \\ \mathcal{Z}^{-1}_{q\to n}\left[\mathcal{F}^{-1}_{\xi\to x}\left[P_{l-}(\xi,q)\Phi_{\ominus}(\xi,q)\right]\right] & x < l. \end{cases} \tag{2.74}$$

Here, $\Phi_{\oplus}(\xi,q)$ and $\Phi_{\ominus}(\xi,q)$ are the Wiener-Hopf factors of $\Phi(\xi,q) = 1 - q\Psi(\xi,\Delta t) = \Phi_{\oplus}(\xi,q)\Phi_{\ominus}(\xi,q)$, $\Psi(\xi,\Delta t)$ being the characteristic function of the transition density $k(x,\Delta t)$, and $P_{l\pm}(\xi,q)$ result from the decomposition around l of $P(\xi,q) = 1/\Phi_{\ominus}(\xi,q) = P_{l+}(\xi,q) + P_{l-}(\xi,q)$. Green et al. (2010) extended the usual definition of the single sided Wiener-Hopf equation, where the range of x is limited to the range of the integral. Using arguments from complex analysis Green showed that the form of the identity for the final value subject to the path being monitored in Eq (2.74) is different depending on whether the value of $X(t_n)$ at the final date is above or below l and this is reflected above. If we link the formulation of this identity for $x < l$ to the Wiener-Hopf method described in Section 2.1.7, then the general idea behind these identities is that rather than solving Eq. (2.48) for $\widehat{f_+}(\xi)$, we instead solve Eq. (2.49) for $\widehat{f_-}(\xi)$. The pricing scheme for barrier options in Fusai et al. (2016) used in Chapters 3 requires the expression in the case that the final value of $X(t_n)$ is above the barrier. However, the pricing methods for Bermudan options described in Chapter 5 require the distribution where the value of $X(t_n) < l$. Eq. (2.73)

is calculated in the Fourier-z domain as

$$p_{lu}(x,n) = \mathcal{Z}_{q\to n}^{-1}\left[\mathcal{F}_{\xi\to x}^{-1}\left[\frac{1-J_{l-}(\xi,q)-J_{u+}(\xi,q)}{\Phi(\xi,q)}\right]\right], \tag{2.75}$$

where J_{l-} and $Ju+$ are the solutions to the coupled equations.

$$\frac{J_{u+}(\xi,q)}{\Phi_{\oplus}(\xi,q)} = \left[\frac{1-J_{l-}(\xi,q)}{\Phi_{\oplus}(\xi,q)}\right]_{u+}, \tag{2.76}$$

$$\frac{J_{l-}(\xi,q)}{\Phi_{\ominus}(\xi,q)} = \left[\frac{1-J_{u+}(\xi,q)}{\Phi_{\ominus}(\xi,q)}\right]_{l-}. \tag{2.77}$$

Eq. (2.75) is analogous to Eq. (2.65) in Section 2.1.7 for the general solution to the Fredholm equation and Eqs. (2.76) and (2.77) are equivalent to the coupled equations for the auxiliary functions in Eqs. (2.64) and (2.65).

For the work on α-quantile options in Chapter 5, we also require the probability distribution of the minimum and maximum of discretely monitored processes, i.e.

$$p_m(x,n)dx = \mathrm{P}\left[\max_{j\le n} X_j \in (x, x+dx)\right], \tag{2.78}$$

$$p_M(x,n)dx = \mathrm{P}\left[\min_{j\le n} X_j \in (x, x+dx)\right]. \tag{2.79}$$

Green (2009) used arguments from complex analysis to caluclate these in the Fourier-z domain as

$$p_m(x,n) = \mathcal{Z}_{q\to n}^{-1}\left[\mathcal{F}_{\xi\to x}^{-1}\left[\frac{1}{\Phi_{\oplus}(0,q)\Phi_{\ominus}(\xi,q)}\right]\right], \tag{2.80}$$

$$p_M(x,n) = \mathcal{Z}_{q\to n}^{-1}\left[\mathcal{F}_{\xi\to x}^{-1}\left[\frac{1}{\Phi_{\oplus}(\xi,q)\Phi_{\ominus}(0,q)}\right]\right]. \tag{2.81}$$

In addition to Spitzer identities for the distribution of the values of processes subject to touching discretely monitored barriers, Green (2009) produced formulations for the PDFs of the first crossing time of discretely monitored barrier. For example, the PDF of the first crossing time of a lower barrier τ_l is obtained by calculating the probability that n is the first date that $X_j \le d\ \forall j \le n$ as

$$\mathrm{P}\left[\tau_l = n\right] = \int_{-\infty}^{d} p_l(x,n)dx, \tag{2.82}$$

where the required expression for $p_l(x,n)$ can be obtained applying the inverse Fourier-z transform to the expression for $x < l$ in Eq. (2.74).

2.1.8.1 Green's formulation for continuous monitoring

Following the work by Baxter and Donsker (1957), it was shown in Green et al. (2010) and the supplementary material of Fusai et al. (2016) that the identities derived above for discretely monitored barriers could be extended to the continuous monitoring case using the relationship between the Laplace and z-transforms described in Eq. (2.19). We repeat this here due to its importance for later chapters of this thesis. Having rearranged Eq. (2.71) in the form of the standard Wiener-Hopf equation, the first step in the solution method outlined in Section 2.1.7 is to factorise $\widehat{l}(\xi) = \lambda - \widehat{k}(\xi)$ which is equivalently defined as $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t)$ using the notation from Section 2.1.8 where $\Psi(\xi, \Delta t)$ is the characteristic function, as defined in Eq. 2.11, of the transition density $k(x|\Delta t)$.

To obtain Green's formulation of the Spitzer identities for continuous monitoring, first the equivalent quantity of $\frac{1}{\Phi(\xi,q)}$ in the Laplace domain $\frac{1}{\Phi^c(\xi,s)}$ must be calculated. This can be found using the result shown in Eq. (2.19), i.e. $\mathcal{L}[f_{\mathrm{c}}(t)] = \lim_{\Delta t \to 0} \Delta t \mathcal{Z}[f_{\mathrm{d}}(n)]$. Defining $\frac{1}{\Phi^c(\xi,s)}$ in terms of the limit of the z-transform gives

$$\begin{aligned}
\frac{1}{\Phi^c(\xi, s)} &= \lim_{\Delta t \to 0} \frac{\Delta t}{\Phi(\xi, q)} \\
&= \lim_{\Delta t \to 0} \frac{\Delta t}{1 - q\Psi(\xi, \Delta t)} \\
&= \lim_{\Delta t \to 0} \frac{\Delta t}{1 - e^{-s\Delta t} e^{\psi(\xi)\Delta t}} \\
&= \lim_{\Delta t \to 0} \frac{\Delta t}{1 - e^{(\psi(\xi) - s)\Delta t}} \\
&= \frac{1}{s - \psi(\xi)},
\end{aligned} \tag{2.83}$$

where the final step uses the Taylor expansion $e^x = \sum_{n=0}^{\infty} \frac{x^n}{n!}$. Having obtained $\Phi^c(\xi, s)$, the solution can be obtained using the Wiener-Hopf method as before, except in this case the results will be the probability distribution of the final value of a process subject to continuous monitoring, expressed in the Fourier-Laplace domain.

2.2 Option pricing

There are many applications for Spitzer identities and Wiener-Hopf equations; examples are described in Sections 2.1.7 and 2.1.8, although these lists are not intended to be exhaustive. The main focus of this thesis is the application of the pricing of financial options which has been a major area within finance literature for the last few decades. Seminal work by Black and Scholes (1973) and Merton (1973) introduced the ideas of the risk neutral measure and replicating portfolio which are still important today (see e.g. Björk, 2009; Hull, 2017; Shreve, 2004). The famous formula by Black and Scholes (1973) gives a closed-form solution for the price of European options and Merton (1973)

described closed-form solutions for both perpetual American options and down-and-out barrier options. However both papers were limited to a log-normal price process for the underlying asset and where barrier monitoring was required, it was assumed to be continuous. Since that time a large area of research has been to increase the sophistication of both the models of the underlying assets and the contracts in order to simulate the real world more accurately. The importance of this work was brought sharply into focus during the 2007–2008 global financial crisis which demonstrated the risks associated with the reliance on models which inaccurately reflected reality.

2.2.1 Stochastic processes for modelling asset prices

In terms of increasing the model sophistication, one approach has been to retain the log-normal form of the stochastic differential equation controlling the underlying asset, i.e.

$$dS(t) = \mu(t)S(t)dt + \sigma(t)S(t)dW(t), \tag{2.84}$$

but then to introduce one or more stochastic process driving the process parameters. Examples for stochastic volatility are the Heston (1993) model, the CEV model (Cox, 1975) and the SABR model (Hagan et al., 2002) with Chen (1996) being the first to publish a model with both stochastic mean and volatility.

Another approach to improve the model of the underlying asset has been to keep the parameters deterministic but to increase complexity by combining a diffusion process, as described in Eq. (2.84), with a jump process. This allows the process to model large sudden movements in the asset price and also produces the excess kurtosis seen in the distribution of real world market returns. A large class of processes that combine diffusion with jumps process are known as Lévy processes (see Cont and Tankov, 2004, for example), and have the following characteristics:

1. $X(0) = 0$.
2. Independence of increments, i.e. for any $0 < t_1 < t_2 < ... < t_n < \infty$, $X(t_n) - X(t_{n-1}), X(t_{n-1}) - X(t_{n-2}), ..., X(t_1) - X(t_2)$ are independent.
3. Stationarity of increments, i.e. for any $s < t$,$X(t) - X(s)$ has the same probability distribution as $X(t - s)$.
4. Continuity in probability, i.e for any $\epsilon > 0$ and $t \geq 0$ it holds that $\lim_{h \to 0} P(|X_{t+h} - X_t| > \epsilon) = 0$.

The simplest examples of Lévy processes are a Wiener process where the increments are normally distributed with a mean of 0 and a variance of $t - s$, or a Poisson process where increments have a Poisson distribution with mean and variance $\lambda(t - s)$. Other Lévy processes are made up of both these components, for example Merton jump-diffusion which

is comprised of a diffusion process and a compound Poisson process. It is often the case that financial assets are modelled as exponential Lévy processes, i.e. $S(t) = S_0 e^{X(t)}$, where S_0 is the initial value of $S(t)$, a simple example of this being geometric Brownian motion. The increased flexibility of general Lévy processes allows for more accurate simulation of the movement of real asset prices. In addition, they have the advantage that they have the same simple form of characteristic function, i.e. $\Psi(\xi, t) = e^{\psi(\xi)t}$, where the characteristic exponent $\psi(\xi)$ is given by the Lévy-Khincine formula as

$$\psi(\xi) = ia\xi - \frac{1}{2}\sigma^2\xi^2 + \int_{\mathbb{R}} (e^{i\xi\eta} - 1 - i\xi\eta \mathbf{1}_{[-1,1]}(\eta))\nu(d\eta). \tag{2.85}$$

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

$$a = r - q - \frac{1}{2}\sigma^2 - \int_{\mathbb{R}} (e^{\eta} - 1 - i\eta \mathbf{1}_{[-1,1]}(\eta))\nu(d\eta), \tag{2.86}$$

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. Specific examples of this are the Gaussian (Schoutens, 2003), Normal inverse Gaussian (NIG) (Barndorff-Nielsen, 1998), CGMY (Carr et al., 2002), Kou double exponential (Kou, 2002), Merton jump-diffusion (Merton, 1976), Lévy alpha stable (Nolan, 2018), variance gamma (VG) (Madan and Seneta, 1990) and Meixner (Schoutens, 2003) processes.

2.2.2 Exotic options

Having initially focused on the problem of pricing vanilla options, where the payoff only depends on the value of the underlying asset at a single predefined expiry date, finance literature is now more concerned with exotic options, especially those where the option payoff depends on the path of the underlying as well as its final value. Fluctuation identities, such as the ones by Spitzer described in Section 2.1.8, have applications in the pricing of many of these contracts such as single and double-barrier, lookback, quantile, perpetual Bermudan and perpetual American options; these are described in more detail in Sections 2.2.2.1–2.2.2.3 below.

2.2.2.1 Barrier options

A barrier option is a financial contract which has a payoff which is a function of the value of the underlying asset at expiry, but which only pays out depending on whether the path of the underlying has touched (knock-in) or not touched (knock-out) a predefined

barrier. They are widely traded in the financial markets, according to Dadachanji (2015) "In particular, barrier options are one of the most traded exotic derivatives in the forex market". Types of single-barrier options are down-and-out, up-and-out, down-and-in and up-and-in. When considering single-barrier options in this thesis we concentrate on the down-and-out type of option, which becomes valueless if the asset touches a lower barrier, but similar techniques can be used for other types of single-barrier options. Closed-form solutions were derived by Merton (1973) and Goldman et al. (1979) for continuously monitored single-barrier options with the underlying asset price driven by geometric Brownian motion. Kunitomo and Ikeda (1992) extended this to the double-barrier case, deriving a solution which is a sum of an infinite series with rapid convergence. However these solutions, although elegant, do not lend themselves to our application as we wish to price discretely monitored barrier options with general Lévy processes and therefore we need to explore numerical techniques. As also quoted by Lian et al. (2017) the requirement for discrete monitoring is explained by Kou (2008) who wrote "due to regulatory and practical issues, most of path-dependent options traded in markets are discrete path-dependent options.... In practice most, if not all, barrier options traded in markets are discretely monitored. In other words, they specify fixed time points for the monitoring of the barrier (typically daily closings)."

When devising such methods we would wish for one with the following attributes: computational complexity which is invariant with the number of monitoring dates, very fast (preferably exponential) convergence with increasing computational time, useful for both single and double-barriers. Of the existing methods for option pricing, one of the simplest to to implement and most intuitive, is Monte Carlo pricing. However, although variance reduction techniques have been published, such as those by Beaglehole et al. (1997) and Dingeç and Hörmann (2012), which improve the convergence with the number of sample paths, the issue of the computational time increasing linearly with the number of monitoring dates remains.

Another possibility, given the existence of some closed-form solutions for barrier options, is whether continuously monitored option prices can be used as an approximation for the discrete case. Such methods would have the advantage that the computation time is invariant with the number of monitoring dates N. Unfortunately, as shown by Broadie et al. (1997), the convergence of discretely monitored solutions to the continuous case with the number of dates has the very slow rate of $O(1/\sqrt{N})$. Moreover, despite the existence of continuously monitored solutions with geometric Brownian motion, there is no solution for general exponential Lévy processes. Indeed, in Chapter 4 we explore the extension of the Spitzer identity calculations to continuous monitoring using option pricing as a motivating example. Whilst continuously monitored barrier options are not usually used in practice, the technique leads directly to the method for perpetual American options described in Chapter 5, which is a realistic example of a contract using continu-

ous time. Furthermore, a method for the solution of the continuous monitoring barrier problem with exponential Lévy processes is useful for other fields. First-passage problems with models based on Markov processes are also ubiquitous in physical, biological, social, actuarial and other sciences. For example, our technique could be used to compute the ruin probability, i.e. the probability that a Lévy process takes value in a set A at a time $T > 0$ given that the process never falls below a barrier B in the interval $[0, T]$, i.e., $P(X(T) \in A, \min_{t\in[0,T]} X(t) > B)$. This a classical problem in actuarial science and applied probability; see for example Klüppelberg et al. (2004). For applications in physics and biophysics, see e.g. the review by Bray et al. (2013). Similar problems also arise in statistics, see for example the classical paper by Chernoff (1961), or in studying when a process reaches for the first time an adverse threshold state (a patient dies, or an industrial device breaks down).

Finite-difference schemes for pricing discretely monitored options have also been explored in the option pricing literature (see Boyle and Tian, 1998, for example) and recently Golbabai et al. (2014) presented a model which was empirically shown to have 4th order accuracy with grid size. However, the results by Golbabai et al. (2014) were only presented for geometric Brownian motion and not for general Lévy processes. Furthermore, finite-difference schemes have the underlying drawback that they must have at least as many time domain grid point as monitoring dates and thus computational load is generally not independent of N.

Quadrature methods use the property that the price process for discretely monitored double-barrier options can be expressed as a recursive integral for each time step, similar to Eq. 2.70

$$v(x, n) = \int_d^u k(x - x'|\Delta t)v(x', n - 1)dx'. \tag{2.87}$$

Here $v(x, n)$ is the value of the option at the n^{th} monitoring date, $k(x - x'|\Delta t)$ is the transition density of the underlying process, and the option value at expiry is the payoff function $\phi(x)$. Quadrature methods can be used to solve this integral at each time step as shown by Andricopoulos et al. (2003) and Fusai and Recchioni (2007), for example. The convergence of these methods is polynomial, albeit often of a very high order, and the technique can be used for general Lévy process. However, the disadvantage is that we must have at least one integration step for each monitoring date and so the computational load increases linearly with the number of monitoring dates.

The last class of methods that we discuss, and the main focus of this thesis, is the use of transforms in option pricing. These use the integral transforms described in Section 2.1 applied to either the time or log-price domain, or both. Due to the importance of these methods to the work herein we also briefly describe a history of the technique applied to option pricing in general before concentrating on its application to barrier options. A

technical overview of the use of Fourier transforms for pricing options is included in Section 2.2.3 below.

Derivative pricing with Fourier transforms was famously investigated by Heston (1993) who looked at calculating the characteristic functions of the probability of options being in the money at expiry. This approach was preceded by Chen and Scott (1992) who used the moment generating function (Laplace transform of the PDF) instead of the characteristic function. The Fourier transform technique was also used by Bakshi and Chen (1997) and Scott (1997), underlining its usefulness when applied to the probability distributions of price movements in financial markets. The next major step forward in the area of option pricing using transforms was by Carr and Madan (1999) who published the first method for pricing European options with both the characteristic function and the payoff in the Fourier domain. More recently Fang and Oosterlee (2008) devised a pricing technique for European options based on the Fourier-cosine expansion and this was extended to general Lévy processes (Lord et al., 2008) and more exotic types of options including discretely monitored barriers (Fang and Oosterlee, 2009, 2011). The Fourier-cosine method exhibits extremely good accuracy having exponential error convergence with the number of grid points in the log-price domain when the transition density of the underlying asset is completely smooth in all orders of differentiation, i.e. $k(x - x'|\Delta t) \in C^{\infty}$. Unfortunately, for discretely monitored options, this technique has the drawback that the computational time increases linearly with the number of monitoring dates. The original techniques also had the disadvantage that for transition densities where $k(x - x'|\Delta t) \notin C^{\infty}$, such as VG, the error converges polynomially. However very recently Ruijter et al. (2015) showed that this can be improved with very little additional computational cost using spectral filters.

For pricing discretely monitored barrier options, the Hilbert transform (King, 2009) was also successfully employed by Feng and Linetsky (2008, 2009) to price barrier options using backward induction in Fourier space. This method showed excellent accuracy having exponential error convergence with log-price grid size down to machine accuracy of 10^{-16}. However, it also shares the two main disadvantages of the original Fourier-cosine series method when applied to discretely monitored barrier options: its computational load increases linearly with the number of monitoring dates and for transition densities where $k(x - x'|\Delta t) \notin C^{\infty}$ the performance is polynomial. We investigate improving this method with the use of spectral filters in Chapter 3 of this thesis.

In order to remove the linear dependence on the number of dates of the computational time inherent in many pricing methods for discretely monitored barrier options, Fusai et al. (2006) looked at collapsing the time domain using the z-transform. This meant that the z-transform of the option price could be written in the form of a Wiener-Hopf equation which was then solved analytically for a log-normal price process. The inverse z-transform by Abate and Whitt (1992a,b) has a computational time independent of the number of dates due to the use of the Euler acceleration (see e.g. O'Cinneide, 1997). Therefore,

applying this method to the z-transform of the price obtained by solving the Wiener-Hopf equation gives a pricing scheme whose computational time is independent of monitoring dates. As described in Section 2.1.8, this technique was adapted by Green et al. (2007); Green (2009); Green et al. (2010) to give the Fourier-z transform of the required probability distributions. These can then be used to price options in the Fourier domain using the Plancherel relation in Eq. (2.12), as described in Section 2.2.3 in more detail. Fusai et al. (2016) provided a numerical implementation of this scheme that used the discrete Hilbert transform from Stenger (1993, 2011) which was so successfully employed by Feng and Linetsky (2008). This implementation is applicable to general exponential Lévy processes, has a computational load which is independent of the number of monitoring dates and achieves exponential error convergence for single-barrier options where $k(x-x'|\Delta t) \in C^\infty$. However, for double-barriers or when $k(x-x'|\Delta t) \notin C^\infty$, only polynomial convergence is achieved. The work in Chapter 3 of this thesis shows that we can improve this performance with the addition of spectral filters in a similar way to the improvement to the Fourier-cosine method achieved by Ruijter et al. (2015); see also Phelan et al. (2018).

Recently Lian et al. (2017) proposed a method for barrier options with general exponential Lévy processes which avoids the use of the inverse z-transform by calculating the price at the N^{th} date through computing the $(N-1)^{\text{th}}$ power of a matrix. However, they did not make use of any standard acceleration techniques for matrix power calculations, such as diagonalisation. Therefore, the computational load associated with calculating the N^{th} power of a matrix is dependent on N and therefore in some ways this method is a retrograde step as it cannot be considered date independent. In other recent literature Farnoosh et al. (2017) and Sobhani and Milev (2018) published methods based on projection and Legendre polynomials that have computational time independent of the number of monitoring dates. However, although interesting, both methods have only been implemented for Brownian motion thus far and therefore currently do not fulfil the requirement for a pricing model for general exponential Lévy processes.

2.2.2.2 Perpetual Bermudan and American options

It is well known that, in contrast to European options which can only be exercised at a single expiry date, American options can be exercised at any time up to expiry. Bermudan options can be considered to be halfway between the two in that they can be exercised at a finite set of dates. Alternatively, they can be thought of as a discretely monitored version of an American option. These can have a fixed expiry date, beyond which the contract is worthless or no expiry date, known as perpetual Bermudan and American options. The valuation of American options is a long-standing problem in mathematical finance (Merton, 1973; Brennan and Schwartz, 1977) as it combines an optimisation problem (that of the early exercise barrier) and a pricing problem, but a closed-form solution has not been found with a finite expiry date.

Perpetual American options do permit a closed-form solution when the underlying is driven by geometric Brownian motion (Merton, 1973), as the perpetual nature of the option means that the optimal exercise barrier is constant rather than a function of time to expiry. However, this closed-form solution cannot be extended to Bermudan options. Moreover, as explained by Boyarchenko and Levendorskii (2002a), the smooth pasting method used for pricing perpetual American options can fail under jump processes such as those in the general Lévy class.

Several approximate methods suggested for finite-horizon American options, such as finite-differences (Brennan and Schwartz, 1977), trees (Cox et al., 1979), Monte Carlo (Rogers, 2002) and recursive Hilbert transforms by Feng and Lin (2013), are inherently discrete and thus lend themselves to Bermudan options with finite expiry. However, they are not particularly accurate or efficient for perpetual Bermudan options as the computational load increases with the number of monitoring dates which is, of course, infinite for perpetual contracts. Moreover, in times of very low interest rates such as the last decade, the existence of potential exercise dates even many years in the future can have a significant effect on the option value and thus the number of dates cannot be approximated with a sufficiently low value for these methods to be computationally efficient for perpetual options.

Boyarchenko and Levendorskii (2002a) published a method for pricing perpetual American options for many Lévy processes using analytic approximations to the Wiener-Hopf factors, this was a major step forward in showing their applicability to pricing perpetual options with Lévy processes. However it is an approximate solution as there is no general closed-form method for calculating the Wiener-Hopf factors. In addition, the proposed method is not applicable to all Lévy processes and specifically excludes the VG process. In Boyarchenko and Levendorskii (2002b) this method was adapted for perpetual Bermudan options. However, the Wiener-Hopf factorisation again requires approximation in some cases and calculations were presented for simple jump-diffusion and NIG processes only.

Mordecki (2002) also devised a pricing approximation for perpetual American put options with Lévy processes which was based on the optimal stopping problem for partial sums by Darling et al. (1972) and therefore intrinsically operates in discrete time and thus is useful for Bermudan options. However, this method had restrictions placed on the jump measure used in the Lévy-Khinchine characteristic exponent in Eq. (2.85) and therefore cannot be used for general Lévy processes.

Green (2009) suggested a pricing method based on his formulation of the Spitzer identities and his expression for prices of first touch and overshoot options. This also uses the residue method and is applicable for general Lévy processes, however an implementation of this method has not been published. In Chapter 5 we construct a pricing scheme based on Green's residue method and also specify an alternate scheme based on the Spitzer

identities which includes a new way to calculate the optimal exercise barrier.

2.2.2.3 Quantile options

Hindsight options are a class of options with a fixed expiry where the payoff at expiry is determined by the path of the option up to that date. Two such examples are lookback options and quantile options. Fixed-strike lookback options have a similar payoff to a European option. However, instead of being a function of the underlying asset price at expiry it is calculated using the maximum or minimum over the monitoring period i.e. $(S_M - K)^+$ for a call and $(K - S_m)^+$ for a put. For a continuously monitored option $S_M = \max_{t\in[0,T]} S(t)$ and $S_m = \min_{t\in[0,T]} S(t)$, but for a discretely monitored lookback option, the maximum or minimum is based on the asset price at set of discrete dates, i.e. $S_M = \max_{n\in[1,N]} S(t_n)$ and $S_m = \min_{n\in[1,N]} S(t_n)$.

Quantile options can be considered an adaption of lookback options because, rather than using the maximum or minimum of an asset price, the payoff is based on the value which the asset price spends $\alpha\%$ of the time below. For this reason they are often described as α-quantile options. They were first suggested by Miura (1992) as as way to design a hindsight option which is less sensitive to very extreme but short lived extrema of the asset price than simple lookback options. Akahori (1995) and Yor (1995) published analytic pricing methods with geometric Brownian motion. Since then, most work on pricing these types of options has been based on the remarkable identity by Dassios (1995) which used work by Port (1963) and Wendel (1960). The Dassios-Port-Wendel identity states that the α-quantile of a Brownian motion over T has the same distribution as the sum of the infinum of a Brownian motion over time $(1-\alpha)T$ plus the supremum of an independent Brownian motion over time αT. That is if $X_m = \min_{t\in[0,(1-\alpha T)]} X_1(t)$ and $X_M = \max_{t\in[0,\alpha T]} X_2(t)$, where $X_1(t)$ and $X_2(t)$, are independent Brownian motions then

$$X_\alpha \stackrel{d}{=} X_m + X_M, \tag{2.88}$$

where X_α is the α-quantile of the Brownian motion. A note by Dassios (2006) also showed that this could also be extended (although not necessarily uniquely) to general Lévy processes. Several pricing methods with Lévy processes have been developed, such as Monte Carlo methods for jump-diffusion processes by Ballotta (2002) and an analytic method for the Kou double exponential process by Cai et al. (2010). For discrete monitoring, Fusai and Tagliani (2001) explored the relationship between continuously and discretely monitored quantile options and whilst they made recommendations for selecting an optimal pricing method depending on the value of Δt, they did not produce a correction term like the one by Broadie et al. (1997) for barrier options.

Atkinson and Fusai (2007) provided closed-form prices for discretely monitored quantile options by using the z-transform to write the problem in terms of a Wiener-Hopf

equation. They solved this analytically for prices driven by geometric Brownian motion and also demonstrated the relationship between their results and the Spitzer identities. This approach was extended to general Lévy processes by Green (2009), who developed direct methods for calculating the distribution of the supremum and infimum of processes based on his formulation of the Spitzer identities. These were implemented for lookback options in Fusai et al. (2016) achieving exponential error convergence for general Lévy processes with a CPU time which was independent of the number of monitoring dates. In Chapter 5 we implement a pricing scheme for α-quantile options based on the numerical methods by Fusai et al. (2016).

2.2.3 Option pricing using Fourier methods

As discussed in Section 2.2.2.1, several different approaches can be used for option pricing with Fourier transforms. For example, Fusai et al. (2006) looked at collapsing the time domain of the price process using the z-transform to form a Wiener-Hopf equation which was then solved for geometric Brownian motion. However, in this thesis we concentrate on methods where the required PDF is calculated separately in the Fourier-z or Fourier domain and then the price calculation is done using the Plancherel relation with a Fourier transform of the payoff function. We therefore provide here a brief technical background to pricing options when both the payoff and PDF are available in the Fourier domain. We use a payoff which is calculated for an exponential random process which is the case for the option pricing methods we explore in this thesis.

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 $\phi(x(T))e^{-\alpha_{\mathrm{d}}x(T)}$ at maturity $t = T$ with respect to an appropriate risk-neutral probability distribution function (PDF) $p(x,T)$ whose initial condition is $p(x,0) = \delta(x)$. The most straightforward method uses convolution and makes the key assumption that the probability distribution function of the underlying process at time T, given it's value at an earlier time t can be represented as a transition density, i.e.

$$p(x(T)|x(t)) = k(x(T) - x(t)). \tag{2.89}$$

That is, the probability distribution of the process is purely dependent upon the different between the current and future values and not on either of their absolute values. We can

then calculate a price using the convolution theorem described in Section 2.1.1

$$\begin{aligned} v(x,t) &= e^{-rT}\mathrm{E}\left[\phi(x(T))e^{-\alpha_\mathrm{d}x(T)}|x(t)=x\right] \\ &= e^{-rT}\int_{-\infty}^{+\infty}\phi(x)e^{-\alpha_\mathrm{d}x}p(x'|x)dx' \\ &= e^{-rT}\int_{-\infty}^{+\infty}\phi(x')e^{-\alpha_\mathrm{d}x}k(x-x')dx \\ &= e^{-rT}\mathcal{F}^{-1}_{\xi\to x}\left[\widehat{\phi}(\xi)\Psi^*_k(\xi+i\alpha_\mathrm{d},T)\right], \end{aligned} \tag{2.90}$$

where we make use of the fact that the third line in Eq. (2.90) is a convolution and is therefore a multiplication in the Fourier domain. Here, $\Psi(\xi,T)$ is the characteristic function of the transition density and so $\Psi^*_k(\xi+i\alpha_\mathrm{d},T)$ is the complex conjugate of the Fourier transform of $e^{-\alpha_\mathrm{d}x}k(x,T)$. To price options using this relation, we need the Fourier transforms of both the damped payoff and the characteristic function. A double-barrier option has the damped payoff

$$\phi(x) = e^{\alpha_\mathrm{d}x}S_0(\theta(e^x-e^{k_S}))^+\mathbf{1}_{[l,u]}(x), \tag{2.91}$$

where $e^{\alpha_\mathrm{d}x}$ 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, $u=\log(U/S_0)$ is the upper log-barrier, $l=\log(L/S_0)$ is the lower log-barrier, U is the upper barrier and L is the lower barrier. For the log-strike we use $k_S=\log(K/S_0)$ where the subscript distinguishes it from the transition density ($k(\cdot,\Delta T)$) and K is the strike price. The Fourier transform of the damped payoff is available analytically:

$$\widehat{\phi}(\xi) = S_0\left(\frac{e^{(1+i\xi+\alpha_\mathrm{d})a}-e^{(1+i\xi+\alpha_\mathrm{d})b}}{1+i\xi+\alpha_\mathrm{d}} - \frac{e^{k_S+(i\xi+\alpha_\mathrm{d})a}-e^{k_S+(i\xi+\alpha_\mathrm{d})b}}{i\xi+\alpha_\mathrm{d}}\right), \tag{2.92}$$

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

For a Lévy process we have seen that the characteristic function is $\Psi(\xi,t)=e^{\psi(\xi)t}$, where the characteristic exponent $\psi(\xi)$ is given by Eq. (2.85) and is, in general available, in closed-form. Whilst Eq. (2.90) is useful for pricing options where the requirement that $p(x(T)|x(t))=k(x(T)-x(t))$ is met, it cannot be used in more general cases. An example of this is pricing barrier options where the required probability distribution of future values depends not only on the distance to the current value but also on its proximity to a barrier. In cases such as these Lewis (2001) showed that this can be done using the Plancherel

relation:

$$\begin{aligned} v(0,0) &= e^{-rT}\mathrm{E}\left[\phi(x(T))e^{-\alpha_{\mathrm{d}}x(T)}|x(0)=0\right] \\ &= e^{-rT}\int_{-\infty}^{+\infty}\phi(x)e^{-\alpha_{\mathrm{d}}x}p(x,T)dx \\ &= \frac{e^{-rT}}{2\pi}\int_{-\infty}^{+\infty}\widehat{\phi}(\xi)\widehat{p}^{*}(\xi+i\alpha_{\mathrm{d}},T)d\xi \\ &= e^{-rT}\mathcal{F}_{\xi\to x}^{-1}\left[\widehat{\phi}(\xi)\widehat{p}^{*}(\xi+i\alpha_{\mathrm{d}},T)\right](0). \end{aligned} \tag{2.93}$$

The expression in Eq. (2.93) looks very similar to Eq. (2.90) but the key difference is that the former is only valid at $x(t)=0$ whereas the latter is valid for all values of $x(t)$. This is acceptable for our pricing application where the current value of $x(t)$ is normalised using $S(t)$ so that it is equal to zero.

2.3 Numerical methods

The analytical methods described in Section 2.1 which are required to price the options described in Section 2.2 are described theoretically. However, many of the calculations involve some some expressions which cannot be solved in closed-form. Therefore their implementation requires the use of numerical approximations which we describe in the following sections.

2.3.1 Discrete Fourier transform and spectral filtering

The forward and inverse Fourier transforms in Eqs. (2.1) and (2.2) 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 numerical schemes both the x and ξ grids are centred around zero and their definition is determined by 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

$$x_j = j\Delta x, \quad j = -\frac{M}{2}, -\frac{M}{2}+1, \ldots, \frac{M}{2}-1. \tag{2.94}$$

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

$$\xi_k = k\Delta\xi, \quad k = -\frac{M}{2}, -\frac{M}{2}+1, \ldots, \frac{M}{2}-1. \tag{2.95}$$

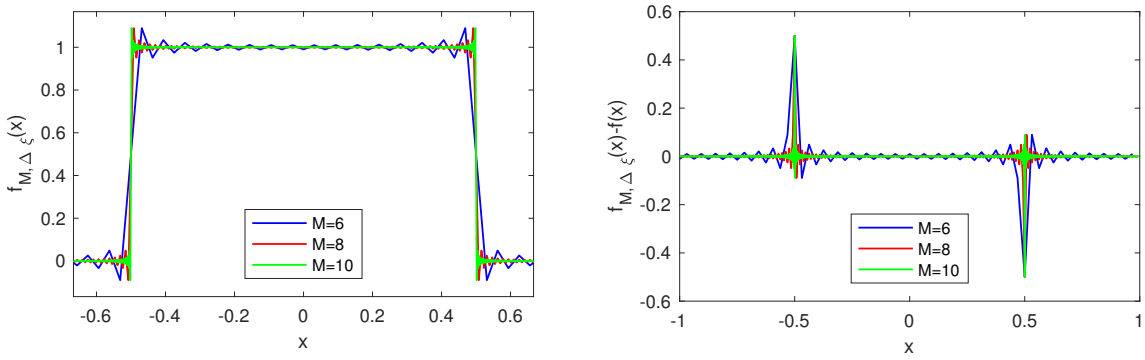

Figure 2.2: Illustration of the effect of the Gibbs phenomenon on a rectangular pulse recovered applying the inverse FFT with grid size M to $\mathrm{sinc}(\xi/2\pi)$. The left-hand plot shows the recovered waveform and the right-hand plot shows the error. On increasing M, the peak error at the discontinuity remains the same, the error away from the discontinuity reduces and the frequency of the oscillations increases.

The discrete Fourier transform is then

$$\widehat{f}_{M,\Delta x}(\xi_k) = \Delta x \sum_{j=-M/2}^{M/2-1} f(x_j)\, e^{ix_j\xi_k} \tag{2.96}$$

$$f_{M,\Delta\xi}(x_j) = \frac{\Delta\xi}{2\pi} \sum_{k=-M/2}^{M/2-1} \widehat{f}(\xi_k)\, e^{-ix_j\xi_k}. \tag{2.97}$$

There are several techniques to improve the efficiency of this calculation such as the fast Fourier transform (FFT) which reduces the computational complexity from $O(M^2)$ to $O(M \log M)$ by restricting the number of FFT points to integer powers of 2. In practice, we perform this calculation using the built-in MATLAB FFT function based on the FFTW library by Frigo and Johnson (1998).

2.3.1.1 Gibbs phenomenon

It can be seen in Eqs. (2.96) and (2.97) 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. (2.97), 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 a way this effect can be manifested is illustrated in Figure 2.2 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 $f(x_\mathrm{d})$, where $|x_d| = 0.5$, 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

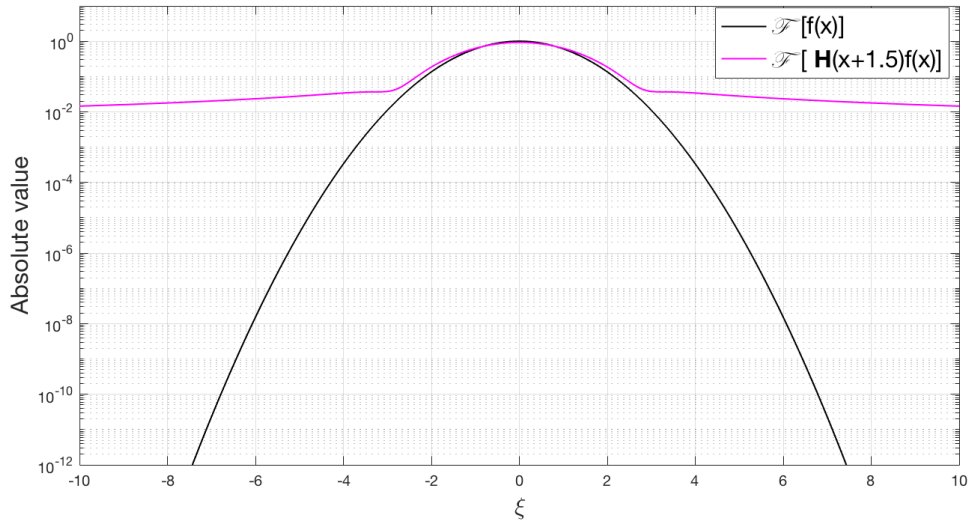

Figure 2.3: Illustration of the effect of the Gibbs phenomenon on a the Fourier transform of a standard normal distribution. Notice how the presence of a small discontinuity in the function changes the decay of the Fourier coefficients from exponential to polynomial.

discontinuity $f_{M,\Delta\xi}(x_\mathrm{d})$ will be the mean of the values immediately before and after the discontinuity, i.e. $f_{M,\Delta\xi}(x_\mathrm{d}) = \frac{1}{2}[f(x_\mathrm{d}^+) + f(x_\mathrm{d}^-)]$, and thus stays the same even as the value of M increases. In contrast, it can be observed from Figure 2.2 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 0^{th} order convergence at the discontinuity and away from the discontinuity we only achieve first order polynomial convergence, i.e. for $x_j \neq x_\mathrm{d}$, $|f(x_j) - f_{M,\Delta\xi}(x_j)| \sim O(1/M)$ (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 $O(1/\xi^k)$. This is illustrated in Figure 2.3 which shows the plot of a Fourier transform of a standard normal distribution and a plot of a Fourier transform of the same function multiplied by the Heaviside unit step function $H(x - x_0)$ with $x_0 = 1.5$. It can be seen that introducing even a very small step in a function has the effect of changing the decay of the Fourier transform to $O(1/\xi)$. 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.

2.3.2 Discrete Hilbert transform

We saw in Section 2.1.3 that 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,

$$\mathcal{H}\big[\widehat{f}(\xi)\big] = -i\,\mathcal{F}_{x\to\xi}\big[\operatorname{sgn}(x)\mathcal{F}^{-1}_{\xi\to x}\widehat{f}(\xi)\big]. \tag{2.98}$$

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) 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

$$\widehat{f}(\xi) = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi)\frac{\sin(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}, \tag{2.99}$$

$$\mathcal{H}\big[\widehat{f}(\xi)\big] = \sum_{k=-\infty}^{+\infty} \widehat{f}(k\Delta\xi)\frac{1-\cos(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}, \tag{2.100}$$

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. (2.99) and (2.100) are approximations whose error decays exponentially as $\Delta\xi$ decreases. In addition to discretisation, the infinite sum in Eq. (2.100) must also be truncated to the grid size M, so that the discrete approximation of the Hilbert transform becomes

$$\mathcal{H}\big[\widehat{f}(\xi)\big] \approx \sum_{k=-M/2}^{+M/2} \widehat{f}(k\Delta\xi)\frac{1-\cos(\pi(\xi-k\Delta\xi)/\Delta\xi)}{\pi(\xi-k\Delta\xi)/\Delta\xi}. \tag{2.101}$$

Feng and Linetsky (2008, 2009) showed that if $\widehat{f}(\xi)$ decays at least exponentially as $\xi\to\infty$, i.e. $\widehat{f}(\xi)\le\kappa\exp(-c|\xi|^\nu)$, then the error in the Hilbert transform and the Plemelj-Sokhotsky relations caused by truncating the infinite sum in Eq. (2.100) is also exponentially bounded. They also showed that if $\widehat{f}(\xi)$ is polynomially bounded as $\xi\to\infty$, i.e. $\widehat{f}(\xi)\le c|\xi|^\nu$, then the error caused by truncating the series is no longer exponentially bounded (Stenger, 1993; Feng and Linetsky, 2008, 2009).

2.3.3 Inverse z-transform

To recover the option price from the Wiener-Hopf and Spitzer identity calculations we must apply the inverse z-transform. This is a complex contour integral as shown in Eq. (2.18) and can be difficult to calculate in closed-form, therefore a numerical inverse transform is required for this step. We use the Fourier series based technique by Abate and Whitt (1992a,b) which has been used with great success in financial applications (see e.g. Fusai et al., 2006, 2011, 2012; Marazzina et al., 2012; Chang et al., 2013; Chen et al., 2014; Fusai et al., 2016).

The first step in the Abate and Whitt method is to approximate the integral numerically as

$$f(n) \approx \frac{1}{2n\rho^n}\left[\widetilde{f}(\rho) + 2\sum_{j=1}^{n-1}(-1)^j \operatorname{Re}\widetilde{f}\left(\rho e^{\frac{\pi j i}{n}}\right) + (-1)^n \widetilde{f}(-\rho)\right]. \tag{2.102}$$

The parameter ρ is selected to control the error; for an accuracy of $10^{-\lambda}$, we require $\rho = 10^{-\lambda/2n}$. In order to perform the inverse z-transform in Eq. (2.102), we must compute the weighted sum of $\widetilde{f}(q)$ for calculated for n distinct values of q. Thus the computation time of the operation described in Eq. (2.102) increases with n. In order to overcome this, Abate and Whitt used the Euler series acceleration which approximates the later coefficients of an alternating convergent series with a binomial average of the earlier points in the series: see the online supplementary material of Fusai et al. (2016) for an illustration of the accelerated convergence using this method.

The Abate and Whitt numerical method to recover $f(n)$ from $\widetilde{f}(q)$ via the inverse z-transform is therefore a two stage process:

1. Compute the partial sums

$$f(l) = \frac{1}{2}\widetilde{f}(\rho) + \sum_{j=1}^{l-1}(-1)^j \operatorname{Re}\widetilde{f}\left(\rho e^{\frac{\pi j i}{n}}\right). \tag{2.103}$$

2. Calculate the binomial average (or Euler transform) of these values from n_{E} to $n_{\mathrm{E}} + m_{\mathrm{E}}$ as

$$f(n) \approx \frac{1}{2^{m_{\mathrm{E}}} n\rho^n} \sum_{l=n_{\mathrm{E}}}^{n_{\mathrm{E}}+m_{\mathrm{E}}} \binom{m_{\mathrm{E}}}{l - n_{\mathrm{E}}} f(l) \tag{2.104}$$

The values of n_{E} and m_{E} are selected to give sufficient accuracy. Fusai et al. (2016) reported that experimental results showed that these could be set to $n_{\mathrm{E}} = 12$ and $m_{\mathrm{E}} = 20$. This means that for a number of dates N greater than 32, the computational time required by the inverse z-transform is independent of n.

2.3.4 Numerical pricing techniques

In Chapters 3 and 4 we adapt existing numerical pricing procedures by Fusai et al. (2016) and Feng and Linetsky (2008). Chapter 3 looks at the effect of the Gibbs phenomenon on the error performance of the pricing method for discrete barrier options. Chapter 4 extends the calculation of the Spitzer identities to continuous monitoring using option pricing as a motivating example. In addition, for the methods based on the pricing procedures by Fusai et al. (2016), we perform a detailed error analysis which gives additional insight into their original method. For a comprehensive background to the techniques we refer the interested reader to the aforementioned papers. However in order to provide a self contained guide to the work done in later chapters we include here brief summaries of the original implementation of the three pricing methods that we modify and improve in this thesis.

2.3.4.1 Spitzer based method for single-barrier options

We describe the pricing procedure devised by Fusai et al. (2016) 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.

1. Compute the characteristic function of the underlying asset process $\Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t)$, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

$$\Phi(\xi, q) := 1 - \Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t) = \Phi_{\oplus}(\xi, q)\Phi_{\ominus}(\xi, q) \tag{2.105}$$

   with q selected for $N-2$ dates according to the criteria specified by Abate and Whitt (1992b) for the inverse z-transform.

3. Decompose with respect to l

$$P(\xi, q) := \frac{\Psi(\xi + i\alpha_{\mathrm{d}})}{\Phi_{\ominus}(\xi, q)} = P_{l+}(\xi, q) + P_{l-}(\xi, q), \tag{2.106}$$

   and calculate

$$\widetilde{\widehat{p}}(\xi, q) := \Psi(\xi + i\alpha_{\mathrm{d}}) \frac{P_{l+}(\xi, q)}{\Phi_{\oplus}(\xi, q)}. \tag{2.107}$$

4. Calculate the option price as

$$v(0, 0) := \mathcal{F}^{-1}_{\xi \to x} \left[\widehat{\phi}^*(\xi) \mathcal{Z}^{-1}_{q \to N-2} \widetilde{\widehat{p}}(\xi, q) \right] (0), \tag{2.108}$$

   where $\widehat{\phi}^*(\xi)$ is the complex conjugate of the Fourier transform of the damped payoff

function given in Eq. (2.92).

Notice that in the numerical implementation described above, the Spitzer identity is calculated for $N-2$ dates, with the inclusion of the characteristic function $\Psi(\xi+i\alpha_{\mathrm{d}},\Delta t)$ in Eqs. (2.106) and (2.107) applying the first and final dates respectively. In Fusai et al. (2016), results were presented showing exponential error convergence for general Lévy processes and in Section 3.1.1 of Chapter 3 we explore the reason for this and develop bounds for the error convergence.

2.3.4.2 Spitzer based method for double-barrier options

The pricing procedure by Fusai et al. (2016) for double-barrier options is very similar to the method for the single-barrier options described in Section 2.3.4.1, 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 are the same as those for single-barrier down-and-out options described in Section 2.3.4.1 with the exception of Step 3 which is now replaced by the fixed-point algorithm

3 (a) Set $J_-(\xi,q)=J_{l-}(\xi,q)=0$.

(b) Decompose with respect to l

$$P_{\mathrm{c}}(\xi,q):=\frac{\Psi(\xi+i\alpha_{\mathrm{d}})-\Phi_{\oplus}(\xi,q)J_{u+}(\xi,q)}{\Phi_{\ominus}(\xi,q)}=P_{l+}(\xi,q)+P_{l-}(\xi,q), \tag{2.109}$$

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

(c) Decompose with respect to u

$$Q_{\mathrm{c}}(\xi,q):=\frac{\Psi(\xi+i\alpha_{\mathrm{d}})-\Phi_{\ominus}(\xi,q)J_{l-}(\xi,q)}{\Phi_{\oplus}(\xi,q)}=Q_{u+}(\xi,q)+Q_{u-}(\xi,q), \tag{2.110}$$

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

(d) Calculate

$$\widetilde{\widehat{p}}(\xi,q):=\Psi(\xi+i\alpha_{\mathrm{d}})\frac{\Psi(\xi+i\alpha_{\mathrm{d}})-\Phi_{\ominus}(\xi,q)J_{l-}(\xi,q)-\Phi_{\oplus}(\xi,q)J_{-}(\xi,q)}{\Phi(\xi,q)}. \tag{2.111}$$

(e) If the difference between the new and the old value of $\widetilde{\widehat{p}}(\xi,q)$ 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.

Similarly to the procedure for single-barrier options described in Section 2.3.4.1, the Spitzer identity is calculated for $N-2$ dates and the first and last date are applied using the

characteristic function in Eqs. (2.109)–(2.111). In contrast to the method for pricing single-barrier options, Fusai et al. (2016) present result showing polynomial error convergence only. The reason for this is explored in more detail in Section 3.1.2 and we present an updated method with exponential error convergence in Section 3.2.

2.3.4.3 Feng and Linetsky method

As shown in Section 2.1, we can use Eqs. (2.26)–(2.28) to obtain the Fourier transform of the part of a function above or below a barrier or between two barriers. This property of 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:

$$v(x, t_{n-1}) = \int_l^u v(x', t_n) k(x - x', \Delta t) dx'. \tag{2.112}$$

Here $v(x, t_N) = \phi(x) e^{-\alpha_d x}$, i.e. the payoff of the option, and $k(\cdot, \Delta t)$ denotes the transition density of the underlying process with step size Δt and $\Psi(\xi, \Delta t)$ is its characteristic function. Using the convolution theorem together with the Hilbert transform the relationship between the price at two successive dates can be expressed as

$$\widehat{v}(\xi, t_{n-1}) = \frac{1}{2} \left\{ \Psi(\xi + i\alpha_d, \Delta t) \widehat{v}(\xi, t_n) + e^{il\xi} i\mathcal{H} \left[e^{-il\xi} \Psi(\xi + i\alpha_d, \Delta t) \widehat{v}(\xi, t_n) \right] \right\} \tag{2.113}$$

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

$$\widehat{v}(\xi, t_{n-1}) = \frac{1}{2} \left\{ e^{il\xi} i\mathcal{H} \left[e^{-il\xi} \Psi(\xi + i\alpha_d, \Delta t) \widehat{v}(\xi, t_n) \right] - e^{iu\xi} i\mathcal{H} \left[e^{-iu\xi} \Psi(\xi + i\alpha_d, \Delta t) \widehat{v}(\xi, t_n) \right] \right\} \tag{2.114}$$

for a double-barrier option.

Feng and Linetsky (2008) presented results for this method showing exponential error convergence for underlying processes with exponential characteristic functions and polynomial error convergence for the VG process. This was explained in detail in the original paper and in Section 3.1.3 we present a new method with improved error performance for the VG process.

2.3.5 Practical implementation of the numerical techniques

The work in this thesis is primarily focused on the challenges associated with the numerical pricing of exotic options driven by exponential Lévy processes. For this we have made the implicit assumption that the parameters controlling the underlying Lévy processes are known and we have concentrated on developing the new numerical methods described herein. Indeed, in Appendix A we list process parameters which are based on values used in some of the existing literature on exotic options, thus allowing a fair comparison with

previous numerical methods.

However, if the methods developed in this thesis are to be used in a trading environment we must first calibrate the Lévy process parameters to market data before using our numerical pricing methods. Although this is not the primary focus of this thesis we direct the reader to the work by Cont and Takov (2004); Cont and Tankov (2006), which deals with the calibration of several exponential Lévy processes including all the ones used herein. Also of interest is the later work by Askari and Krichene (2008) which investigates calibrating oil price dynamics to Merton jump-diffusion and VG processes and by Ramezani and Zeng (2007) which considers the double exponential jump process by Kou (2002) as a model for the NASDAQ and S&P500 indices.

Chapter 3

Hilbert transform, spectral filters and option pricing

In this chapter we modify the pricing method for discretely monitored barrier options by Fusai et al. (2016) and Feng and Linetsky (2008), described in Section 2.3.4. In the following we refer to the method by Fusai et al. (2016) as FGM and the method by Feng and Linetsky (2008) as FL. We first perform a detailed error examination into the existing techniques and then propose improved methods based on spectral filtering. Lastly, we present results showing that the improved method now achieves exponential convergence whilst retaining the advantages of the original method, i.e. a CPU time independent of the number of monitoring dates for the FGM method and applicability to general Lévy processes for both methods.

3.1 Error performance of the pricing procedures

In this section we investigate the error convergence of the calculations making up the original pricing procedures without spectral filters and show bounds for the individual steps. As described in Chapter 2, these methods are based on the sinc-based fast Hilbert transform by Stenger (1993), who showed that the discretisation error is exponentially convergent with grid step size, under conditions which are met, described in more detail in Section 2.3.2. Therefore, here, our error analysis is concerned with the truncation errors of each step as it was shown by Feng and Linetsky (2008) that this can vary with the shape of the function in the Fourier domain. As part of the error investigation, the effect of each step in the procedure on the shape of the output function as $|\xi| \to \infty$ in the Fourier domain is also considered, as this largely determines the truncation error 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, thus we concentrate our investigation on the later steps in the pricing procedures.

In the following calculations, the damping factor α_{d} is omitted to make the notation more concise. This is appropriate as the value of $i\alpha_{\mathrm{d}}$ becomes insignificant as $|\xi| \to \infty$. In the error calculations below, c_{n} (where $\mathrm{n} = 1, 2, 3...$) are included as positive constants.

3.1.1 Pricing single-barrier options with the VG process using the Spitzer identity

Following the calculation of the characteristic function, the next step in the pricing procedure described in Section 2.3.4.1 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 VG process, as $|\xi| \to \infty$, $q\Psi(\xi, \Delta t) \sim qe^{-\Delta t \xi^2}$ which quickly becomes very small. Thus we can say that as $|\xi| \to \infty$, $|\log[1 - q\Psi(\xi, \Delta t)]| < ce^{-\kappa \Delta t \xi^2}$ 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), described in Section 2.3.2, the output of the decomposition of $\log \Phi(\xi, q)$ has exponential error performance for exponentially decaying characteristic functions.

By contrast the characteristic function for the VG process is

$$\Psi(\xi, t) = \left(1 - i\nu\theta\xi + \frac{1}{2}\nu^2\xi^2\sigma^2\right)^{-t/\nu}. \tag{3.1}$$

When the value of ξ is very large then $\Psi(\xi, \Delta t)$ is dominated by $\xi^{-2\Delta t/\nu}$, so when $|\xi| \to \infty$, $|\log[1 - q\Psi(\xi, \Delta t)]| < c\xi^{-2\Delta t/\nu}$ and we can use this to 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. (2.100) and $f_{\Delta\xi,M}(\xi)$ as the output of the truncated sum from Eq. (2.101),

$$\begin{aligned}|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| <& c_1\Delta\xi \left| \sum_{k>M/2} \frac{(k\Delta\xi)^{-2\Delta t/\nu}\left(1-\cos\left(\pi\frac{\xi-k\Delta\xi}{\Delta\xi}\right)\right)}{\xi - k\Delta\xi} \right. \\ &\left. + \sum_{k<-M/2} \frac{(k\Delta\xi)^{-2\Delta t/\nu}\left(1-\cos\left(\pi\frac{\xi-k\Delta\xi}{\Delta\xi}\right)\right)}{\xi - k\Delta\xi} \right| \\ <& c_1\Delta\xi \sum_{k>M/2} (k\Delta\xi)^{-2\Delta t/\nu} \left| \frac{1-\cos\left(\pi\frac{\xi-k\Delta\xi}{\Delta\xi}\right)}{\xi - k\Delta\xi} + \frac{1-\cos\left(\pi\frac{\xi+k\Delta\xi}{\Delta\xi}\right)}{\xi + k\Delta\xi} \right|. \end{aligned} \tag{3.2}$$

The two cosines are equal because the difference of their arguments is $2k\pi$ and so

$$\begin{aligned}|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| <& c_2\Delta\xi \sum_{k>M/2} (k\Delta\xi)^{-2\Delta t/\nu} \left| \frac{1}{\xi - k\Delta\xi} + \frac{1}{\xi + k\Delta\xi} \right| \\ <& c_3\Delta\xi \sum_{k>M/2} \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{|\xi^2 - (k\Delta\xi)^2|} < c_4\Delta\xi \sum_{k>M/2} \frac{(k\Delta\xi)^{-2\Delta t/\nu}}{(k\Delta\xi)^2} \\ <& c_5 \int_{M\Delta\xi/2}^{+\infty} \xi_k^{-\left(\frac{2\Delta t}{\nu}+2\right)} d\xi_k < c_6 (M\Delta\xi)^{-\left(\frac{2\Delta t}{\nu}+1\right)}. \end{aligned} \tag{3.3}$$

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

$$\left| \frac{e^{f_{\Delta\xi}(\xi)} - e^{f_{\Delta\xi,M}(\xi)}}{e^{f_{\Delta\xi}(\xi)}} \right| < c_7 \left[1 - e^{c_8 (M\Delta\xi)^{-\left(\frac{2\Delta t}{\nu}+1\right)}} \right]. \tag{3.4}$$

For large M this converges as $O\left(M^{-(2\Delta t/\nu+1)}\right)$; thus the error convergence of the factorisation is polynomial. The expression in Eq. (3.3) gives the error at fixed values of ξ, i.e. the chosen grid points; therefore ξ can be absorbed into c_2. Moreover, in the final price calculation the Spitzer identities are multiplied by the payoff and the characteristic function which both decay as $|\xi|$ increases and therefore the errors close to $\xi = 0$ will have the largest influence on the final error of the solution. However, as M increases, our range of ξ increases and so we should consider the effect of errors at large values of ξ on the error of the next step in the calculation. As explained below, we multiply the input to the subsequent Hilbert transform by the characteristic function and if the number of grid points increases from M^* to M, where $M^* \ll M$, then the additional error from these points is bounded as

$$c_9 \sum_{j=M^*}^{M} \frac{j\Delta\xi M^{-(4\Delta t/\nu+1)}}{\xi - j\Delta\xi} < c_{10} \sum_{j=M^*}^{M} M^{-(4\Delta t/\nu+1)} < c_{11} M^{-(4\Delta t/\nu)}. \tag{3.5}$$

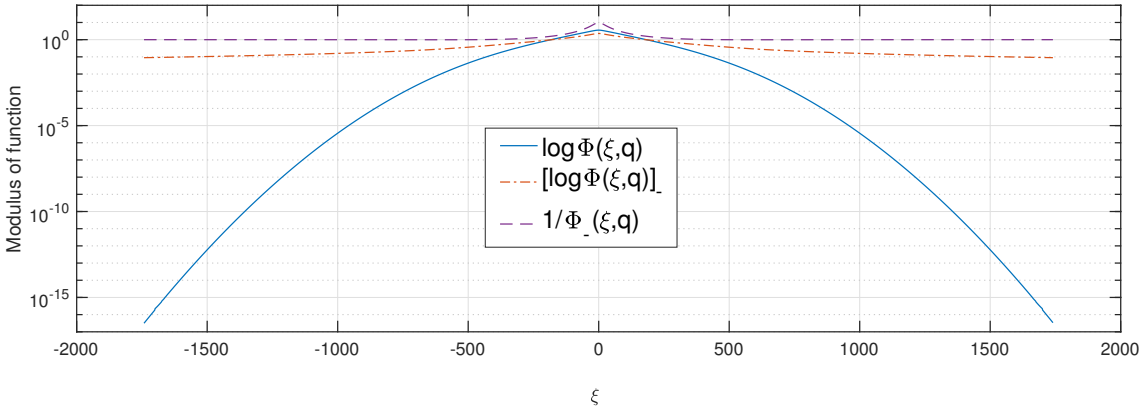

Figure 3.1: Input and output functions for the factorisation of $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t)$ with the Kou process for $q = \rho$.

Depending on the parameters of the VG process, this decreases more or less rapidly than the original bound, and if we were to select our parameters according to the requirement in Feng and Linetsky (2008) of $2\Delta t/\nu > 1$ then $O\left(M^{-(2\Delta t/\nu+1)}\right)$ will dominate. However, regardless of parameter selection, the error converges polynomially with M.

The requirement to multiply the output of the the factorisation by the characteristic function is due to its shape in the Fourier domain as this will influence the error performance of the subsequent step. Figure 3.1 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 VG process is used then the input to the decomposition is only polynomially decaying and thus we again have polynomial error convergence for this stage.

3.1.2 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 3.1.1, with 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 $|\xi| \to \infty$. 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 $\xi \to \infty$. 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. (2.100) to give Eq. (2.101) is no longer exponentially bounded. A bound for this error is

$$\begin{aligned} |f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< c_1 \Delta\xi \sum_{k>M/2} \frac{1}{(k\Delta\xi)^2} \\ &< c_2 \int_{M\Delta\xi/2}^{+\infty} \frac{1}{\xi_k^2} d\xi_k < c_3 \frac{1}{M\Delta\xi}. \end{aligned} \tag{3.6}$$

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. (3.6) 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.

3.1.3 Feng and Linetsky pricing method with the VG process

The FL method is described in Eqs. (2.113) and (2.114), which show how the Hilbert transform is applied for each monitoring date. As explained in Section 3.1.2, 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 $\xi \to \infty$. 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 VG process, then a polynomially convergent error will be achieved.

3.2 Updated pricing procedures

In the previous section we showed how the Gibbs phenomenon causes polynomial error convergence in the original pricing techniques. In this section we show how we can use spectral filtering to achieve exponential convergence whilst retaining the advantages of the original methods.

3.2.1 Spectral filtering

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). David Gottlieb, Chi-Wang Shu and their colleagues published several results showing that it is possible to solve the Gibbs phenomenon by transforming the Fourier series coefficients to a series which converges more rapidly using other basis functions such as Gegenbauer polynomials (Gottlieb et al., 1992; Gottlieb and Shu, 1995a, 1996, 1995b, 1997). 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 was defined as a function $\sigma(\eta)$ supported on $\eta \in [-1, 1]$ with the following properties:

$$\begin{aligned} &\text{a) } \sigma(0) = 1,\ \sigma^{(l)}(0) = 0, \\ &\text{b) } \sigma(\eta) = 0 \text{ for } |\eta| = 1, \\ &\text{c) } \sigma(\eta) \in C^{p-1}. \end{aligned} \tag{3.7}$$

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

$$\sigma(\eta) = e^{-\vartheta \eta^p}, \tag{3.8}$$

where p is even and positive. This does not strictly meet criterion b in Eq. (3.7) 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.

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

$$\sigma(\eta) = \begin{cases} 0, & \eta \le \eta_1, \qquad \eta_1 = -1, \\ \frac{1}{e^{\,z(\eta)}+1},\ z(\eta) = \frac{\eta_2-\eta_1}{\eta-\eta_1} + \frac{\eta_2-\eta_1}{\eta-\eta_2}, & \eta_1 < \eta < \eta_2,\ \eta_2 = \epsilon - 1, \\ 1, & \eta_2 \le \eta \le \eta_3,\ \eta_3 = 1-\epsilon, \\ \frac{1}{e^{\,z(\eta)}+1},\ z(\eta) = \frac{\eta_3-\eta_4}{\eta-\eta_3} + \frac{\eta_3-\eta_4}{\eta-\eta_4}, & \eta_3 < \eta < \eta_4,\ \eta_4 = 1, \\ 0, & \eta \ge \eta_4. \end{cases} \tag{3.9}$$

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 3.2.

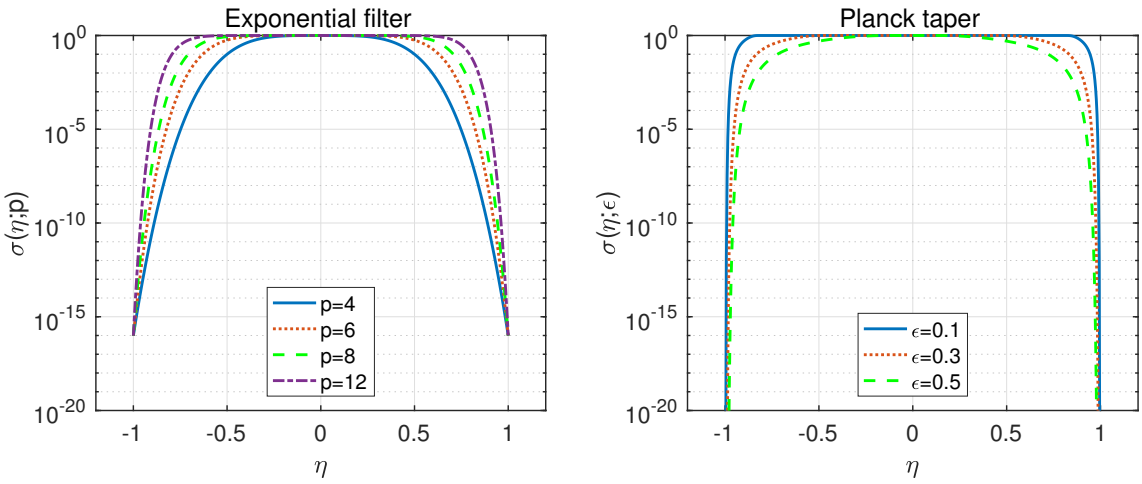

Figure 3.2: Shape of the exponential filter (left) and Planck taper (right) with different parameter values.

3.2.2 Single-barrier options with Spitzer identities

Fusai et al. (2016) showed that the Spitzer-based method for single-barrier options could achieve exponential convergence with a wide range of Lévy processes. However, the performance of the method with the VG process only achieved polynomial convergence. This is consistent with the error behaviour of the discrete Hilbert transform with the VG process, as explained in Section 2.3.2 above. Section 3.1 explains in more detail how the error performance is bounded when this process is used.

In order to improve the result, we modified the method in Section 2.3.4.1 by multiplying the characteristic function by a spectral filter $\sigma(\eta)$ so that the input to both the factorisation and decomposition steps decay exponentially. The modified pricing procedure is

1. Compute the characteristic function $\Psi(\xi + i\alpha_\mathrm{d})$, where α_d is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

torise

$$\Phi(\xi, q) := 1 - \Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t)\sigma\left(\frac{\xi}{\xi_{\max}}\right) = \Phi_{\oplus}(\xi, q)\Phi_{\ominus}(\xi, q), \tag{3.10}$$

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

3. Decompose with respect to l

$$P(\xi, q) := \frac{e^{-il\xi}\Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t)\sigma\left(\frac{\xi}{\xi_{\max}}\right)}{\Phi_{\ominus}(\xi, q)} = P_{l+}(\xi, q) + P_{l-}(\xi, q), \tag{3.11}$$

and calculate

$$\widetilde{\widehat{p}}(\xi, q) := \Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t)\sigma\left(\frac{\xi}{\xi_{\max}}\right)\frac{P_{l+}(\xi, q)}{\Phi_{\oplus}(\xi, q)}. \tag{3.12}$$

4. Calculate the option price as

$$v(0, 0) := \mathcal{F}^{-1}_{\xi\to x}\left[\widehat{\phi}^{*}(\xi)\mathcal{Z}^{-1}_{\xi\to x}\widetilde{\widehat{p}}(\xi, q)\right](0), \tag{3.13}$$

where $\widehat{\phi}^{*}(\xi)$ is the complex conjugate of the Fourier transform of the damped payoff function given in Eq. (2.92).

3.2.3 Double-barrier options

Unlike the direct method for single-barrier options described in Section 2.3.4.1, the iterative method described in Section 2.3.4.2 is limited to polynomial error convergence for all processes and in Section 3.1 above, we have shown 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. Step 3 in the single-barrier pricing procedure in Section 2.3.4.1 is replaced by the filtered fixed-point algorithm

3 (a) Set $J_{u+}(\xi, q) = J_{l-}(\xi, q) = 0$.

(b) Decompose with respect to l

$$P(\xi, q) := \sigma\left(\frac{\xi}{\xi_{\max}}\right)\left[\frac{\Psi(\xi + i\alpha_{\mathrm{d}}) - \Phi_{\oplus}(\xi, q)J_{u+}(\xi, q)}{\Phi_{\ominus}(\xi, q)}\right] = P_{l+}(\xi, q) + P_{l-}(\xi, q), \tag{3.14}$$

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

(c) Decompose with respect to u

$$Q(\xi, q) := \sigma\left(\frac{\xi}{\xi_{\max}}\right)\left[\frac{\Psi(\xi + i\alpha_{\mathrm{d}}) - \Phi_{\ominus}(\xi, q)J_{l-}(\xi, q)}{\Phi_{\oplus}(\xi, q)}\right] = Q_{u+}(\xi, q) + Q_{u-}(\xi, q), \tag{3.15}$$

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

(d) Calculate

$$\widetilde{\widehat{p}}(\xi,q) := \Psi(\xi + i\alpha_{\mathrm{d}}) \frac{\Psi(\xi + i\alpha_{\mathrm{d}}) - \Phi_{\ominus}(\xi,q) J_{l-}(\xi,q) - \Phi_{\oplus}(\xi,q) J_{u+}(\xi,q)}{\Phi(\xi,q)}. \quad (3.16)$$

(e) If the difference between the new and the old value of $\widetilde{\widehat{p}}(\xi,q)$ 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.

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 VG process, this method will also be subject to the same limitations on accuracy as described in Section 3.1.1 for single-barrier options. Therefore, if we wish to use this scheme with the VG process, we must also apply filtering to the factorisation step as shown in Eq. (3.10). Therefore, instead, we replace Step 3 in the single-barrier pricing procedure in Section 3.2.2 by the filtered fixed-point algorithm above. Numerical results with the updated method for the double-barrier case are shown in Section 3.3.

3.2.4 Feng and Linetsky method

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 by Feng and Linetsky (2008), explained in Section 2.3.4.3. 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 $|\xi| \to \infty$. Therefore, poor error performance is achieved for the VG process which has a characteristic function which only reduces polynomially as $|\xi| \to \infty$. 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. (2.113) and (2.114) by

$$\begin{aligned}\widehat{v}(\xi,t_{n-1}) =&\frac{1}{2}\left\{\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi+i\alpha_{\mathrm{d}},\Delta t)\widehat{v}(\xi,t_n)\right.\\ &\left.+e^{il\xi}i\mathcal{H}\left[e^{-il\xi}\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi+i\alpha_{\mathrm{d}},\Delta t)\widehat{v}(\xi,t_n)\right]\right\}, \qquad (3.17)\\ \widehat{v}(\xi,t_{n-1}) =&\frac{1}{2}\left\{e^{il\xi}i\mathcal{H}\left[e^{-il\xi}\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi+i\alpha_{\mathrm{d}},\Delta t)\widehat{v}(\xi,t_n)\right]\right.\\ &\left.-e^{iu\xi}i\mathcal{H}\left[e^{-iu\xi}\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi+i\alpha_{\mathrm{d}},\Delta t)\widehat{v}(\xi,t_n)\right]\right\}. \qquad (3.18)\end{aligned}$$

3.2.5 Removing the first and last dates from the pricing procedure

As described in Sections 2.3.4.1 and 2.3.4.2, the first and last dates were removed from the Spitzer scheme and applied via convolution. In the case of single-barrier options this has the effect of ensuring that the Fourier coefficients are exponentially bounded on the input to the decomposition and the final inverse Fourier transform. However, this does not work for the fixed-point algorithm, thus necessitating the use of additional filtering at every iteration. With regards to the discretely monitored barrier options, the removal of the characteristic function from the pricing procedure is not especially advantageous. It would enable the scheme to be used for contracts with 2 dates or less, but these are extreme cases which would anyway probably be more suited to the use of the FL method due to the very low number of monitoring dates. However, in order to move to the continuously monitored case, as described in Chapter 4, we do not have a characteristic function to use as $\Delta t = 0$ and therefore it is useful to show that the pricing procedure works well, even without the use of the characteristic function to provide additional smoothing.

There are three main alterations associated with this change. First we must remove the characteristic function in Eqs. (3.14), (3.15) and (3.16) and so any smoothing on the input to the next step is solely provided by the spectral filter. Secondly the the inverse z-transform will now go from q to N, rather than $N-2$, dates. The third change concerns the form of the payoff function. When the final date step is applied by convolution, then the barrier is not applied to the probability distribution at the final date. This does not cause problems as the the barrier is applied via the payoff function so the pricing method gives the correct value. However, when all N dates are inside the Spitzer scheme then the barrier is applied by the probability distribution as illustrated in Figure 3.3. It was found that if the barrier is also applied by the payoff, then only polynomial error performance is achieved. However, if the barrier for the payoff is moved to the truncation value of x, rather than the upper barrier, then exponential performance is recovered.

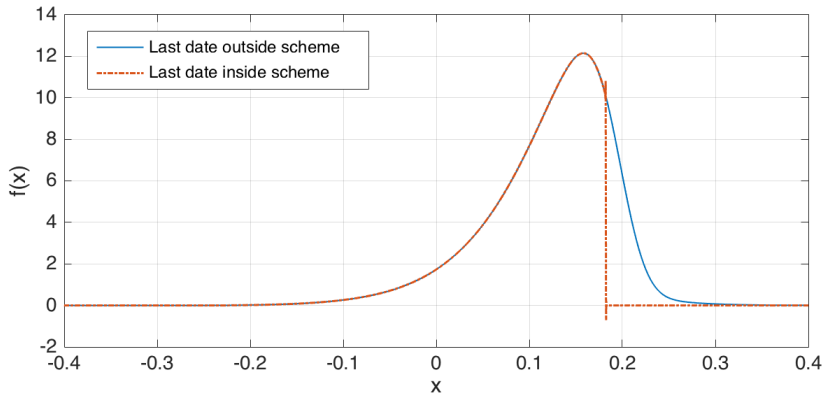

Figure 3.3: Comparison of the probability density with and without the final date inside the Spitzer scheme.

3.2.6 Error performance with filtering on the sinc-based Hilbert transform

The multiplication by a filter with exponentially decaying coefficients as $|\xi| \to \infty$ 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 3.3 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) suggested that the faster convergence they see may be due to the alternating nature of the Fourier coefficients.

3.3 Numerical results

We performed numerical tests using the pricing schemes updated to include filtering, as described in Section 3.2. The results for the FGM method for double-barrier options with exponentially decaying characteristic functions are presented in Section 3.3.1 and Section 3.3.2 contains results for all methods with the VG process. Details of the contract and the model parameters are included in Table A.1 in Appendix A.

3.3.1 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 3.2.3. 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 3.4 shows results for the Kou process and Figure 3.5 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 3.1 and 3.2 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 3.4 and 3.5 show how the convergence of the numerical techniques changes with the grid size and Figures 3.6 and 3.7 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 due to an increase in accuracy of the calculation, as shown in Figures 3.4 and 3.5. However, this also improves the algorithm computationally as it now reaches the required accuracy in a smaller number of iterations than the original FGM 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 greater than 10^{-8}; for higher number of dates, the filtered FGM method shows the best performance for errors greater than 10^{-8}. 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 3.2.6: 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.

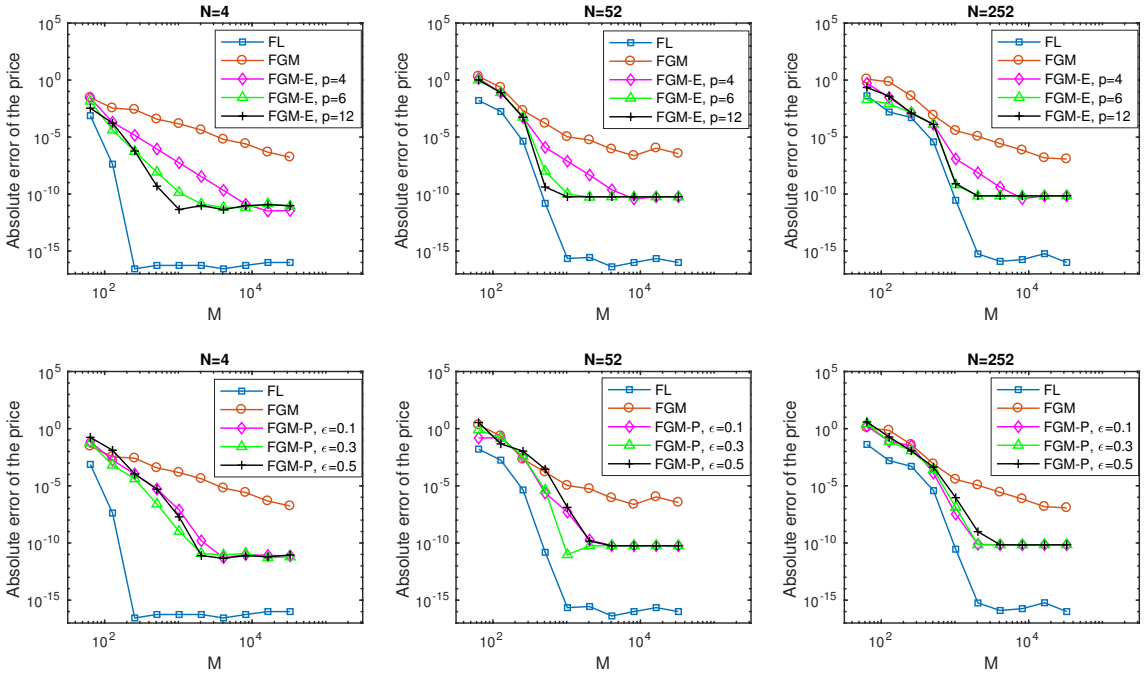

Figure 3.4: Error vs. grid size M for the Kou process and varying number of monitoring dates N. The filter improves the performance of the FGM method from polynomial to exponential. The best results are obtained with an exponential filter of order $p = 12$.

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

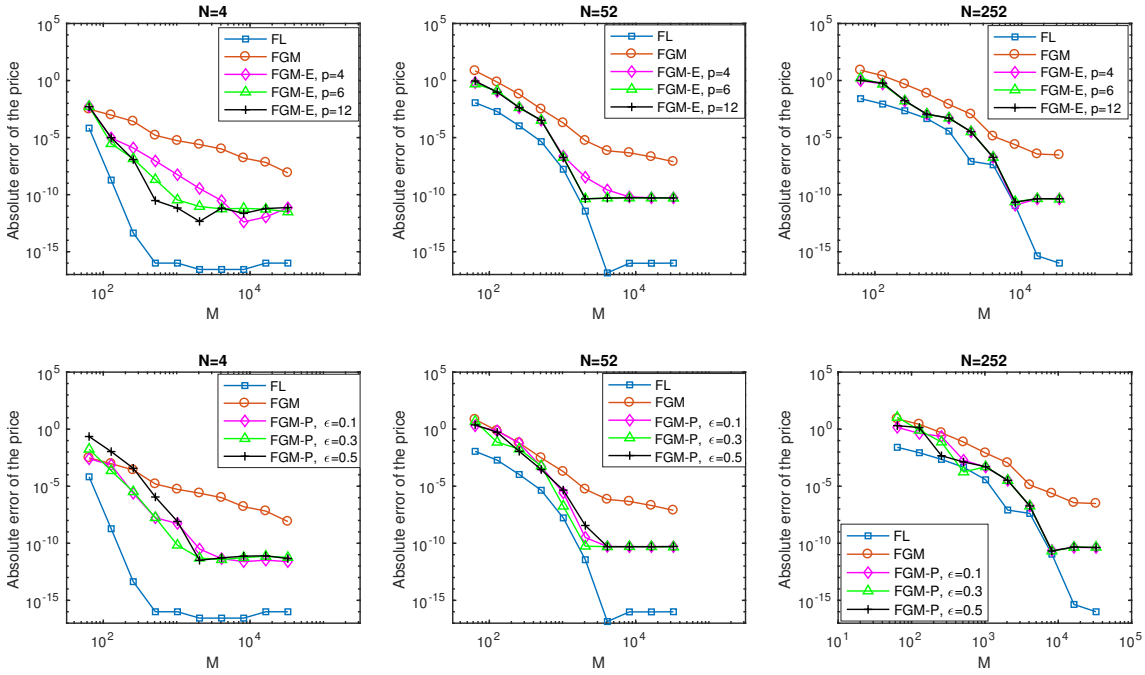

Figure 3.5: Error vs. grid size M with the NIG process and varying number of monitoring dates N. The filter improves the performance of the FGM method from polynomial to exponential. The best results are obtained with an exponential filter of order $p = 12$.

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

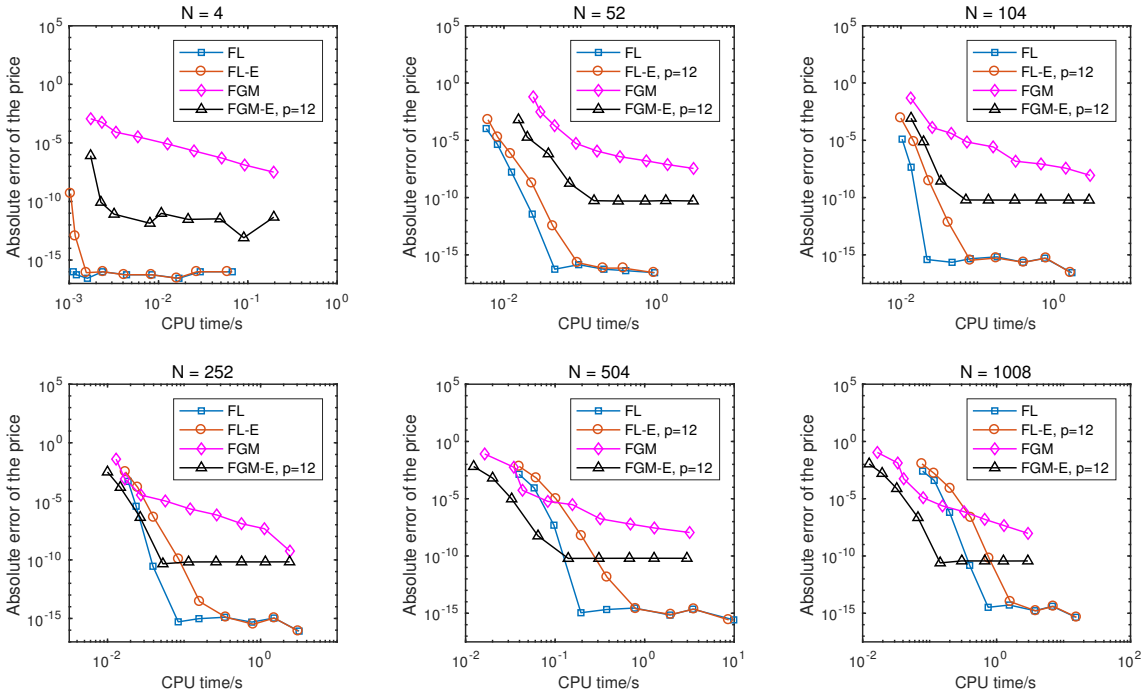

Figure 3.6: Error vs. CPU time for a double-barrier option with the Kou process and varying numbers of monitoring dates N. The filter improves the FGM method for all N; FGM-F is the fastest method for an error of 10^{-8} with $N > 252$.

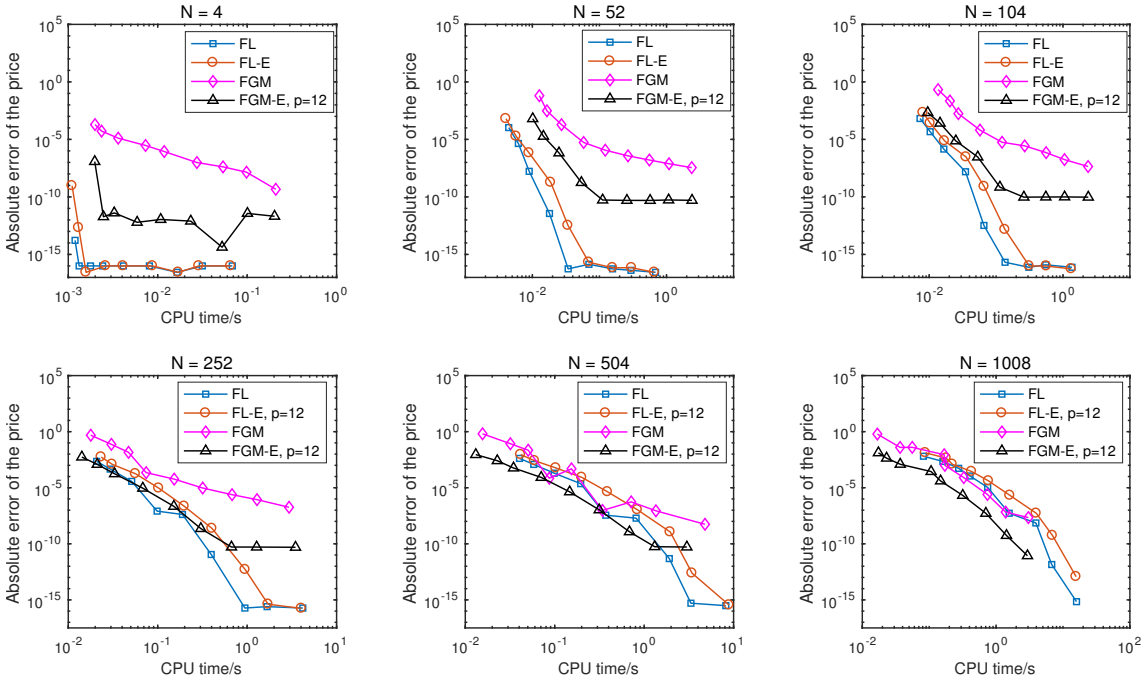

Figure 3.7: Error vs. CPU time for a double-barrier option with the NIG process and varying numbers of monitoring dates N. The filter improves the FGM method for all N; FGM-F is the fastest method for an error of 10^{-8} with $N \geq 504$.

3.3.2 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 VG process. Figures 3.8 and 3.9 show the results of tests for single and double-barrier options where we have applied exponential filtering, as described in Section 3.2.

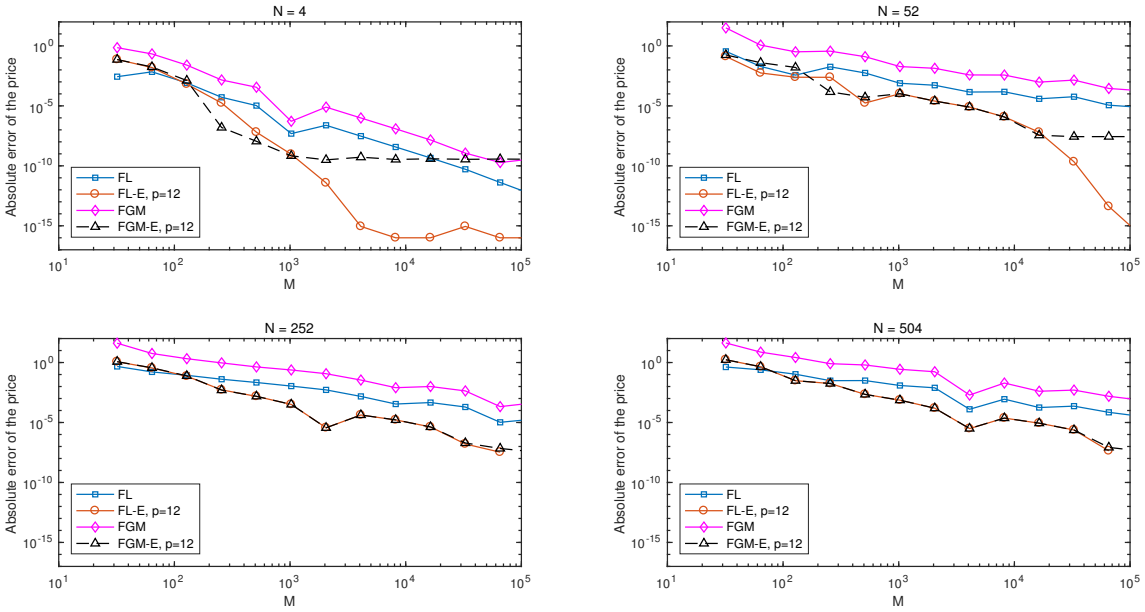

Figure 3.8: Error vs. grid size M for a single-barrier down-and-out option with the VG process and varying numbers of monitoring dates N. The filter improves both the FGM and FL methods, with the FL-F method performing best at low numbers of dates.

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 3.3.1. 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 3.10 and 3.11 where 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.

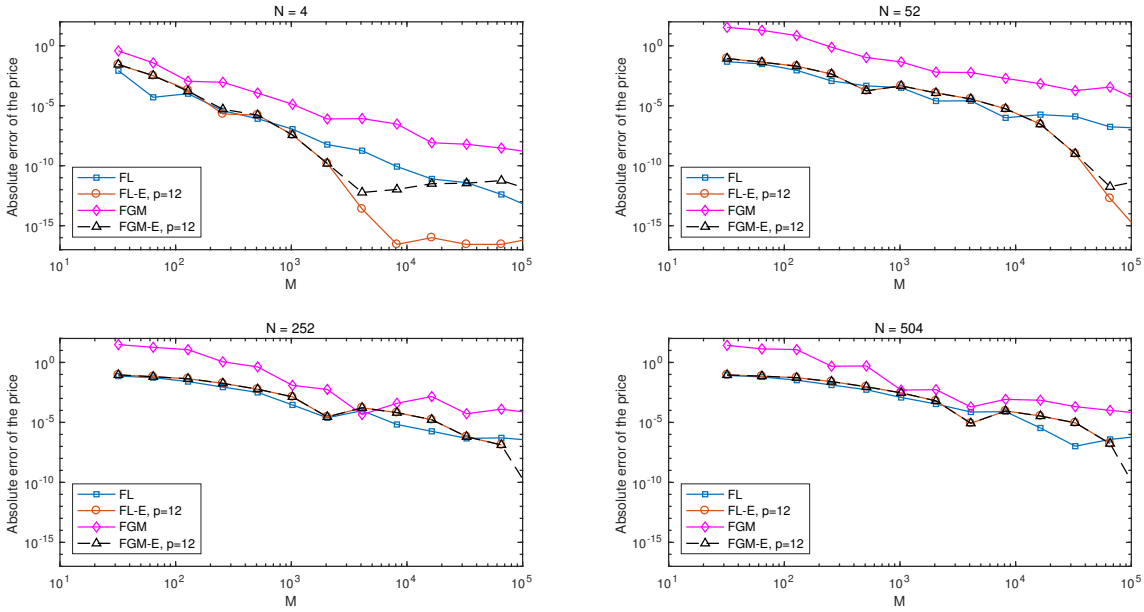

Figure 3.9: Error vs. grid size M for a double-barrier option with the VG process and varying numbers of monitoring dates N. The filter improves both the FGM methods for all numbers of monitoring dates and improved the FL method for low numbers of dates.

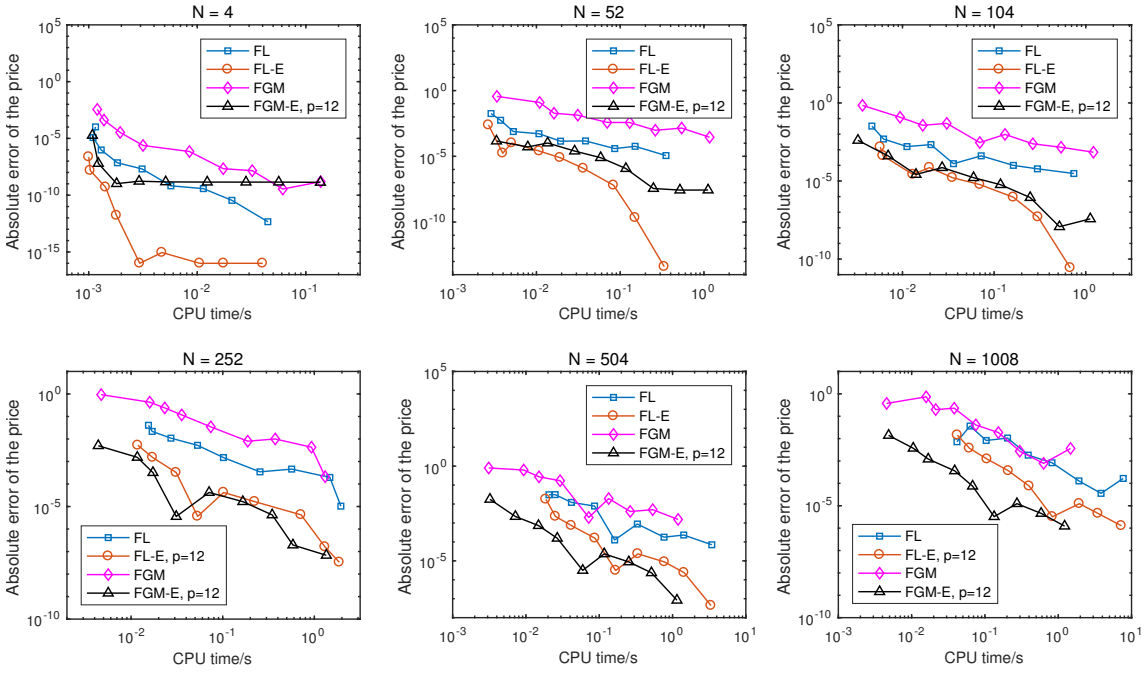

Figure 3.10: Error vs. CPU time for a single-barrier option with the VG process and varying numbers of monitoring dates N. The best performance of the new filtered methods, FL-F and FGM-F, either equals or exceeds the performance of the existing methods over all numbers of dates.

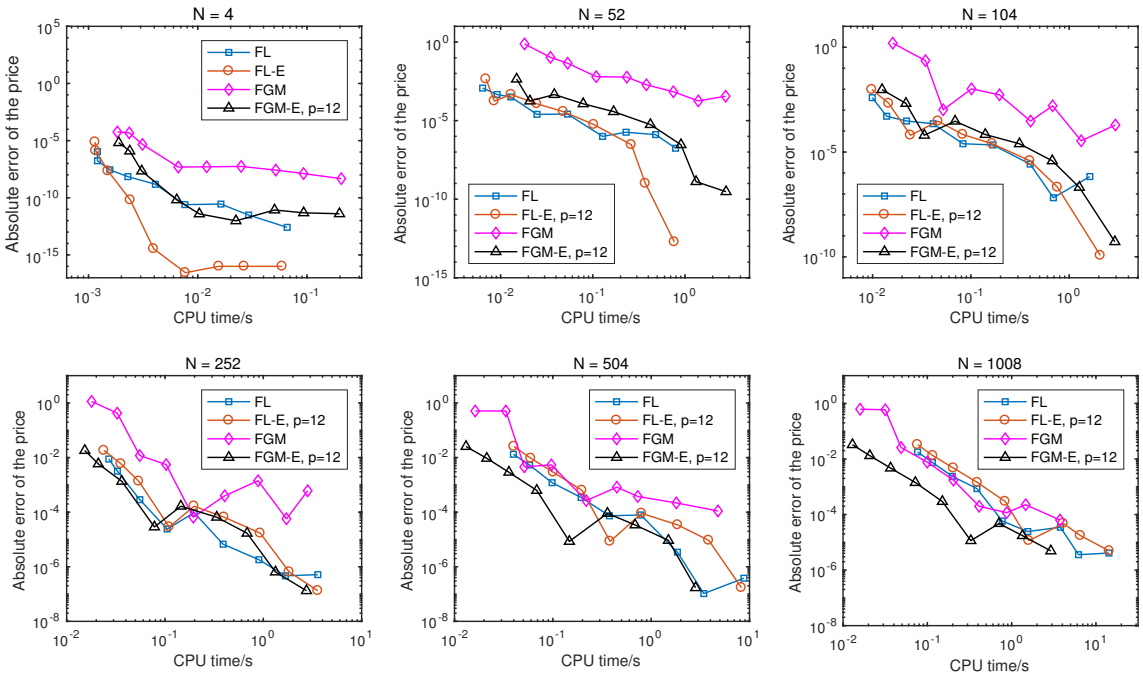

Figure 3.11: Error vs. CPU time for a double-barrier option with the VG process and varying numbers of monitoring dates N. The best performance of the new filtered methods, FL-F and FGM-F, either equals or exceeds the performance of the existing methods over all numbers of dates.

3.3.3 Summary of results

Table 3.3 shows a summary of the best performing methods in terms of CPU time for different processes and types of options. The table gives the fastest method for an error of 10^{-8} for the NIG and Kou processes and, due the slower convergence of all methods with the VG process, entries marked with an asterisk show the quickest method for an error of 10^{-5}. The entries in the table are coloured to give an at a glance guide to the improvements achieved by spectral filtering. Green entries mean that a filtered method provides the best performance. Blue entries mean that the performance of the filtered methods equals, but does not exceed, the best performance of an existing method. The remaining red entries are the few cases where an existing method performs best. Therefore for 75% of the test cases, the new methods equalled or improved upon the performance of the existing method. However for options traded in the market it most common for contracts to be written over several years with daily, or latterly even intra-day, monitoring. Therefore in terms of the options that are traded on the market, the proportion for which our new methods provide the best performance is likely to be far in excess of this percentage.

	Single barrier	Double barrier		
Dates	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 3.3: Quickest method for an error of 10^{-8}. Due to the slower convergence of all methods with the VG 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.

3.4 Conclusions

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 VG process, even for single-barrier options, can also be improved by spectral filters.

Chapter 4

Continuously monitored barrier options

As described in Section 2.1.8.1, Green et al. (2010) showed that the numerical method for calculating discretely monitored Spitzer identities can be extended to the continuous case by substituting the Laplace transform for the z-transform. In this chapter we use this extension to price continuous barrier options. Whilst continuously monitored barrier options may be uncommon in finance, they are useful as a motivating example as it allows us to compare results with the method for discrete monitoring described in Chapter 3. Moreover, the work in this chapter serves as background to the pricing of American options in Chapter 5 which are examples of options which operate in continuous time.

We first describe the new pricing procedure and detail the inverse Laplace transform by Abate and Whitt (1995) which replaces the inverse z-transform used for discrete monitoring. We then perform an in-depth error examination which shows how the error performance is affected by each step in the pricing procedure. Lastly we present results which closely match the error bound and describe how the performance of the method with continuous monitoring represents a limit of the performance for discrete monitoring with a very large number of dates.

4.1 Pricing procedure with continuous monitoring

We adapt the pricing procedures for discretely monitored barrier options to continuous monitoring by replacing $\Phi(\xi, q) = 1 - q\Phi(\xi, \Delta t)$ in the Fourier-z domain with $\Phi_{\mathrm{c}} = s - \psi(\xi)$. Step-by-step methods are presented for both single and double-barrier options.

4.1.1 Pricing procedure: single-barrier options

1. Compute the characteristic exponent $\psi(\xi + i\alpha_{\mathrm{d}})$, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

$$\Phi_{\mathrm{c}}(\xi, s) := s - \psi(\xi + i\alpha_{\mathrm{d}}) = \Phi_{\mathrm{c}\oplus}(\xi, s)\Phi_{\mathrm{c}\ominus}(\xi, s) \tag{4.1}$$

for all $s = \frac{A+2k\pi i}{2t}$ required for the inverse Laplace transform, these are specified in Eq. (4.9) below.

3. Decompose with respect to l

$$P_{\mathrm{c}}(\xi, s) := \frac{\sigma(\xi/\xi_{\max})}{\Phi_{\mathrm{c}\ominus}(\xi, s)} = P_{cl+}(\xi, s) + P_{cl-}(\xi, s), \tag{4.2}$$

and calculate

$$\widetilde{\widehat{p}}(\xi, s) := \frac{P_{cl+}(\xi, s)}{\Phi_{c+}(\xi, s)}, \tag{4.3}$$

where $\sigma(\xi/\xi_{\max})$ is an exponential filter of order p (see Section 4.3.2).

4. Calculate the option price as

$$v(0, 0) := \mathcal{F}^{-1}_{\xi\to x}\left[\widehat{\phi}^*(\xi)\mathcal{L}^{-1}_{s\to T}\widetilde{\widehat{p}}(\xi, s)\right](0), \tag{4.4}$$

where $\widehat{\phi}^*(\xi)$ is the complex conjugate of the Fourier transform of the damped payoff function which is given in Eq. (2.92).

4.1.1.1 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 4.1.1, in that it uses Wiener-Hopf factorisation and decomposition to compute the appropriate Spitzer identity. However, similarly to the pricing method for discretely monitored options in Chapter 3, 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 are the same as those for single-barrier down-and-out options described in Section 4.1.1 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

$$P_{\mathrm{c}}(\xi, s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right)\frac{1 - \Phi_{\mathrm{c}\oplus}(\xi, s)J_{cu+}(\xi, s)}{\Phi_{\mathrm{c}\ominus}(\xi, s)} = P_{cl+}(\xi, s) + P_{cl-}(\xi, s), \tag{4.5}$$

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

(c) Decompose with respect to u

$$Q_{\mathrm{c}}(\xi,s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right)\frac{1-\Phi_{\mathrm{c}\ominus}(\xi,s)J_{cl-}(\xi,s)}{\Phi_{\mathrm{c}\oplus}(\xi,s)} = Q_{cu+}(\xi,s)+Q_{cu-}(\xi,s), \tag{4.6}$$

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

(d) Calculate

$$\widetilde{\widehat{p}}(\xi,s) := \sigma\left(\frac{\xi}{\xi_{\max}}\right)\frac{1-\Phi_{\mathrm{c}\ominus}(\xi,s)J_{cl-}(\xi,s)-\Phi_{\mathrm{c}\oplus}(\xi,s)J_{cu+}(\xi,s)}{\Phi_{\mathrm{c}}(\xi,s)}. \tag{4.7}$$

(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.

4.2 Inverse Laplace transform

The Spitzer identities with continuous monitoring provide the Laplace transform of the required characteristic function, so to calculate the option price using Eq. (4.4), 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. (2.16) with

$$f(t) \approx \frac{e^{A/2}}{2t}\mathrm{Re}\widetilde{f}\left(\frac{A}{2t}\right)+\frac{e^{A/2}}{t}\sum_{k=1}^{\infty}(-1)^k\mathrm{Re}\widetilde{f}\left(\frac{A+2k\pi i}{2t}\right), \tag{4.8}$$

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

$$b_k = \frac{e^{A/2}}{2t}\mathrm{Re}\widetilde{f}\left(\frac{A}{2t}\right)+\frac{e^{A/2}}{t}\sum_{j=1}^{k}(-1)^j\mathrm{Re}\widetilde{f}\left(\frac{A+2j\pi i}{2t}\right) \tag{4.9}$$

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

$$f(t) \approx \frac{1}{2^{m_{\mathrm{E}}}}\sum_{k=0}^{m_{\mathrm{E}}}\binom{m_{\mathrm{E}}}{k}b_{n_{\mathrm{E}}+k}. \tag{4.10}$$

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 in-

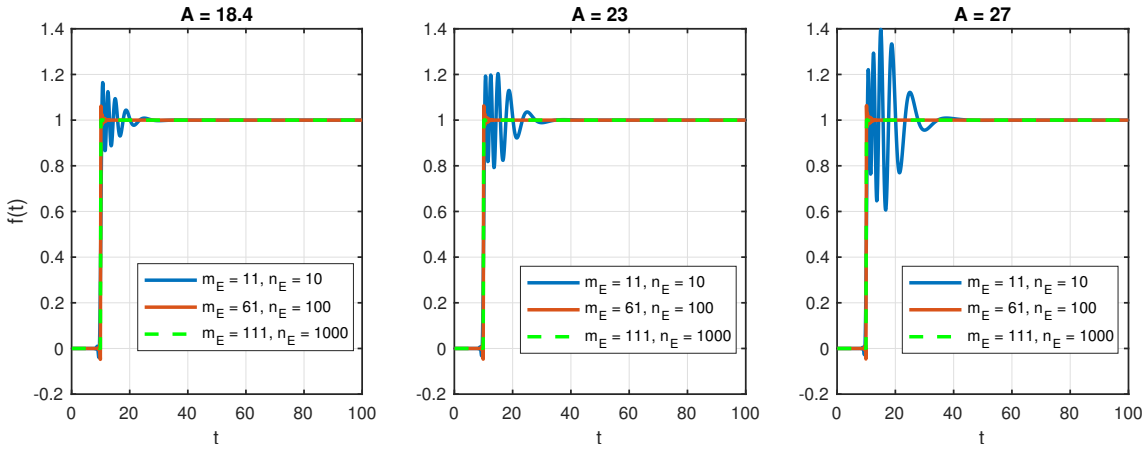

Figure 4.1: Output of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing n_{E} and m_{E} reduces the size of the oscillations but they are unaffected by increasing A.

verting the Laplace transform of a delayed unit step step function $\widetilde{f}(s) = e^{-s\tau}/s$ where the delay τ is set to 10. This is an extreme test case as it has a jump discontinuity and Abate and Whitt (1992a) stated 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 4.1 and the errors are shown in Figures 4.2 and 4.3. The empirical results in Figure 4.3 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 4.3 and 4.4 that the error bounds and observed results for the pricing procedure are limited by the performance of the sinc-based Hilbert transform and so are not limited by the accuracy of the inverse Laplace 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 empirical results. From Figures 4.1 and 4.2 we can see that the size of the oscillations due to the discontinuity are predominantly affected by m_{E} and n_{E}. The error floor is controlled by the selection of A; the values of 18.4, 23 and 27 in Figures 4.1–4.3 correspond to errors of approximately 10^{-8}, 10^{-10} and 10^{-12} respectively. However, Figure 4.3 shows that the variation of 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.

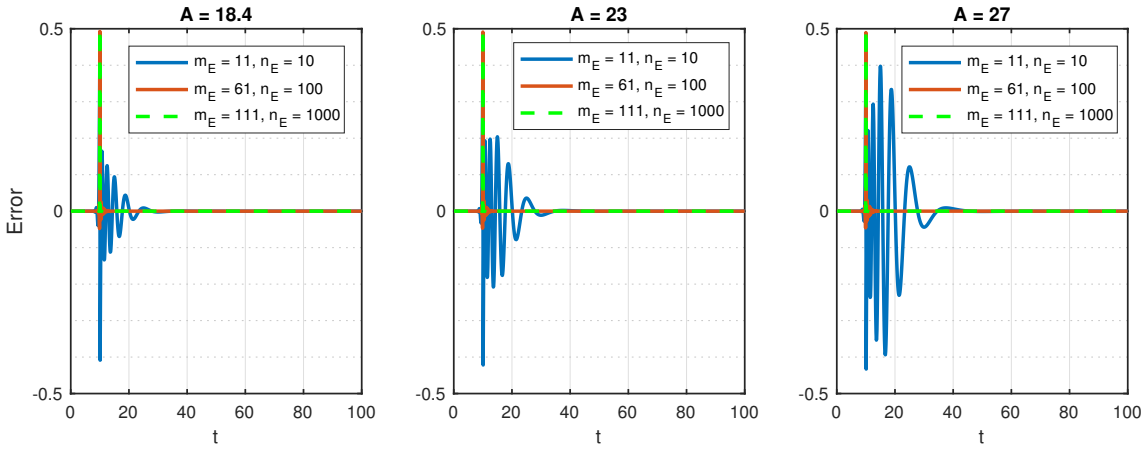

Figure 4.2: Error of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing n_{E} and m_{E} reduces the size of the errors due to the oscillations but they are unaffected by increasing A.

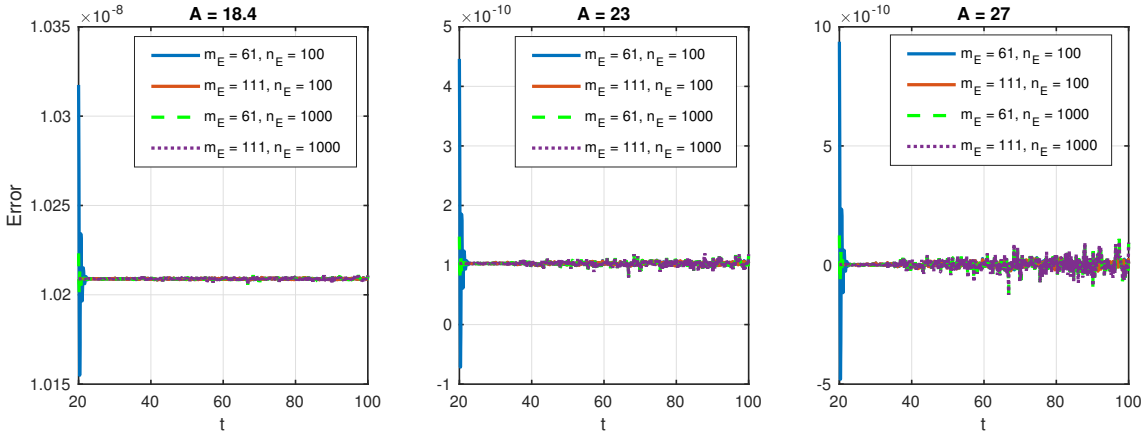

Figure 4.3: Error of the inverse Laplace transform of $\widetilde{f}(s) = \frac{e^{-10s}}{s}$. Increasing the value of A decreases the error floor but it is unaffected by increasing n_{E} and m_{E}. The noise on the error floor is $\approx 10^{-10}$.

4.3 Error convergence of the pricing procedure

We examine the performance of each stage of the pricing procedure and discuss their 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 of the sinc-based fast Hilbert transform in Eq. (2.100) is exponentially convergent when the function $f(\xi)$ is analytic in a strip of the complex plane including the real axis. Therefore, similarly to the error calculations in Section 3.1, our concern is with the truncation error from the approximation in Eq. (2.101). As before, the truncation error using the sinc-based Hilbert transform depends on the behaviour of the characteristic exponent as $|\xi| \to \infty$; Table 4.1 shows the characteristic exponents of five Lévy processes. The damping factor α_{d}

is omitted to make the notation more concise which is appropriate as the value of $i\alpha_{\mathrm{d}}$ becomes insignificant as $|\xi| \to \infty$.

Process	Characteristic exponent $\psi(\xi)$	Rational
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 4.1: Characteristic exponent of some Lévy processes.

4.3.1 Factorisation

After calculating the characteristic exponent, the next step in the pricing procedure is the numerical factorisation of $\Phi_{\mathrm{c}}(\xi, s) = s - \psi(\xi)$. In order to understand the error convergence we must consider the way that the function behaves for large values of $|\xi|$. The characteristic exponents of the processes listed in Table 4.1 will take very high negative values which will dominate $\Phi_{\mathrm{c}}(\xi, s)$ so that as $|\xi| \to \infty$ we can approximate $s - \psi(\xi) \approx -\psi(\xi)$. The function to be decomposed in the factorisation stage is therefore $\approx \log[-\psi(\xi)]$. This function increases in magnitude 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. (2.101) for positive and negative values of k individually, we obtain two infinite summations that do not converge. However, Table 4.1 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

$$\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| =& \Delta\xi \left| \sum_{k<-M/2} \log \Phi_c(k\Delta\xi, s) \frac{1 - \cos\left(\pi \frac{\xi - k\Delta\xi}{\Delta\xi}\right)}{\pi(\xi - k\Delta\xi)} \right. \\
&\left. + \sum_{k>M/2} \log \Phi_c(k\Delta\xi, s) \frac{1 - \cos\left(\pi \frac{\xi - k\Delta\xi}{\Delta\xi}\right)}{\pi(\xi - k\Delta\xi)} \right| \\
<& \Delta\xi \sum_{k>M/2} \left| \frac{\log \Phi_c(k\Delta\xi, s)}{\pi(\xi - k\Delta\xi)} \left(1 - \cos\left(\pi \frac{\xi - k\Delta\xi}{\Delta\xi}\right)\right) \right. \\
&\left. + \frac{\log \Phi_c(-k\Delta\xi, s)}{\pi(\xi + k\Delta\xi)} \left(1 - \cos\left(\pi \frac{\xi + k\Delta\xi}{\Delta\xi}\right)\right) \right| \\
\leq& \frac{2\Delta\xi}{\pi} \sum_{k>M/2} \left| \frac{\xi\big(\log \Phi_c(k\Delta\xi, s) + \log \Phi_c(-k\Delta\xi, s)\big)}{\xi^2 - k^2\Delta\xi^2} \right. \\
&\left. + \frac{k\Delta\xi\big(\log \Phi_c(k\Delta\xi, s) - \log \Phi_c(-k\Delta\xi, s)\big)}{\xi^2 - k^2\Delta\xi^2} \right|, \qquad (4.11)
\end{aligned}$$

where $f_{\Delta\xi}(\xi)$ is the value of the infinite summation shown in Eq. (2.100) and $f_{\Delta\xi,M}(\xi)$ is the result of the truncated summation in Eq. (2.101). The two cosines are equal because the difference of their arguments is $2k\pi$; thus $|1 - \cos(.)| \leq 2$ can be factored out.

The next step in bounding the error convergence is to show that the expression in Eq. (4.11) 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. (4.11). We then calculate a bound for the error based on the first summation term in Eq. (4.11). These steps are carried out in a slightly different way depending on the form of the characteristic exponents shown in Table 4.1.

4.3.1.1 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 4.1. The parameters μ and σ are specific to the distribution. We can therefore approximate the second expression in the summation by

$$\begin{aligned}
\frac{k\Delta\xi\big(\log \Phi_c(k\Delta\xi, s) - \log \Phi_c(-k\Delta\xi, s)\big)}{\xi^2 - k^2\Delta\xi^2} &= \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2} \log \frac{\Phi_c(k\Delta\xi, s)}{\Phi_c(-k\Delta\xi, s)} \\
&\approx \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2} \log \frac{\sigma^2(k\Delta\xi)^2/2 + i\mu(k\Delta\xi)}{\sigma^2(k\Delta\xi)^2/2 - i\mu(k\Delta\xi)} \\
&= \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2} \log \frac{1 + 2i\mu/(\sigma^2 k\Delta\xi)}{1 - 2i\mu/(\sigma^2 k\Delta\xi)} \qquad (4.12)
\end{aligned}$$

The logarithm in Eq. (4.12) 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

$$\frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\log\frac{1+2i\mu/(\sigma^2 k\Delta\xi)}{1-2i\mu/(\sigma^2 k\Delta\xi)} \approx \frac{k\Delta\xi}{\xi^2 - k^2\Delta\xi^2}\frac{i4\mu}{\sigma^2 k\Delta\xi} = \frac{i4\mu}{\sigma^2(\xi^2 - k^2\Delta\xi^2)} \tag{4.13}$$

gives an approximation for the second term in Eq. (4.11). Due to the denominator, this is $O(1/k^2)$. Thus, as $\log\Phi_{\mathrm{c}}(k\Delta\xi, s)+\log\Phi_{\mathrm{c}}(-k\Delta\xi, s)$ is increasing in k, the error is indeed dominated by the first term in Eq. (4.11).

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

$$\begin{aligned}|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< \frac{c_1\Delta\xi}{\pi}\sum_{k>M/2}\left|\frac{\log\Phi_{\mathrm{c}}(k\Delta\xi, s)}{\xi^2 - k^2\Delta\xi^2}\right| \\ &< c_2\Delta\xi\sum_{k>M/2}\frac{\log(k^2\Delta\xi^2)}{k^2\Delta\xi^2},\end{aligned} \tag{4.14}$$

Here, as Eq. (4.14) 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 $\Phi_{\mathrm{c}\pm}(\xi, s)$ which decays as $|\xi| \to \infty$ and the use of a filter on the input to the next step as described in Section 4.3.2. These effects combine to mean that the error as a proportion of the value of $\Phi_{\mathrm{c}\pm}(\xi, s)$ 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 $\Phi_{\mathrm{c}\pm}(\xi, s)$ for low values of $|\xi|$. Approximating the summation in Eq. (4.14) by an integral with $M' = M/2$, we obtain

$$\begin{aligned}|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &< c_3\int_{M\Delta\xi/2}^{+\infty}\frac{\log\xi_k}{\xi_k^2}d\xi_k \\ &= c_3\left[\frac{\log\xi_k}{\xi_k} + \frac{1}{\xi_k}\right]_{+\infty}^{M\Delta\xi/2} \\ &= c_3\left(\frac{\log(M\Delta\xi/2)}{M\Delta\xi/2} + \frac{1}{M\Delta\xi/2}\right).\end{aligned} \tag{4.15}$$

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. (4.15), we can bound the truncation error of the Wiener-Hopf factors, and by extension the total error as the truncation error dominates, as

$$\left|\frac{\Phi_{\Delta\xi,\mathrm{c}\pm}(\xi) - \Phi_{\Delta\xi,M',\mathrm{c}\pm}(\xi)}{\Phi_{\Delta\xi,\mathrm{c}\pm}(\xi)}\right| < \left|1 - (eM'\Delta\xi)^{\frac{\kappa}{M'\Delta\xi}}\right|, \tag{4.16}$$

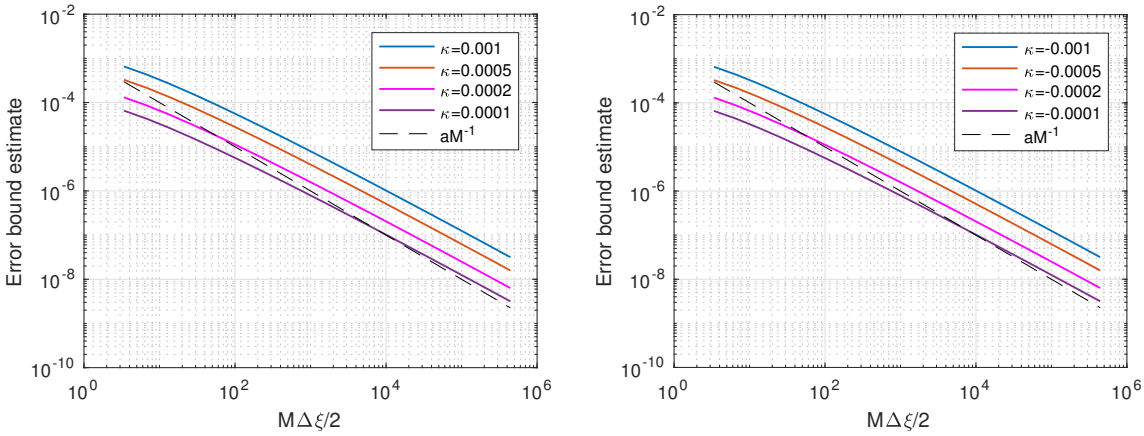

Figure 4.4: Right-hand side of Eq. (4.16) plotted for different values of κ (left: positive values; right: negative values) to show the estimate of the error bound on the sinc-based numerical factorisation of $\Phi_{\mathrm{c}}(\xi, s)$. Notice that the predicted error bound from the factorisation has 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 4.3.1.1, 4.3.1.2 and 4.3.1.3 show that this bound applies for the normal, NIG, Kou, Merton and VG processes.

where κ is some constant. Here, $\Phi_{\Delta\xi,\mathrm{c}\pm}(\xi)$ denotes the (theoretical) Wiener-Hopf factors generated using the infinite summation in Eq. (2.100) and $\Phi_{\Delta\xi,M',\mathrm{c}\pm}(\xi)$ denotes the Wiener-Hopf factors calculated using the truncated summation in Eq. (2.101).

Figure 4.4 shows Eq. (4.16) 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.

4.3.1.2 NIG process

In the case of the NIG process the characteristic exponent is

$$\psi(\xi) = \delta(\sqrt{\alpha^2 - (\beta + i\xi)^2} - \sqrt{\alpha^2 - \beta^2}). \tag{4.17}$$

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. (4.12) is $\frac{1}{2}\log\frac{1+i2\beta/(k\Delta\xi)}{1-i2\beta/(k\Delta\xi)}$. Furthermore, $\Phi_{\mathrm{c}}(k\Delta\xi, s)$ and $\Phi_{\mathrm{c}}(-k\Delta\xi, s)$ 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. (4.15) is the size of the constants.

4.3.1.3 VG process

The characteristic exponent of the VG process is

$$\psi(\xi) = -\frac{1}{\nu}\log\left(1 - i\xi\theta\nu + 0.5\nu\sigma^2\xi^2\right). \tag{4.18}$$

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

$$\begin{aligned}
\log\frac{\Phi_c(k\Delta\xi,s)}{\Phi_c(-k\Delta\xi,s)} &\approx \log\frac{\log(-ik\Delta\xi\theta\nu + \nu\sigma^2(k\Delta\xi)^2/2)}{\log\left(ik\Delta\xi\theta\nu + \nu\sigma^2(k\Delta\xi)^2/2\right)} \\
&= \log\frac{\log\left(1 - \frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2}\right) + \log(\nu\sigma^2(k\Delta\xi)^2/2)}{\log\left(1 + \frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2}\right) + \log\left(\nu\sigma^2(k\Delta\xi)^2/2\right)} \\
&\approx \log\frac{-\frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2} + \log(\nu\sigma^2(k\Delta\xi)^2/2)}{\frac{2ik\Delta\xi\theta\nu}{\nu\sigma^2(k\Delta\xi)^2} + \log\left(\nu\sigma^2(k\Delta\xi)^2/2\right)} \\
&= \log\frac{1 - \frac{2i\theta}{\sigma^2 k\Delta\xi\log(\nu\sigma^2(k\Delta\xi)^2/2)}}{1 + \frac{2i\theta}{\sigma^2 k\Delta\xi\log(\nu\sigma^2(k\Delta\xi)^2/2)}} \\
&\approx \frac{-4i\theta}{\sigma^2 k\Delta\xi\log\left(\nu\sigma^2(k\Delta\xi)^2/2\right)}.
\end{aligned} \tag{4.19}$$

This decreases quicker than $O(1/k)$ and thus Eq. (4.11) is dominated by the first term. The equivalent expression to Eq. (4.14) for the VG process is

$$\begin{aligned}
|f_{\Delta\xi} - f_{\Delta\xi,M}| &< c\sum_{k>M/2}\frac{\Delta\xi}{\pi}\frac{\log\Phi_c(k\Delta\xi,s)}{\xi^2 - k^2\Delta\xi^2}, \\
&< c_1\Delta\xi\sum_{k>M/2}\frac{\log\log(k^2\Delta\xi^2)}{k^2\Delta\xi^2}.
\end{aligned} \tag{4.20}$$

As $|\log\log(x)|$ is bounded by $|\log x|$ as $x\to\infty$, the factorisation error with the VG process is also bounded by Eq. (4.16).

4.3.2 Decomposition error

The output of the factorisation is shown in Figure 4.5. The next step in the calculation is to find the positive part of $P_c(\xi,s) = \frac{e^{-il\xi}}{\Phi_{c\ominus}(\xi,s)}$. We can attempt to bound the truncation error of this calculation by combining the errors from the positive and negative truncations

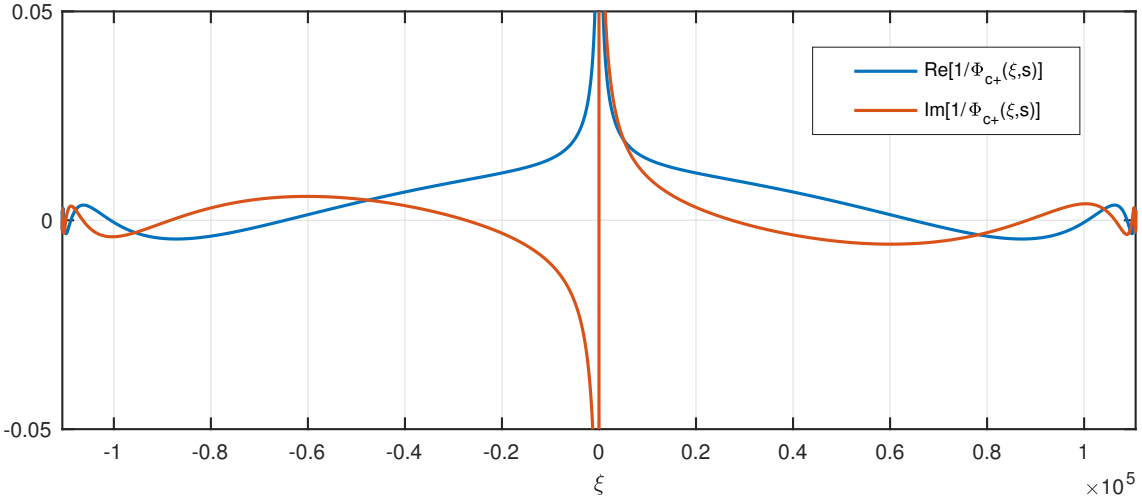

Figure 4.5: Example plot of the real and imaginary parts of $\Phi_{c\oplus}(\xi, s)$ plotted against ξ with s chosen as $A/2t$, as specified for the Abate and Whitt inverse Laplace transform. Notice that although the value of $|\Phi_{c\oplus}(\xi, s)|$ is bounded by a constant as $|\xi| \to \infty$, the rate of decay is very slow and we have not been able to determine a decreasing bound.

as before:

$$\begin{aligned}
|f_{\Delta\xi}(\xi) - f_{\Delta\xi,M}(\xi)| &= \frac{\Delta\xi}{\pi}\left|\sum_{k>M/2}\frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi} + \sum_{k<-M/2}\frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi}\right| \\
&= \frac{\Delta\xi}{\pi}\left|\sum_{k>M/2}\frac{P_c(k\Delta\xi)}{\xi - k\Delta\xi} + \frac{P_c(-k\Delta\xi)}{\xi + k\Delta\xi}\right| \\
&= \frac{\Delta\xi}{\pi}\left|\sum_{k>M/2}\frac{\xi[P_c(k\Delta\xi) + P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2} + \frac{k\Delta\xi[P_c(k\Delta\xi) - P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2}\right|.
\end{aligned} \tag{4.21}$$

Figure 4.5 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. (4.21) as

$$\sum_{k>M/2}\frac{k\Delta\xi[P_c(k\Delta\xi) - P_c(-k\Delta\xi)]}{\xi^2 - (k\Delta\xi)^2} < c\sum_{k>M/2}\frac{k\Delta\xi}{\xi^2 - (k\Delta\xi)^2} \tag{4.22}$$

where c is some positive constant, this does not converge. We can compare this with the discretely monitored version from Fusai et al. (2016), which we modified in Chapter 3, where the first and last dates are taken out of the scheme. They are then applied by convolution which means 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 of ξ which again means that the truncation error of

the sinc-based discrete Hilbert transform is exponentially bounded. In Section 3.2.5 we described how we can remove the first and last dates from the discrete pricing pricing procedure by replacing the characteristic function with a spectral filter. For continuous monitoring, we use the exponential filter which we used with discretely monitored options in Chapter 3 and has previously achieved good results in other option pricing applications (Ruijter et al., 2015; Cui et al., 2017). The filter is described by Eq. (3.8) and its shape is shown in Figure 3.2. 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. (4.16) and shown in Figure 4.4.

4.4 Results

We present results for the Spitzer based 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 step between them.

4.4.1 Results for Spitzer-Laplace method for continuously monitored options

The error convergence for single-barrier down-and-out options is shown in Figure 4.6, the results with the NIG and VG processes conform very closely to the calculated error bound. The results with the Kou process deviate from this slightly for lower values of M but show the typical sub-polynomial error convergence for higher values of M. Figure 4.7 shows the results for double-barrier options, the absolute error is worse than that for the single-barrier option but the speed of the error convergence conforms very closely to the calculated error bound. The computed prices are given in Tables 4.2 and 4.3 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 4.4. The details of the contract and underlying processes are shown in Table A.2 in Appendix A.

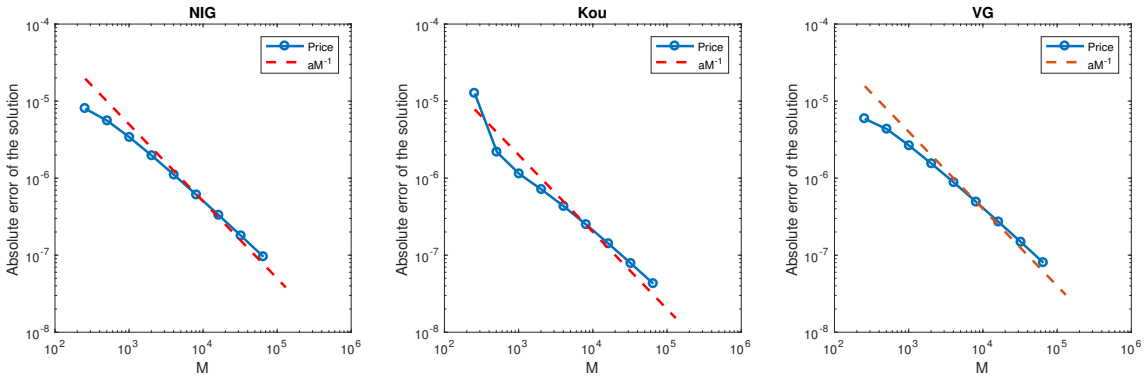

Figure 4.6: Error convergence for a continuously monitored single-barrier option. The results with the NIG and VG processes conform very closely to the calculated error bound. The results with the Kou process deviate from this slightly for lower values of M but show the typical sub-polynomial error convergence for higher values of M.

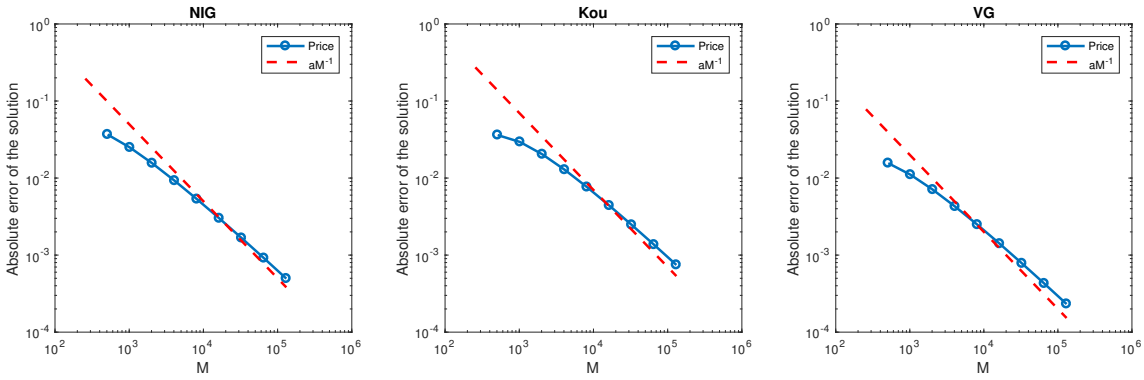

Figure 4.7: Error convergence for a continuously monitored double-barrier option. The absolute error is worse than that for a single-barrier option but the speed of the error convergence conforms very closely to the calculated error bound.

4.4.2 Comparison with the error convergence of Spitzer based pricing methods for discretely monitored options

In Section 2.1.8.1 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 single-barrier options in Fusai et al. (2016) and double-barrier options in Chapter 3, with the number of monitoring dates up to $\approx 10^3$, 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) and Broadie et al. (1997) the error between the discretely and continuously monitored prices was shown to be bounded as $O(1/\sqrt{N})$, where N is the

Process	Price
NIG	4.77403523401E-2
Kou	4.32042632202E-2
VG	4.70627023105E-2

Table 4.2: Prices calculated for single-barrier options with the contract details and process parameters described in Table A.2 in Appendix A and $M = 2^{17}$.

Process	Price
NIG	2.78787488E-2
Kou	3.30368034E-2
VG	2.82666693E-2

Table 4.3: Prices calculated for double-barrier options with the contract details and process parameters described in Table A.2 in Appendix A and $M = 2^{17}$.

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 Chapter 3 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 4.9 and 4.10. 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 $\Delta t \to 0$. In simple terms 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 can be intuitively considered as a limit of the error convergence for the discrete case as $\Delta t \to 0$. The relationship between the error convergence of the methods with discrete and continuous monitoring is examined in more detail in Section 4.4.2.2 below.

Computed prices for continuously and discretely monitored options are plotted against M in Figures 4.11 and 4.12. In addition, computation times for the pricing methods for the discrete and continuously monitored methods are shown in Tables 4.4 and 4.5. Figures 4.11 and 4.12 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 4.4 and 4.5 we take the lowest time where the convergence error for the discretely monitored method is significantly (about ten times) lower 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 single-barrier and $\approx 10^{-2}$ for double-barrier, the discretely monitored method is slightly quicker. However, the discretely monitored method is unable to achieve the lower errors, $\approx 10^{-6}$ for single-barrier and $\approx 10^{-4}$ for double-barrier, which are attained by the continuously monitored method and therefore is not a sufficiently accurate approximation.

However, it should be considered whether 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 the convergence of the discrete method to the continuous method with increasing monitoring dates to be $O(1/\sqrt{N})$. From Figure 4.8, which displays the difference between the prices for discrete and continuous monitoring at the maximum grid size of 2^{17}, 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 4.4 and 4.5 by 100 times for the NIG and VG processes and by 200^2 times for the Kou process. Figures 4.9 and 4.10 show the error for each N (including continuous monitoring) calculated against the price for the same number of dates with 2^{18} grid points. These results demonstrate 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. Moreover, Figures 4.9 and 4.10 also illustrate the fact that as $\Delta t \to 0$ the slope of the error convergence of the discretely monitored scheme approaches that of the continuously monitored scheme. Thus the only possibility to improve the error of the discrete method 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.

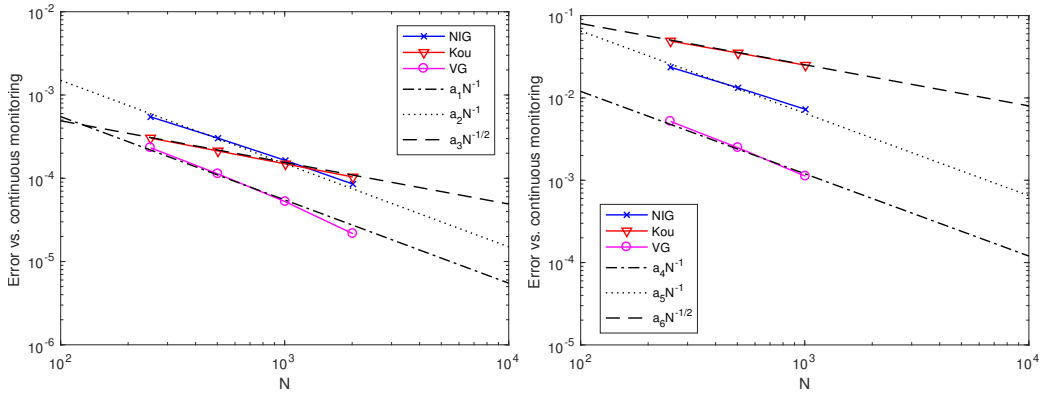

Figure 4.8: Error for discretely monitored barrier options used as an approximation of the price for the continuously monitored case, plotted as a function of the number of monitoring dates. The error is calculated as the difference between the prices for discrete and continuous monitoring at the maximum 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 $O(1/N)$.

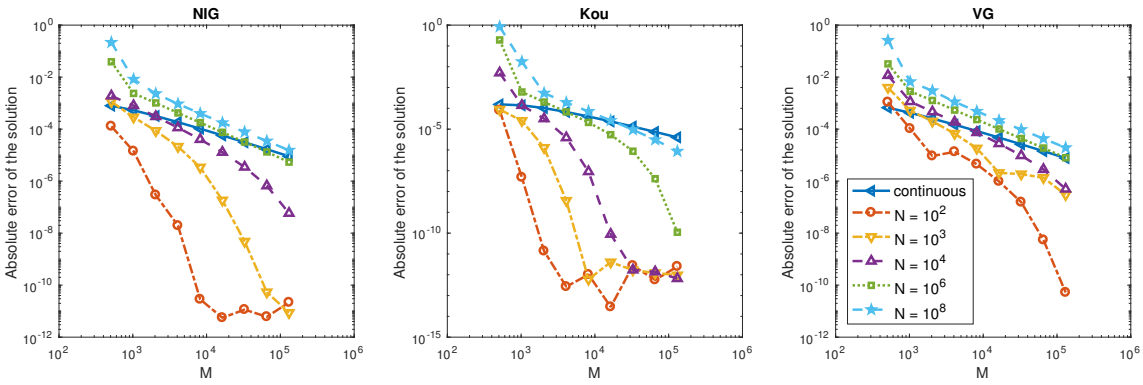

Figure 4.9: Error as a function of the grid size M for continuously monitored single-barrier options compared to discretely monitored options as the number of monitoring dates N increases. The error for each number of dates (including continuous monitoring) is calculated against the price for the same 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.

Process	Dates	Error	Time	M
NIG	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
VG	252	2.31E-04	0.04	2048
	cont	2.57E-04	0.06	1024
	504	1.13E-04	0.11	4096
	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
Kou	252	3.02E-04	0.01	512
	cont	1.57E-04	0.02	256
	504	2.03E-04	0.01	512
	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 4.4: 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.

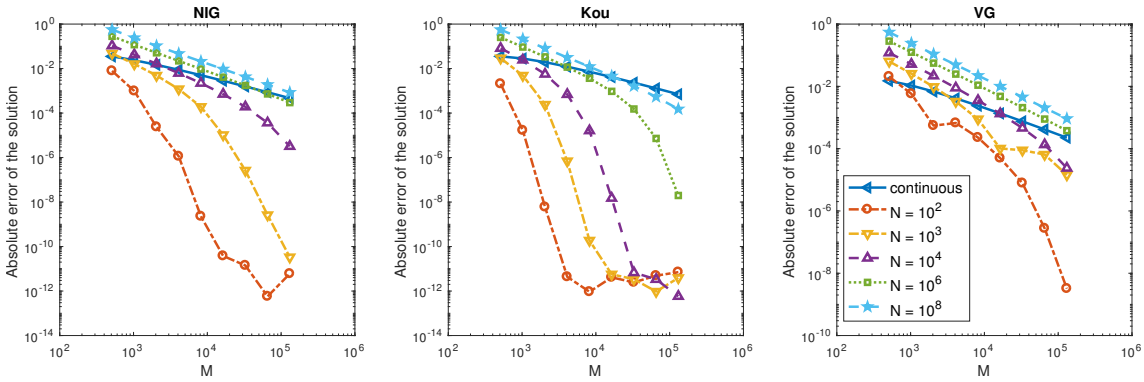

Figure 4.10: Error as a function of the grid size M for continuously monitored double-barrier options compared to discretely monitored options as the number of monitoring dates N increases. The error for each number of dates (including continuous monitoring) is calculated against the price for the same 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.

Process	Dates	Error	Time	M
NIG	252	2.35E-02	0.16	4096
	cont	2.38E-02	0.11	512
	504	1.32E-02	0.14	4096
	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
VG	252	5.10E-03	0.08	2048
	cont	6.86E-03	0.24	1024
	504	2.51E-03	0.13	4096
	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
Kou	252	4.90E-02	0.03	1024
	cont	3.54E-02	0.07	256
	504	3.51E-02	0.07	2048
	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 4.5: 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.

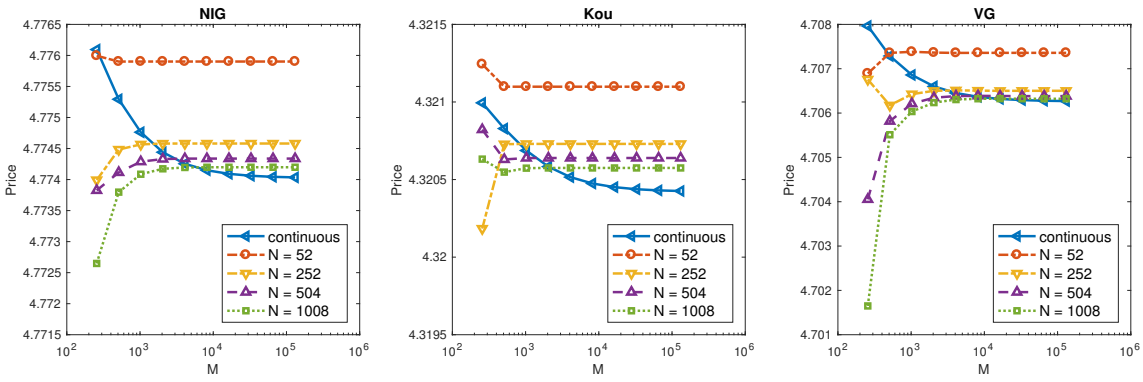

Figure 4.11: Price plotted against the grid size M for continuously monitored single-barrier options compared to discretely monitored options as the number of monitoring dates N increases. Note that the larger the value of N, the closer the price is to the continuously monitored option price.

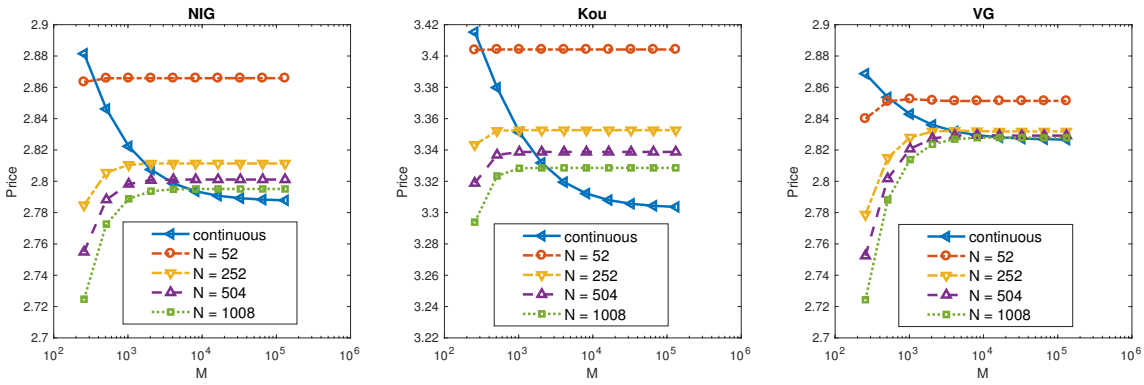

Figure 4.12: Price plotted against grid size M for continuously monitored double-barrier options compared to discretely monitored options as the number of monitoring dates N increases. Note that the larger the value of N, the closer the price is to the continuously monitored option price.

4.4.2.1 Effect of a spectral filter on the error convergence

The numerical method described in the main paper has three main stages of calculation: an initial factorisation, a fixed-point algorithm to solve a pair of coupled integral equations with decompositions, and the final stage where the product of the Fourier transforms of the probability density and payoff functions is transformed back into the log-price domain.

As described in Section 4.3.2 we use require filtering on the decomposition step and we apply a filter before the final inverse Fourier transform. Using the variance gamma (VG) process as an example, Figure 4.13 shows results with no filtering, filtering only one of the two stages, and filtering on both stages. The results show that a filter improves the convergence of the fixed-point algorithm and of the final step, so that the error convergence becomes dominated by the initial factorisation and follows the bound for the latter as described in Section 3.1 of the main paper. Specifically, a filter improves convergence for low values of the grid size M, but this effect fades away for larger numbers of grid points; moreover a filter changes the direction of convergence, so that the convergence of the price and of the error on M is now monotonic rather than oscillating.

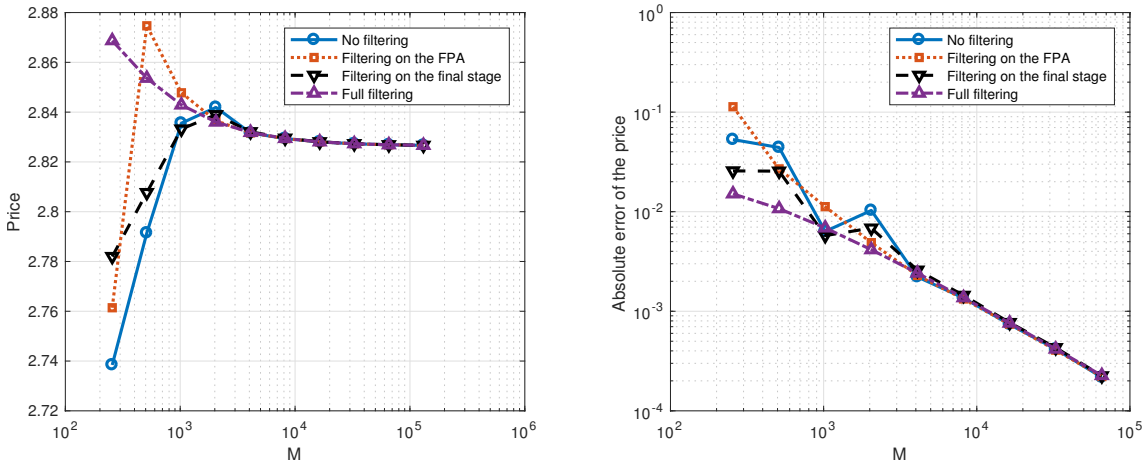

Figure 4.13: Effect of a spectral filter for the variance gamma process. The results show that a filter improves the convergence of the fixed-point algorithm and of the final step, so that the error convergence becomes dominated by the initial factorisation.

4.4.2.2 Detailed comparison of the error convergence with continuous and discrete monitoring

Substituting $q = e^{-s\Delta t}$ in $\Phi(\xi, q) = 1 - q\Psi(\xi, \Delta t)$ to exploit the conformal mapping between the z domain and the Laplace domain, the function factorised in the discretely monitored scheme becomes

$$\Phi_{\mathrm{d}}(\xi, s) = 1 - e^{-s\Delta t}\Psi(\xi, \Delta t) = 1 - e^{(\psi(\xi)-s)\Delta t}, \tag{4.23}$$

where $\psi(\xi) = \frac{1}{\Delta t}\log\Psi(\xi, \Delta t)$ is the characteristic exponent of the Lévy process.

To factorise $\Phi_{\mathrm{d}}(\xi, s)$, we decompose

$$\log\Phi_{\mathrm{d}}(\xi, s) = \log\left\{1 - e^{[\psi(\xi)-s]\Delta t}\right\} \tag{4.24}$$

and exponentiate the result. Figure 4.14 shows how, As Δt becomes smaller, the shape of $|\log\Phi_{\mathrm{d}}(\xi, s)|$ changes in two significant ways. Firstly the value of $|\log\Phi_{\mathrm{d}}(0, s)|$ increases and secondly the gradient of $|\log\Phi_{\mathrm{d}}(\xi, s)|$ as $|\xi| \to \infty$ becomes shallower. At the limit $\Delta t = 0$,

$$\forall\,\xi \quad |\log\Phi_{\mathrm{d}}(\xi, s)| = \infty, \tag{4.25}$$

so the function cannot be decomposed.

Although the factorisation for $\Delta t = 0$ cannot be performed using $\Phi_{\mathrm{d}}(\xi, s)$ directly, when we apply the Laplace transform we multiply $1/\Phi_{\mathrm{d}}(\xi, s)$ by Δt, which recovers a well defined function $\Phi_{\mathrm{c}}(\xi, s) = s - \psi(\xi)$. However, this affects the shape of the function under decomposition, i.e. $\log\Phi_{\mathrm{c}}(\xi, s)$, and we can understand this further comparing the

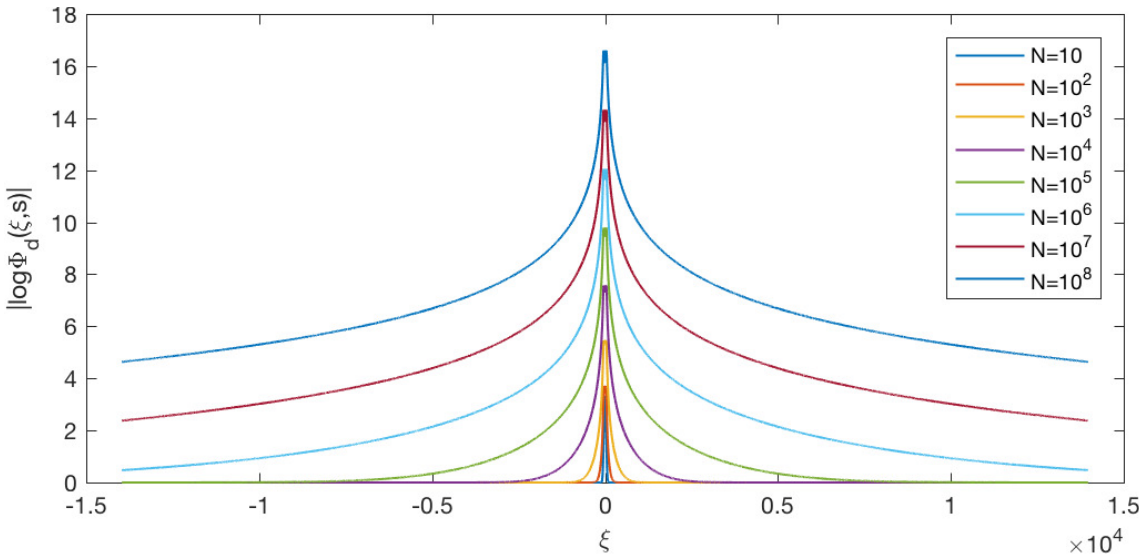

Figure 4.14: Plot of $|\log \Phi_{\mathrm{d}}(\xi, s)|$ versus ξ for maturity $T = 1$ and different numbers of monitoring dates N. As $\Delta t = T/N$ decreases, the value of the function at $\xi = 0$ increases and the slope of the function becomes shallower.

functions $\log \Phi_{\mathrm{d}}(\xi, s)$ and $-\log \frac{\Delta t}{\Phi_{\mathrm{d}}(\xi,s)}$ with discrete monitoring. As $|\xi| \to \infty$, $\Phi_{\mathrm{d}}(\xi, s) \to 1$, so clearly $\log \Phi_{\mathrm{d}}(\xi, s) \to 0$; however as $|\xi| \to \infty$, $\frac{\Delta t}{\Phi_{\mathrm{d}}(\xi,s)} \to \Delta t$, so rather than decreasing to 0, $-\log \frac{\Delta t}{\Phi_{\mathrm{d}}(\xi,s)}$ instead increases to $-\log \Delta t$. Clearly $\log \Phi_{\mathrm{c}}(\xi, s)$ represents the limit of this for $\Delta t \to 0$ and thus increases to $-\log 0 = +\infty$ as $|\xi| \to \infty$.

As described by Stenger (1993) and Feng and Linetsky (2008), the shape of the function to be decomposed determines the truncation error with the sinc-based discrete Hilbert transform. Section 3.1, showed that $|\log \Phi_{\mathrm{d}}(\xi, s)|$ decreases with ξ and that the speed of error convergence is determined by the rate of decrease. Moreover, Figure 4.14 shows that the slope of $|\log \Phi_{\mathrm{d}}(\xi, s)|$ becomes shallower with increasing number of dates N, so we would expect the convergence to become slower with decreasing Δt, which is indeed observed in Figures 4.9–4.12. In contrast $\left|\log \frac{\Delta t}{\Phi_{\mathrm{d}}(\xi,s)}\right|$ and $|\log(s - \psi(\xi))|$ are both increasing functions of $|\xi|$ and so we must use the error calculations in Section 4.3, which again match the error results above.

4.5 Conclusions

We showed that the numerical method for calculating the discretely monitored Spitzer identities described by Fusai et al. (2016), and improved with spectral filtering in Chapter 3, 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 can achieve 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 the discretely monitored case achieves exponential convergence as seen for the single-barrier case in Fusai et al. (2016) and for the double-barrier case in Chapter 3, but the continuous case described here does not. However, we demonstrate 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.

Chapter 5

Methods for other exotic options

Following the pricing of barrier options in Chapters 3 and 4, we describe here the use of the Spitzer identity formulation by Green (2009) in order to price other exotic options. We look at quantile options with discrete and continuous monitoring and perpetual early exercise options, also with discrete and continuous monitoring, i.e. options of Bermudan and American type.

The pricing of quantile options is based on the Dassios-Port-Wendel identity in Eq. (2.88) and uses Green's formulation for hindsight options which were implemented for general Lévy processes by Fusai et al. (2016) in order to price fixed-strike Lookback options. We use two methods for pricing perpetual Bermudan and American options, one by Green (2009) which is based on residue theorem and another which is based on the fluctuation identities for pricing barrier options in Chapters 3 and 4 and includes a novel method of calculating the optimal exercise barrier.

5.1 Quantile options

The α-quantile of a random process $X(t)$ (X_α) is the value which it is below $\alpha\%$ of the time; α-quantile options have a payoff which is calculated as a function of this value. They were first designed by Miura (1992) as a variation on lookback options which would be less susceptible to the effects of very large, but short lived, swings in the price of the underlying asset. They are considered to be more resistant to market manipulation as it is easier to cause a brief peak or trough in an asset price compared to a longer term price movement. We propose a pricing method for α-quantile options with general exponential Lévy processes.

5.1.1 Pricing discretely monitored quantile options

For α-quantile options the form of the payoff is the same as in Eq. (2.91), i.e.

$$\phi(x) = e^{\alpha_{\mathrm{d}} x} S_0(\theta(e^x - e^{k_S}))^+ \mathbf{1}_{[l,u]}(x), \tag{5.1}$$

but in this case it is calculated as a function of $x = X_\alpha$ rather than the value at expiry, $x = X(T)$. Therefore, if we have the characteristic function of X_α then we can price an α-quantile option in the Fourier domain using the Plancherel relation in Eq. (2.93) with the Fourier transform of the damped payoff function $\widehat{\phi}(\xi)$ from Eq. (2.92).

The Dassios-Port-Wendel identity states that the α-quantile of a Brownian motion over T has the same distribution as the sum of the infimum of a Brownian motion over time $(1-\alpha)T$ and the supremum of an independent Brownian motion over time αT. That is if $X_m = \min_{t\in[0,(1-\alpha)T]} X_1(t)$ and $X_M = \max_{t\in[0,\alpha T]} X_2(t)$, where $X_1(t)$ and $X_2(t)$ are independent Brownian motions, then

$$X_\alpha \overset{d}{=} X_m + X_M, \tag{5.2}$$

where X_α is the α-quantile of the Brownian motion.

We can understand the link between the value of X_α and the supremum and infimum quite intuitively. Firstly, if a process is split into a section for $t \in [0, \alpha T]$ and one for $t \in (\alpha T, T]$ then, by the property of independent increments, they represent two independent processes over $t \in [0, \alpha T]$ and $t \in (0, (1-\alpha)T]$. Moreover, the supremum of $X_1(t)$ is the value that the process has spent αT time below and conversely the infimum of $X_2(t)$ is the value which this process has spent $(1-\alpha)T$ time above. Although the basic idea behind the relationship in Eq. (5.2) is quite clear, the mathematical proof is quite involved and we refer the interested reader to the original paper by Dassios (1995) for the details. A note by Dassios (2006) also showed that this identity could also be extended (although not uniquely) to general Lévy processes.

Green et al. (2010) devised Spitzer-based formulations for the probability distributions of the maximum and minimum of a process which were used by Fusai et al. (2016) to price fixed-strike lookback options with exponential error convergence. These are defined in the Fourier-z domain as

$$\widetilde{\widehat{p_{X_M}}}(\xi, q) = \frac{1}{\Phi_\oplus(\xi, q)\Phi_\ominus(0, q)}, \tag{5.3}$$

$$\widetilde{\widehat{p_{X_m}}}(\xi, q) = \frac{1}{\Phi_\oplus(0, q)\Phi_\ominus(\xi, q)}, \tag{5.4}$$

where $\Phi_\oplus(\xi, q)$ and $\Phi_\ominus(\xi, q)$ are the Wiener-Hopf factors of $1 - q\Psi(\xi, \Delta t)$ as described in

Section 2.1.8. The inverse z-transform can be applied to obtain

$$\widehat{p_{X_M}}(\xi, j) = \mathcal{Z}^{-1}_{q\to j}[\widetilde{\widehat{p_{X_M}}}(\xi, q)], \tag{5.5}$$

$$\widehat{p_{X_m}}(\xi, N-j) = \mathcal{Z}^{-1}_{q\to N-j}[\widetilde{\widehat{p_{X_m}}}(\xi, q)], \tag{5.6}$$

in the Fourier-domain, where N is the number of discrete monitoring dates and $j = \alpha N$ to the nearest integer,. As X_M and X_m are the supremum and infimum of mutually independent processes it is clear that X_M and X_m are independent random variables. It is a basic result in probability theory that the PDF of the sum of two independent random variables is equal to the convolution of their respective PDFs, i.e.

$$p_{X_\alpha}(x) = \int_{-\infty}^{+\infty} p_{X_M}(x')p_{X_m}(x-x')dx'. \tag{5.7}$$

Therefore, by convolution theorem, we can multiply them point-wise in the Fourier domain, as described in Section 2.1.1, i.e.

$$\widehat{p_{X_\alpha}}(\xi) = \widehat{p_{X_M}}(\xi, j)\widehat{p_{X_m}}(\xi, N-j). \tag{5.8}$$

The option price can then be calculated using the Plancherel relation in Eq. (2.93).

The calculations for discretely monitored price for loopback options as used by Fusai et al. (2016) are based on a distribution of the maximum (or minimum) of the process at $t > 0$. Similarly to the schemes for barrier options described in Section 2.3.4, the first date is taken out of the Spitzer-based scheme and the result for $N-1$ dates is multiplied by the characteristic function. This gives a smooth probability distribution, i.e. $p_{X_{M'}} \in C^\infty$, as illustrated on the left-hand plot of Figure 5.1, where $X_{M'}$ is used to denote that we are using the maximum for $t > 0$. As shown by Ruijter et al. (2015), Fourier based pricing methods using this PDF will therefore not be negatively affected by the Gibbs phenomenon and can thus achieve exponential convergence with the number of log-price grid points.

However, for the α-quantile options we price in this chapter, we require the distribution for the maximum (minimum) for $t \geq 0$. As all Lévy processes have the property that $X(0) = 0$, the value of the maximum for $t \geq 0$ cannot go below 0. Therefore obtaining $p_{X_M}(x)$ using the Spitzer-based scheme with the full number of dates alters the PDF so that it now has an abrupt discontinuity and a large spike at $t = 0$ as shown on the right-hand side of Figure 5.1. The large spike corresponds to a single probability mass equal to $\int_{-\infty}^{0} p_{X'_M}(x)dx$. We also note that the introduction of the discontinuity and spike has caused oscillations in the plot of the pdf via the Gibbs phenomenon.

The existence of the discontinuity means that, as described by Boyd (2001) and Gottlieb and Shu (1997) for example, we would no longer obtain exponential error convergence with grid size using these distributions to price options. We therefore use spectral filtering

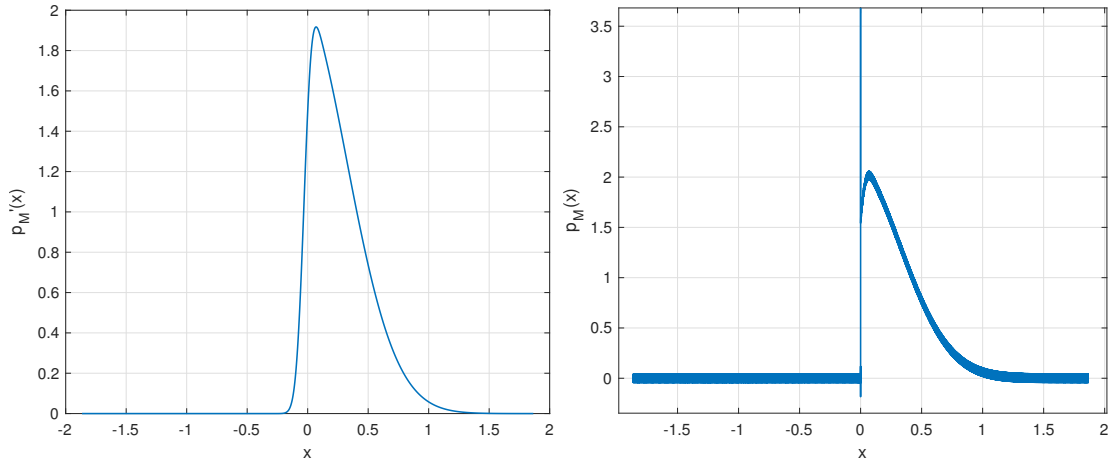

Figure 5.1: PDF of the maximum of a discretely monitored Brownian motion for $t > 0$ (left-hand side) and $t \geq 0$ (right-hand side) for 50 monitoring dates over 1 year, σ=0.4, risk free rate $r = 0.05$.

as successfully implemented in Chapter 3 to improve the error convergence for double-barrier options.

5.1.2 Numerical procedure for discretely monitored α-quantile options

In order to calculate the price of discretely monitored α-quantile options using the Fourier-z transform, we must express the time for the two independent random processes in terms of the number of monitoring dates. For N monitoring dates we calculate $j = \alpha N$ to the nearest integer. The pricing procedure is then

1. Compute the characteristic function $\Psi(\xi + i\alpha_{\mathrm{d}}, \Delta T)$ of the underlying transition density, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91) and $\Delta T = T/N$.

2. Use the Plemelj-Sokhotsky relations with the sinc-based Hilbert transform to factorise
$$\Phi(\xi, q) := 1 - q\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t) = \Phi_{\oplus}(\xi, q)\Phi_{\ominus}(\xi, q), \tag{5.9}$$
where $\sigma(\eta)$ is an exponential filter of order 12, as defined in Eq. (3.7) and q is selected according to the criteria specified by Abate and Whitt (1992b) for the inverse z-transform.

3. Calculate
$$\widetilde{\widetilde{p_{X_M}}}(\xi, q) = \frac{1}{\Phi_{\oplus}(\xi, q)\Phi_{\ominus}(0, q)} \tag{5.10}$$
$$\widetilde{\widetilde{p_{X_m}}}(\xi, q) = \frac{1}{\Phi_{\oplus}(0, q)\Phi_{\ominus}(\xi, q)} \tag{5.11}$$

4. Apply the inverse z-transform for j and $N-j$ dates respectively

$$\widehat{p_{X_M}}(\xi, j) = \mathcal{Z}^{-1}_{q\to j}\left[\widetilde{\widehat{p_{X_M}}}(\xi, q)\right] \tag{5.12}$$

$$\widehat{p_{X_m}}(\xi, N-j) = \mathcal{Z}^{-1}_{q\to N-j}\left[\widetilde{\widehat{p_{X_m}}}(\xi, q)\right] \tag{5.13}$$

5. Calculate the Fourier transform of the required probability distributions for X_M and X_m by obtaining the real parts of $\widehat{p_{X_M}}(\xi, j)$ and $\widehat{p_{X_m}}(\xi, N-j)$ directly in the Fourier domain using

$$\widehat{p^{\Re}_{X_M}}(\xi, j) = \frac{1}{2}\left[\widehat{p_{X_M}}(\xi, j) + \widehat{p_{X_M}}^{*}(-\xi, j)\right] \tag{5.14}$$

$$\widehat{p^{\Re}_{X_m}}(\xi, N-j) = \frac{1}{2}\left[\widehat{p_{X_m}}(\xi, N-j) + \widehat{p_{X_m}}^{*}(-\xi, N-j)\right] \tag{5.15}$$

where $\cdot^*$ denotes the complex conjugate.

6. Calculate the Fourier transform of the probability distribution for X_α over N monitoring dates as

$$\widehat{p_{X_\alpha}}(\xi, N) = \widehat{p^{\Re}_{X_M}}(\xi, j)\widehat{p^{\Re}_{X_m}}(\xi, N-j) \tag{5.16}$$

7. Calculate the price of the discretely monitored α-quantile option as

$$v(0,0) = \mathcal{F}^{-1}_{\xi\to x}\left[\sigma\left(\frac{\xi}{\xi_{\max}}\right)\widehat{p_{X_\alpha}}(\xi, N)\widehat{\phi}^{*}(\xi)\right](0) \tag{5.17}$$

where $\sigma(\eta)$ is an exponential filter of order 12, as defined in Eq. (3.7).

5.1.3 Comparison of results with Monte Carlo results

In Chapter 3, the results from the FL method from Feng and Linetsky (2008) and the FGM method from Fusai et al. (2016) provide a benchmark for each other. However, here we do not either have closed-form results or a highly accurate numerical method to provide a comparison.

We therefore use Monte Carlo pricing procedures with underlying assets driven by Gaussian, VG and Merton jump-diffusion processes as these have paths which can be simply constructed using the built-in Matlab commands for standard normal, gamma and Poisson random variables. Indeed, the ease of construction of these paths is the reason for selecting these processes which are different to the ones used in the work in Chapters 3 and 4 and which were selected due to their use by Fusai et al. (2016).

In all cases the risk neutral drift was calculated using $\mu_{\text{RN}} = \psi(-i)$, as described by Feng and Linetsky (2008), where $\psi(\xi)$ is the characteristic exponent of the Lévy process.

VG PROCESS. A single step over time Δt for the VG process was simulated using the following procedure:

1. Calculate the gamma distributed random variable dG which has a probability distribution $\Gamma(\frac{1}{\nu\Delta t}, \nu)$.

2. Calculate the process step:

$$dX = \mu_{\mathrm{RN}}\Delta t + \theta dG + \sigma\sqrt{dG}\zeta \tag{5.18}$$

   where ζ is a standard normal random variable.

MERTON JUMP-DIFFUSION PROCESS. A single step over time Δt for the Merton jump-diffusion process was simulated using the following procedure

1. Calculate the Poisson distributed random variable N_{P} with parameter $\Delta t\lambda$.

2. Calculate the process step:

$$dX = \mu_{\mathrm{RN}}\Delta t + \sigma\sqrt{dt}\zeta_1 + \mu_{\mathrm{J}}N_{\mathrm{P}} + \sigma_{\mathrm{J}}\sqrt{N_{\mathrm{P}}}\zeta_2 \tag{5.19}$$

   where ζ_1 and ζ_2 are independent standard normal random variables.

We used two methods to calculate the value for the α-quantile. For the first we simulated a path of N points and then found the j^{th} smallest value. For the second we combined the Dassios-Port-Wendel identity with the Monte Carlo method as in Ballotta and Kyprianou (2001). Two independent paths of length αN and $(1-\alpha)N$ dates are simulated and the sum of their respective infimum and supremum are calculated to provide an estimate of the α-quantile.

5.1.4 RESULTS FOR DISCRETELY MONITORED α-QUANTILE OPTIONS

Figures 5.2–5.4 show results for $\alpha = 0.75$ for N of 52, 252 and 1008 for the Gaussian, NIG, VG and Kou processes. In Table 5.1 we show the variation of the price with α for 252 monitoring dates and compare the results with that for the Monte Carlo method; this shows that we are within two standard deviations (corresponding to an approximate 95% confidence interval).

The error convergence for this method is extremely fast; we were able to achieve a CPU time of 2×10^{-2} or better for an error of 10^{-8} using the Gaussian process and Merton processes, and and 10^{-6} using the VG processes. When assessing the results, the error floor of 10^{-11} caused by the inverse z-transform means that there are too few data points above the noise floor to allow us to assess whether the error convergence is exponential or high order polynomial. This is not a concern for using the method in practice. However,

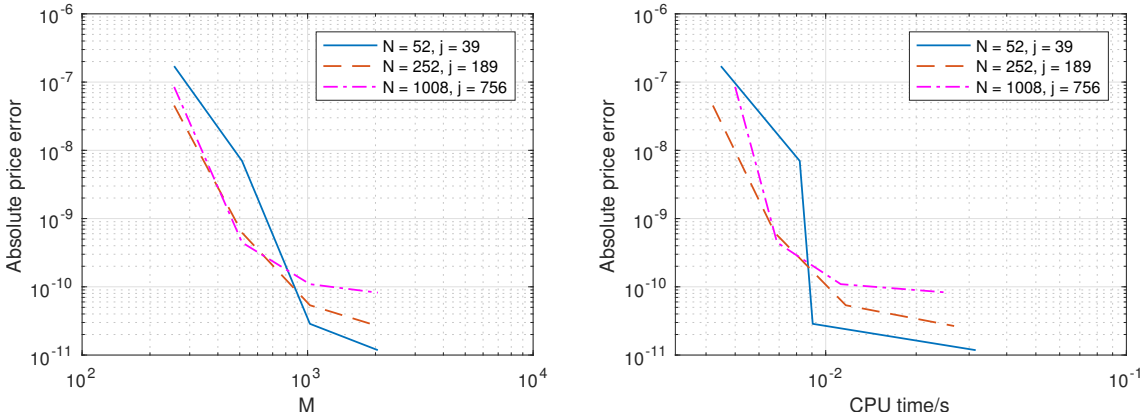

Figure 5.2: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Gaussian process.

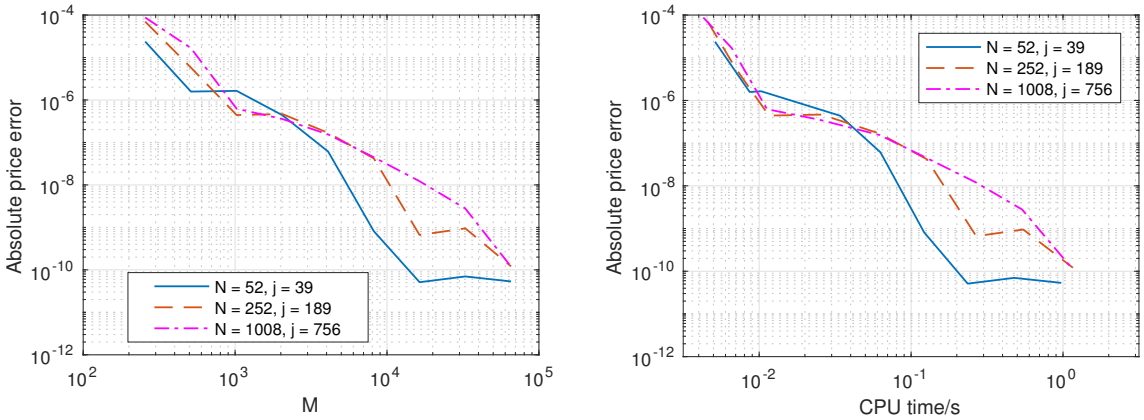

Figure 5.3: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a VG process.

a z-transform with a lower noise floor would allow us to understand the convergence of the techniques better. Table 5.1 shows that we are within two standard deviations for all test cases.

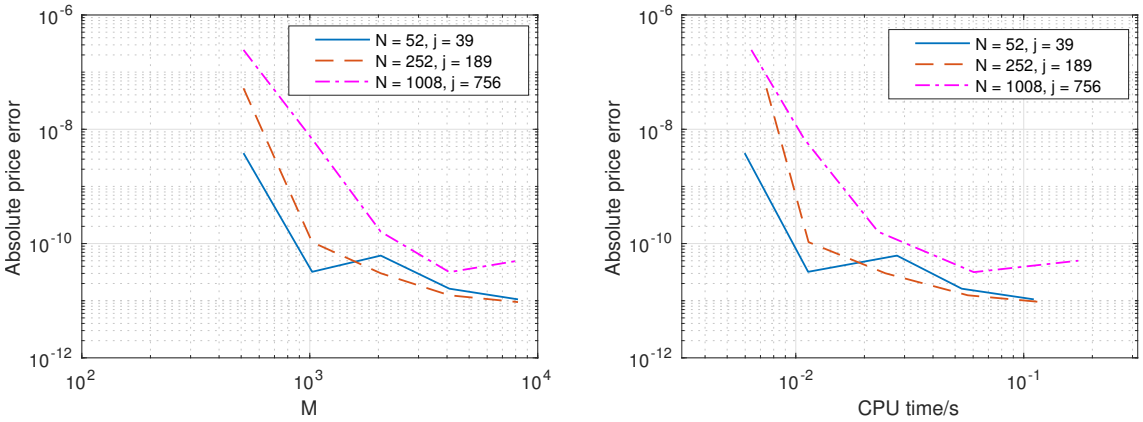

Figure 5.4: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Merton process.

Gaussian							
	Spitzer	Monte Carlo			Monte Carlo + DWP		
α	price	price	2σ	difference	price	2σ	difference
$\frac{2}{3}$	0.20847	0.20863	0.00044	-0.00016	0.20874	0.00046	0.00027
$\frac{3}{4}$	0.24346	0.24365	0.00059	-0.00019	0.24338	0.00058	7.8E-05
$\frac{5}{6}$	0.28541	0.2850	0.00057	0.00037	0.28520	0.00058	0.00021
VG							
	Spitzer	Monte Carlo			Monte Carlo + DWP		
α	price	price	2σ	difference	price	2σ	difference
$\frac{2}{3}$	0.15894	0.15888	0.00021	5.8E-05	0.15886	0.00022	7.4E-05
$\frac{3}{4}$	0.17649	0.17656	0.00024	-6.8E-05	0.17658	0.00023	-9.5E-05
$\frac{5}{6}$	0.19604	0.19589	0.00025	0.00015	0.19596	0.00025	8.5E-05
Merton jump-diffusion							
	Spitzer	Monte Carlo			Monte Carlo + DWP		
α	price	price	2σ	difference	price	2σ	difference
$\frac{2}{3}$	0.15957	0.15960	0.00022	-3.3E-05	0.15971	0.00024	-0.00014
$\frac{3}{4}$	0.17757	0.17762	0.00027	-4.5E-05	0.17757	0.00025	4.7E-07
$\frac{5}{6}$	0.19777	0.19779	0.00027	-1.9E-05	0.19778	0.00027	-1.8E-05

Table 5.1: Comparison between the value of a α-quantile option with 252 dates and 2^{16} price grid points compared with the value for the same contract using a Monte-Carlo approximation and a Monte-Carlo approximation combined with the Dassios-Port-Wendel identity. Notice that the prices calculated using the Spitzer-based method are within two standard deviations of the Monte-Carlo price.

5.1.5 Continuously monitored α-quantile options

We use the same relationship between the Spitzer identities with discrete and continuous monitoring which we exploited in Chapter 4 to price continuously monitored barrier op-

tions. Specifically, we utilise the conversion between $\Phi(\xi, q)$ in the Fourier-z domain and $\Phi_{\mathrm{c}}(\xi, s)$ in the Fourier-Laplace domain in Eq. (2.83). We also use the Abate and Whitt inverse Laplace transform which is described in Section 4.2.

5.1.5.1 Pricing procedure for continuously monitored α-quantile options

We adapt the pricing procedure for discretely monitored options to continuous monitoring by replacing Steps 1–3 in Section 5.1.2 with

1. Compute the characteristic exponent $\psi(\xi + i\alpha_{\mathrm{d}})$ of the underlying transition density, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

$$\Phi_{\mathrm{c}}(\xi, s) := s - \psi(\xi + i\alpha_{\mathrm{d}}) = \Phi_{\mathrm{c}\oplus}(\xi, s)\Phi_{\mathrm{c}\ominus}(\xi, s), \tag{5.20}$$

   where s is selected according to the criteria specified by Abate and Whitt (1995) for the inverse Laplace transform.

3. Calculate

$$\widetilde{\widetilde{p_{X_M}}}(\xi, s) = \frac{1}{\Phi_{\mathrm{c}\oplus}(\xi, s)\Phi_{\mathrm{c}\ominus}(0, s)}, \tag{5.21}$$

$$\widetilde{\widetilde{p_{X_m}}}(\xi, s) = \frac{1}{\Phi_{\mathrm{c}\oplus}(0, s)\Phi_{\mathrm{c}\ominus}(\xi, s)}. \tag{5.22}$$

The pricing procedure then continues as before with the inverse z-transform in Step 4 replaced by the inverse Laplace transform. Note that because $|\log \Phi_{\mathrm{c}}(\xi, s)|$ is increasing in $|\xi|$, unlike $\log \Phi(\xi, q)$ which is decreasing, similarly to the work in Chapter 4, a filter cannot be used for the factorisation in Step 1.

5.2 Results for continuously monitored α-quantile options

Figures 5.5–5.7 show results for a time to expiry of $T = 1$ and $\alpha = 0.75$ for the Gaussian, variance gamma and Merton processes.

The error convergence for this method is considerably inferior to that of the discretely monitored method. The convergence is approximately $O(1/M)$, with the rate increasing with M; this is consistent with the error convergence bounds for the factorisation of $\Phi_{\mathrm{c}}(\xi, s)$ in Chapter 4, although the absolute error is much higher as discussed further in Section 5.5.2. We must therefore revisit the question of whether it is better to use a discretely monitored scheme with a very small monitoring interval.

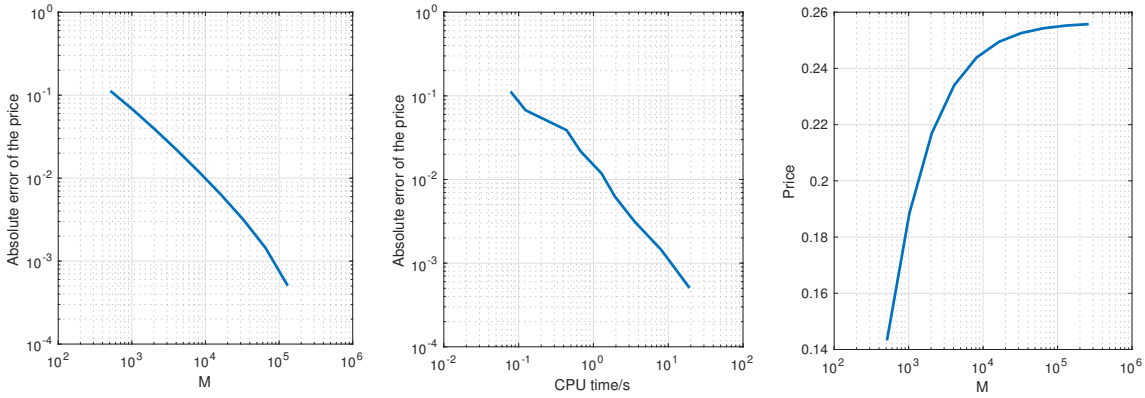

Figure 5.5: Results for the pricing error vs. M (left), pricing error vs. CPU time (centre) and price vs. M (right) with an underlying asset driven by a Gaussian process. Notice the convergence is approximately $O(1/M)$, with the rate increasing with M; this is consistent with the error convergence bounds for the factorisation of $\Phi_c(\xi, s)$ in Chapter 4.

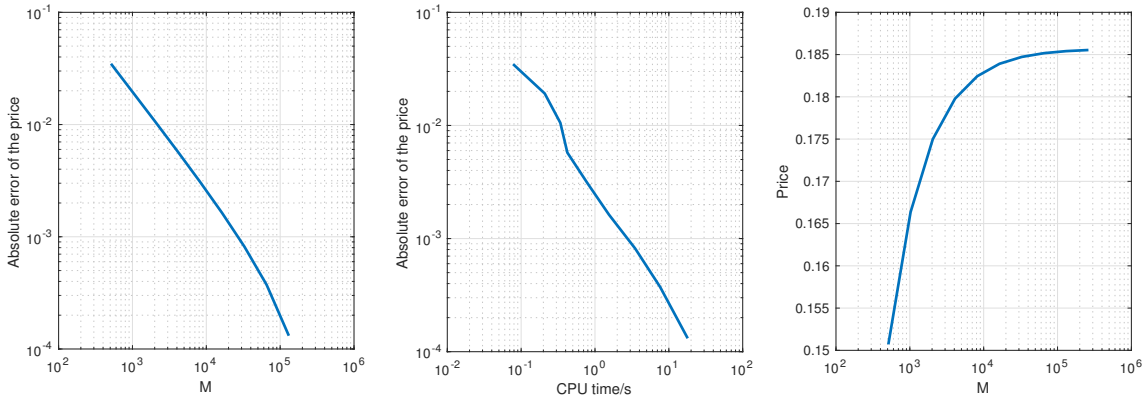

Figure 5.6: Results for the pricing error vs. M (left), pricing error vs. CPU time (centre) and price vs. M (right) with an underlying asset driven by a VG process. Notice the convergence is approximately $O(1/M)$, with the rate increasing with M; this is consistent with the error convergence bounds for the factorisation of $\Phi_c(\xi, s)$ in Chapter 4.

5.2.1 Comparison of the discretely monitored scheme with the continuously monitored pricing scheme

We perform a similar analysis to the one for barrier options in Section 4.4.2 to determine whether it is preferable to use a discrete pricing scheme with a high number of monitoring dates to approximate the continuously monitored pricing scheme. Figure 5.8 shows how the scheme converges with the number of monitoring dates for the Gaussian, NIG, VG and Kou processes. It can be seen that the convergence is approximately $O(1/N)$ for all these processes. As, in this case, the error for continuous monitoring is so high, we can achieve better results for the discretely monitored scheme with approximately 1000 monitoring dates for the Gaussian process and 300 monitoring dates for the other processes we consider. This is far lower than for barrier options where the performance of the method

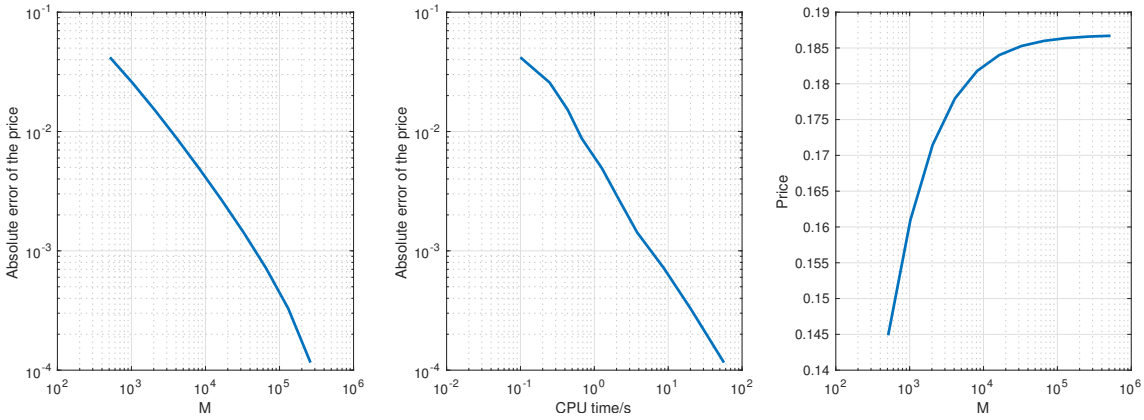

Figure 5.7: Results for the pricing error vs. M (left), pricing error vs. CPU time (centre) and price vs. M (right) with an underlying asset driven by a Merton process. Notice the convergence is approximately $O(1/M)$, with the rate increasing with M; this is consistent with the error convergence bounds for the factorisation of $\Phi_c(\xi, s)$ in Chapter 4.

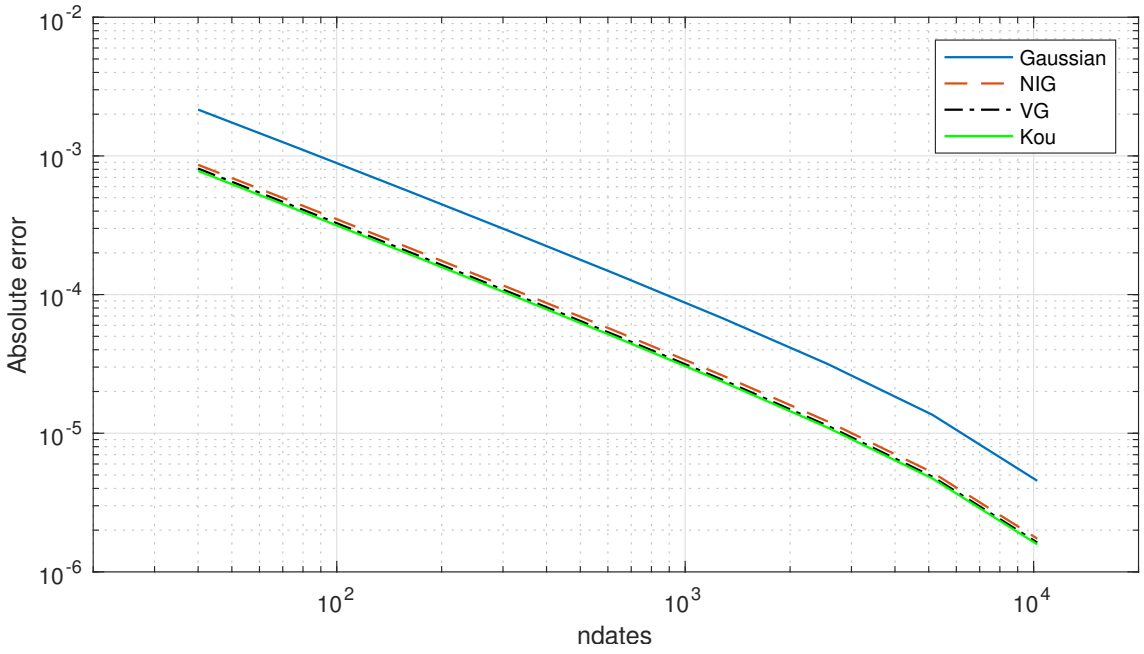

Figure 5.8: Convergence of the method for discrete monitoring to the result for continuous monitoring with an increasing number of monitoring dates for the Gaussian, NIG, VG and Kou processes. It can be seen that the convergence is approximately $O(1/N)$.

for continuous monitoring means that extremely high numbers of monitoring dates must be considered.

We can refer to the results in Figures 5.5–5.7 which show that the error with a CPU time of around $10^{-2}s$ is already far lower than both the error between the discrete and continuous methods and the continuous method itself. Therefore in this case it is preferable to use the method for discrete monitoring with a high number of dates as an approximation to the continuous case.

5.3 Perpetual Bermudan options

Bermudan options have the same payoff as European options but they can be exercised at a discrete set of dates rather than simply at a single expiry date. They can also be thought of a discrete version of American options and, indeed, the prices for Bermudan options are often used as a proxy for the value of American options (see e.g Feng and Lin, 2013). Perpetual Bermudan and American options have no expiry date and are therefore "live" until they are exercised. Pricing perpetual options is an easier problem than the valuation of those with a fixed expiry date as the infinite time horizon means that the optimal exercise boundary is constant rather than a function of time. Indeed, closed-form formulas exist for perpetual American options with simple processes, whereas there are none for finite expiry options.

We look at two methods for pricing perpetual Bermudan options. Firstly we implement the method by Green (2009) which uses residue calculus. We also implement a new method which uses Spitzer identities and show a novel way to calculate the optimal exercise barrier. New numerical truncation bounds for the log-price domain are specified; these are required due to the infinite time horizon.

For pricing single-barrier options in Chapters 3 and 4, we used the version of the Spitzer identity in Eq. (2.74) for $x \geq l$. In contrast, both the methods for pricing Bermudan options described here require the version of the identity for $x < l$. That is we require the distribution of $X(t_n)$ subject to t_n being the first time it has crossed the discretely monitored barrier l. This is illustrated in Figure 5.9 for a process with 10 monitoring dates and a barrier of -0.2. The value of $X(t_{10})$ for paths 1 and 2 would contribute to the distribution as they stay above l for monitoring dates 0–9 but go below l at the 10^{th} monitoring date. Note that path 1 is acceptable as, even though its path does go below l between dates 3 and 4, it is above l at the actual monitoring dates. Path 3 does not count towards the distribution as it is above l at the 10^{th} monitoring date and path 4 does not count as it goes below l at an earlier monitoring date.

Therefore, in the numerical procedures described below, we use $\widetilde{\widehat{p_{l-}}}(\xi, q) := P_{l-}(\xi, q)\,\Phi_{\ominus}(\xi, q)$ instead of $\widetilde{\widehat{p}}_l(\xi, q) := \frac{P_{l+}(\xi,q)}{\Phi_{\oplus}(\xi,q)}$, where $\widetilde{\widehat{p_{l-}}}(\xi, q)$ is used to distinguish the identity for $x < l$.

5.3.1 Green's residue method

This Spitzer identity for $\widetilde{\widehat{p_{l-}}}(\xi, q)$ is important for Green's pricing scheme for a perpetual Bermudan put option as this method is based on a combination of a first touch option which pays $K - D$ the first time the underlying price crosses a barrier D and an overshoot option which pays the difference between the barrier D and the underlying asset $S(t) = e^{X(t)}$, the first time the barrier is crossed. The advantage of using this formulation, as Green (2009) showed, is that these pricing formula can be written in such a way that the inverse

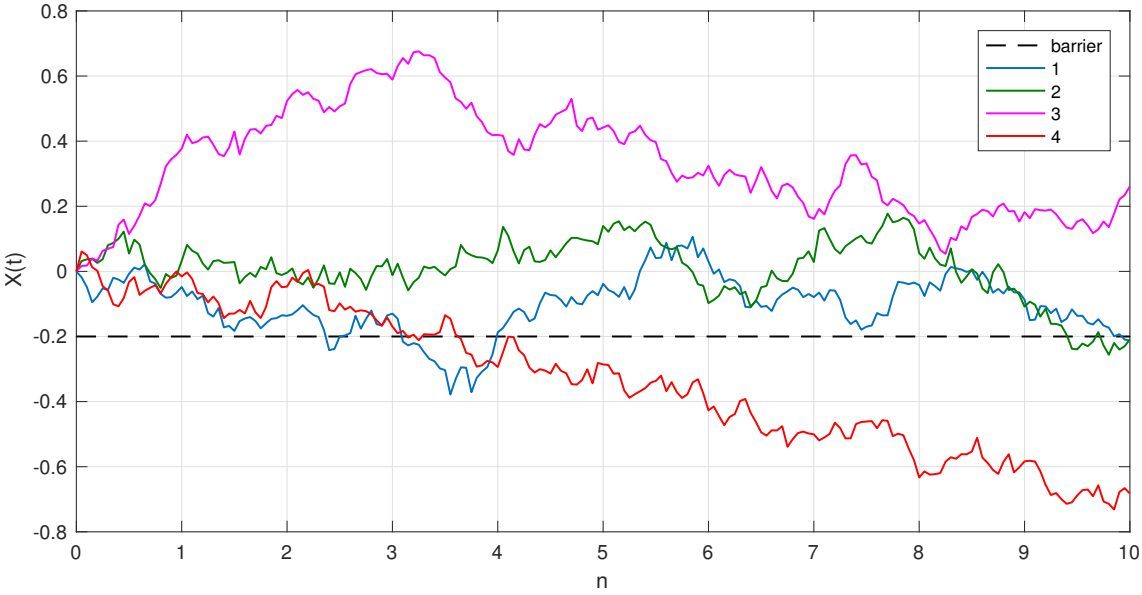

Figure 5.9: Examples of discretely monitored continuous random paths with 10 monitoring dates. Notice that only paths 1 and 2 would contribute to the calculated distribution.

Fourier and inverse z-transforms are no longer required.

A first touch option requires the probability that the first time the underlying asset crosses the barrier l is time n, as shown in Eq. (2.82), is

$$\mathrm{P}\left[\tau_l = n\right] = \int_{-\infty}^{l} p_{l-}(x,n)dx. \tag{5.23}$$

Substituting $\widetilde{\widetilde{p_{l-}}}(\xi,q)$ into this we get

$$\mathrm{P}\left[\tau_l = n\right] = \int_{-\infty}^{l} \mathcal{Z}_{q\to n}^{-1}\left[\mathcal{F}_{\xi\to x}^{-1}\left[\widetilde{\widetilde{p_{l-}}}(\xi,q)\right]\right] dx. \tag{5.24}$$

However it is also true that $\mathcal{F}_{\xi\to x}^{-1}\widetilde{\widetilde{p_{l-}}}(\xi,q)$ is, by definition, "$-$" with respect to l and so we can extend the integral to the entire range of x, i.e.

$$\mathrm{P}\left[\tau_l = n\right] = \mathcal{Z}_{q\to n}^{-1}\int_{-\infty}^{+\infty} \mathcal{F}_{\xi\to x}^{-1}\left[\widetilde{\widetilde{p_{l-}}}(\xi,q)\right] dx. \tag{5.25}$$

Moreover, it is straightforward to see that the integral of the entire range of x of a function is the same as its Fourier transform at $\xi = 0$, i.e. $\int_{-\infty}^{+\infty} f(x)dx = \int_{-\infty}^{+\infty} e^{ix0} f(x)dx = \mathcal{F}_{x\to\xi}[f(x)](0)$ so we now have

$$\mathrm{P}\left[\tau_l = n\right] = \mathcal{Z}_{q\to n}^{-1}\left[\widetilde{\widetilde{p_{l-}}}(0,q)\right]. \tag{5.26}$$

Green (2009) then obtained the price of the option as the expectation of the discounted

payoff at $n = 0$ as

$$v_{\mathrm{F}}(0,0) = (K-D)\sum_{n=1}^{\infty} e^{-r\Delta t n}\mathcal{Z}^{-1}_{q\to n}\left[\widetilde{\widehat{p_{l-}}}(0,q)\right], \tag{5.27}$$

where Δt is the time step between monitoring dates. The next insight by Green (2009) is that we can remove the summation and inverse z-transform, if we use $e^{-r\Delta t}$ as our z-transform parameter because the summation is the same as a forward z-transform. (We know that there is no payoff at the current date, i.e. $n = 0$, so we start the summation at $n = 1$). Then the price for a first touch option is

$$v_{\mathrm{F}}(0,0) = (K-D)\widetilde{\widehat{p_{l-}}}\left(0, e^{-r\Delta t}\right) = (K-D)P_{l-}\left(0, e^{-r\Delta t}\right)\Phi_{\ominus}\left(0, e^{-r\Delta t}\right). \tag{5.28}$$

The value of the payoff of an overshoot option at time n is the expected overshoot $D - S(t_n)$ conditional on n being the first time that the underlying asset process falls below l, τ_l and can be calculated as

$$E[e^{-r\tau_l}(D - S(\tau_l))|\tau_l = n\Delta t] = e^{-rn\Delta t}\int_{-\infty}^{l}(D - S_0e^x)p_{l-}(x,n)dx. \tag{5.29}$$

We can calculate the option value by using tower property to take the expectation over all discrete monitoring dates to obtain

$$v_{\mathrm{O}}(0,0) = \sum_{n=1}^{\infty} e^{-rn\Delta t}\int_{-\infty}^{l}(D - S_0e^x)p_{l-}(x,n)dx. \tag{5.30}$$

Once again, we can substitute $\widetilde{\widehat{p_{l-}}}(\xi,q)$ into this for

$$v_{\mathrm{O}}(0,0) = \sum_{n=1}^{\infty} e^{-rn\Delta t}\int_{-\infty}^{l}(D - S_0e^x)\mathcal{Z}^{-1}_{q\to n}\mathcal{F}^{-1}_{\xi\to x}\widetilde{\widehat{p_{l-}}}(\xi,q)dx. \tag{5.31}$$

By setting $q = e^{-r\Delta t}$, the summation and the inverse z-transform cancel each other as before and we extend the range of the integral to $\pm\infty$

$$v_{\mathrm{O}}(0,0) = \int_{-\infty}^{+\infty}(D - S_0e^x)^+\mathcal{F}^{-1}_{\xi\to x}\widetilde{\widehat{p_{l-}}}(\xi, e^{-r\Delta t})dx. \tag{5.32}$$

By using the Plancherel relation we can can express this as

$$v_{\mathrm{O}}(0,0) = \frac{1}{2\pi}\int_{-\infty}^{+\infty}\widehat{\phi}(-\xi)\widetilde{\widehat{p_{l-}}}(\xi, e^{-r\Delta t})dx. \tag{5.33}$$

If we use Eq. (2.92) for $\widehat{\phi}(\xi)$ for a put option with $a = -\infty$, $b = l$, $k = l$, $\alpha_{\mathrm{d}} > 0$ and

substituting $\widetilde{\widetilde{p_{l-}}}(\xi, q) := P_{l-}(\xi, q)\Phi_{\ominus}(\xi, q)$ then this can be simplified as

$$v_{\mathrm{O}}(0,0) = \frac{1}{2\pi}\int_{-\infty}^{+\infty} -D\frac{P_{l-}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t})}{\xi^2 + i\xi}dx. \tag{5.34}$$

Using the residue method the integral can be solved to give

$$v_{\mathrm{O}}(0,0) = DP_{l-}(0, e^{-r\Delta t})\Phi_{\ominus}(0, e^{-r\Delta t}) - DP_{l-}(-i, e^{-r\Delta t})\Phi_{\ominus}(-i, e^{-r\Delta t}). \tag{5.35}$$

As the price of a perpetual Bermudan option is the sum of a first touch option with payoff $(K - D)$ from Eq. (5.28) and the price of an overshoot option from Eq. (5.35) we have

$$\begin{aligned} v_{\mathrm{B}}(0,0) &= v_{\mathrm{F}}(0,0) + v_{\mathrm{O}}(0,0) \\ &= KP_{l-}(0, e^{-r\Delta t})\Phi_{\ominus}(0, e^{-r\Delta t}) - DP_{l-}(-i, e^{-r\Delta t})\Phi_{\ominus}(-i, e^{-r\Delta t}). \end{aligned} \tag{5.36}$$

Green (2009) showed that, using the idea that the optimal exercise barrier gives a maximum price, differentiating Eq. (5.36) and solving for D when the gradient is zero gives

$$D_{\mathrm{opt}} = K\frac{\Phi_{\ominus}(0, e^{-r\Delta t})}{\Phi_{\ominus}(-i, e^{-r\Delta t})} \tag{5.37}$$

for the optimal exercise barrier. Note that the optimal exercise barrier must be calculated first as this provides the value of l which $P(\xi, e^{-r\Delta t})$ must be decomposed around. Then

$$v_{\mathrm{B}}(0,0) = K\Phi_{\ominus}(0, e^{-r\Delta t})[P_{l-}(0, e^{-r\Delta t}) - P_{l-}(-i, e^{-r\Delta t})]. \tag{5.38}$$

gives the price for a perpetual Bermudan option.

5.3.2 New formulation based on Spitzer identities

Due to the success of the methods for pricing barrier options described in Chapter 3 we devised an alternate method for pricing perpetual Bermudan options, including a new way of calculating the optimal exercise boundary.

We first recognise that the expected return from exercising a perpetual Bermudan put (subject to us being below the optimal exercise barrier) at monitoring date n can be expressed as

$$E[e^{-rn\Delta t}(K - S(t_n))^+] = e^{-rn\Delta t}\int_{-\infty}^{l}(K - S_0e^x)p_{l-}(x, n)dx. \tag{5.39}$$

To obtain the value of the option we sum over all monitoring dates and we can also

substitute $\widetilde{p_{l-}}(\xi,q) := P_{l-}(\xi,q)\Phi_{\ominus}(\xi,q)$ for the probability distribution $p_{l-}(x,n)$

$$v_{\mathrm{B}}(0,0) = \sum_{n=1}^{\infty} e^{-rn\Delta t} \int_{-\infty}^{l} (K - S_0 e^x) \mathcal{Z}^{-1}_{q\to n} \mathcal{F}^{-1}_{\xi\to x} \widetilde{p_{l-}}(\xi,q) dx. \tag{5.40}$$

We can again use the trick by Green (2009) of using $q = e^{-r\Delta t}$ so that the summation cancels with the inverse z-transform to give

$$v_{\mathrm{B}}(0,0) = \int_{-\infty}^{l} (K - S_0 e^x) \mathcal{F}^{-1}_{\xi\to x} \widetilde{p_{l-}}(\xi, e^{-r\Delta t}) dx. \tag{5.41}$$

This integral can then be expressed in the Fourier domain using the Plancherel relation

$$v_{\mathrm{B}}(0,0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \widehat{\phi}^*(\xi) \widetilde{p_{l-}}(\xi, e^{-r\Delta t}) dx = \mathcal{F}^{-1}_{\xi\to x} \left[\widehat{\phi}^*(\xi) \widetilde{p_{l-}}(\xi, e^{-r\Delta t}) \right](0), \tag{5.42}$$

where $\phi(\xi)$ is the damped payoff for a put option from Eq. (2.92) with l being the optimal exercise boundary in the log-price domain.

5.3.2.1 Optimal exercise boundary Calculation

For this method, the calculation of the optimal exercise boundary is based on the idea that if the underlying asset is exactly at the optimal exercise boundary, i.e. $S_0 = S(0) = D_{\mathrm{opt}}$, then the value of the payoff from exercising the option is equal to the expected value from continuing to hold the option. Furthermore, the boundary used to calculate $p_{l-}(x,n)$ via $\widetilde{p_{l-}}(\xi,q) := P_{l-}(\xi,q)\Phi_{\ominus}(\xi,q)$ is $l = \log(D_{\mathrm{opt}}/S_0))$ and so $l = 0$ and therefore we can rewrite Eq. (5.39) as

$$\begin{aligned} v_{\mathrm{B},S_0=D}(0,0) &= \int_{-\infty}^{0} (K - S_0 e^x) \mathcal{F}^{-1}_{\xi\to x} P_{0-}(\xi, e^{-r\Delta t}) \Phi_{\ominus}(\xi, e^{-r\Delta t}) dx \\ &= \mathcal{F}^{-1}_{\xi\to x} \left[\widehat{\phi}^*(\xi) P_{0-}(\xi, e^{-r\Delta t}) \Phi_{\ominus}(\xi, e^{-r\Delta t}) \right](0). \end{aligned} \tag{5.43}$$

If we differentiate the expression on the first line of Eq. (5.43) with respect to S_0, we obtain

$$\frac{\partial v_{\mathrm{B},S_0=D}(0,0)}{\partial S_0} = -\int_{-\infty}^{0} e^x \mathcal{F}^{-1}_{\xi\to x} P_{0-}(\xi, e^{-r\Delta t}) \Phi_{\ominus}(\xi, e^{-r\Delta t}) dx \tag{5.44}$$

which is constant. Therefore, if we were to plot Eq. (5.43) against S_0 we obtain a straight line and the point where that line crosses the payoff function $(K - S_0)$ represents the optimal exercise barrier. This is illustrated in Figure 5.10. Therefore, by calculating Eq. 5.43 for two values of S_0 (D_1 and D_2) we can obtain the corresponding straight line equation with gradient m and y-axis intercept c. We can then calculate the optimal

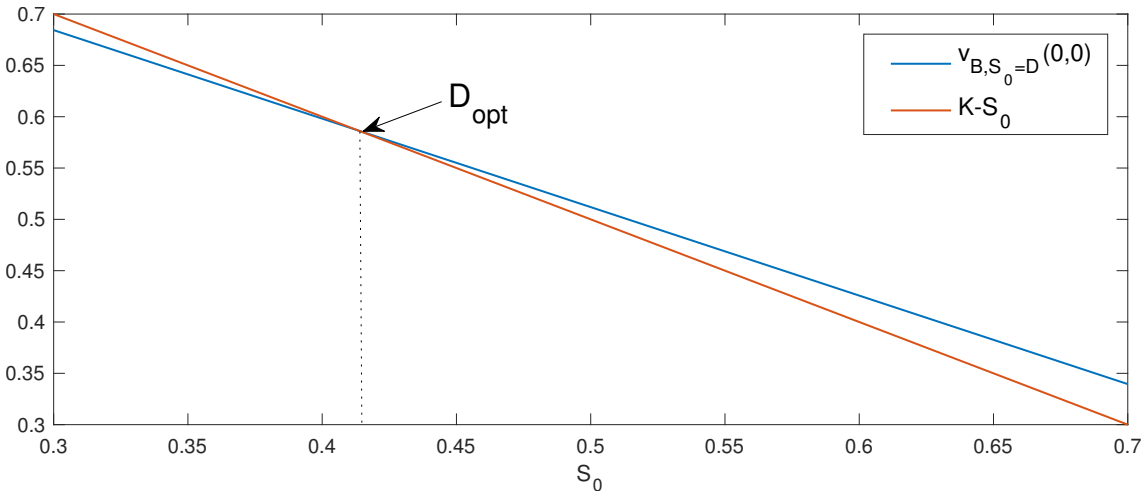

Figure 5.10: Illustration of the crossing point of the $(K - S_0)$ and $v_{\mathrm{B},S_0=D}(0,0)$ lines being used to calculate the optimal exercise boundary D_{opt}.

exercise boundary D_{opt}, corresponding to the point where the lines cross, as

$$D_{\mathrm{opt}} = \frac{K - c}{m + 1}. \tag{5.45}$$

We can also speed up the computational time by noting that if we set $D_1 = 0$ in Eq. (5.43) we obtain a price of

$$v_{\mathrm{B},S_0=D_1}(0,0) = \int_{-\infty}^{0} K\mathcal{F}^{-1}_{\xi\to x} P_{0-}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t})dx = KP_{0-}(0, e^{-r\Delta t})\Phi_{\ominus}(0, e^{-r\Delta t}) \tag{5.46}$$

which means that we only need to perform the inverse Fourier transform for the other calibration point(D_2). To avoid computational errors, rather than calculating the Spitzer identity with $l = 0$ we select $l = l_\epsilon$, where l_ϵ is the log-price domain grid step size spacing. This value does not depend on D and so the calculation of the gradient in Eq. (5.44) still returns a constant. The option price can then be calculated using Eq. (5.42) with $l = D_{\mathrm{opt}}/S_0$.

5.3.3 Truncation limits

For the Fourier based methods used for finite expiry boundary options, the range of the log-price domain grid for the numerical calculations was based on the cumulants of the distribution over a single time step and, for the parameters used in Chapters 3 and 4, were calculated as approximately ± 2. However, for perpetual options we must consider the shape of the probability distribution far in the future. This is especially true when the risk-free rate r is low as the contribution from future dates is discounted away extremely slowly. Therefore, the truncation limits used for finite expiry options are far too narrow

for this application. We base our calculation of the new bounds on the idea that the we should truncate the log-price domain at the value where the discount factor means any distortion of the distribution of the underlying asset process will have negligible effect on the final price calculation. We select a value that we wish the error to be below, i.e. $10^{-\lambda}$, and calculate the time it will take for the discount factor to be below this value

$$T_{\text{bound}} = \lambda \log(10)/r. \tag{5.47}$$

We then approximate the standard deviation of the underlying process at this time with

$$\sigma_{\text{bound}} = \sigma\sqrt{T_{\text{bound}}}, \tag{5.48}$$

where σ is the volatility of the underlying process, normalised to 1 year. The new bound for perpetual options is then

$$b = 6\sigma_{\text{bound}}. \tag{5.49}$$

5.3.4 Numerical procedure for Green's residue method

1. Compute the characteristic function $\Psi(\xi + i\alpha_{\text{d}})$ of the underlying transition density, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

$$\Phi(\xi, q) := 1 - e^{-r\Delta t}\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi + i\alpha_{\text{d}}, \Delta t) = \Phi_{\oplus}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t}), \tag{5.50}$$

   where $\sigma(\eta)$ is an exponential filter of order 12, as defined in Eq. (3.7). From this we can directly obtain $\Phi_{\ominus}(0, e^{-r\Delta t})$

3. Use the shift theorem to calculate $\Phi_{\ominus}(i, e^{-r\Delta t})$ as

$$\Phi_{\ominus}(i, e^{-r\Delta t}) := \mathcal{F}_{x\to\xi}\left[e^{-x}\mathcal{F}^{-1}_{\xi\to x}\left[\Phi_{\ominus}(\xi, e^{-r\Delta t})\right]\right](0) \tag{5.51}$$

4. Calculate the optimal exercise boundary

$$D_{\text{opt}} := K\frac{\Phi_{\ominus}(0, e^{-r\Delta t})}{\Phi_{\ominus}(i, e^{-r\Delta t})}, \tag{5.52}$$

   and compute $l = \log(D_{\text{opt}}/S_0)$.

5. Decompose $P(\xi, e^{-r\Delta t})$ around l

$$P(\xi, e^{-r\Delta t}) = P_{l+}(\xi, e^{-r\Delta t}) + P_{l-}(\xi, e^{-r\Delta t}), \tag{5.53}$$

and directly obtain $P_{l-}(0, e^{-r\Delta t})$.

6. Use the shift theorem to calculate $P_{l-}(i, e^{-r\Delta t})$ as

$$P_{l-}(i, e^{-r\Delta t}) := \mathcal{F}_{xi\to\xi}\left[e^{-x}\mathcal{F}^{-1}_{\xi\to x}\left[P_{l-}(0, e^{-r\Delta t})\right]\right](0). \tag{5.54}$$

7. Calculate the option price as

$$v(0,0) = K\Phi_{\ominus}(0, e^{-r\Delta t})\left[P_{l-}(0, e^{-r\Delta t}) - P_{l-}(i, e^{-r\Delta t})\right]. \tag{5.55}$$

5.3.5 Numerical procedure for new Spitzer based method

1. Compute the characteristic function $\Psi(\xi + i\alpha_{\mathrm{d}})$, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

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

$$\Phi(\xi, q) := 1 - e^{-r\Delta t}\sigma\left(\frac{\xi}{\xi_{\max}}\right)\Psi(\xi + i\alpha_{\mathrm{d}}, \Delta t) = \Phi_{\oplus}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t}), \tag{5.56}$$

where $\sigma(\eta)$ is an exponential filter of order 12, as defined in Eq. (3.7).

3. Decompose $P(\xi, e^{-r\Delta t})$ around l_ϵ (the log-price domain grid step)

$$P(\xi, e^{-r\Delta t}) = P_{l_\epsilon+}(\xi, e^{-r\Delta t}) + P_{l_\epsilon-}(\xi, e^{-r\Delta t}). \tag{5.57}$$

4. Calculate the PDF for the calibration as

$$\widehat{p_{l_\epsilon-}}(\xi, e^{-r\Delta t}) = P_{l_\epsilon-}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t}). \tag{5.58}$$

5. Setting D_1=0 and $D_2 \in (0, K)$ calculate

$$v_{\mathrm{D}_1}(0,0) = K\widehat{p_{l_\epsilon-}}(0, e^{-r\Delta t}), \tag{5.59}$$

$$v_{\mathrm{D}_2}(0,0) = \mathcal{F}^{-1}_{\xi\to x}\left[\widehat{\phi}^*(\xi)\widehat{p_{l_\epsilon-}}(0, e^{-r\Delta t})\right](0). \tag{5.60}$$

here, $\widehat{\phi}(\xi)$ is the Fourier transform of the damped payoff $\phi(x)$ where $x = \log(S(t)/D_2)$.

6. Calculate the optimal exercise barrier as

$$D_{\mathrm{opt}} = \frac{\left[K - v_{\mathrm{D}_1}(0,0)\right]D_2}{v_{\mathrm{D}_2}(0,0) - v_{\mathrm{D}_1}(0,0) + D_2}, \tag{5.61}$$

and compute $l = \log(D_{\mathrm{opt}}/S_0)$.

7. Decompose $P(\xi, e^{-r\Delta t})$ around l

$$P(\xi, e^{-r\Delta t}) = P_{l+}(\xi, e^{-r\Delta t}) + P_{l-}(\xi, e^{-r\Delta t}). \tag{5.62}$$

8. Calculate the PDF for the price as

$$\widehat{p_{l-}}(\xi, e^{-r\Delta t}) = P_{l-}(\xi, e^{-r\Delta t})\Phi_{\ominus}(\xi, e^{-r\Delta t}) \tag{5.63}$$

9. Calculate the option price as

$$v(0,0) = \mathcal{F}^{-1}_{\xi\to x}\left[\widehat{\phi}^*(\xi)\widehat{p_{l-}}(0, e^{-r\Delta t})\right](0). \tag{5.64}$$

where $\widehat{\phi}(\xi)$ is the Fourier transform of the damped payoff $\phi(x)$ where $x = \log(S(t)/S_0)$.

5.3.6 Results for perpetual Bermudan options with the Gaussian process

Figures 5.11–5.16 show results for the two methods for pricing perpetual Bermudan options. The results labelled "Green's" are from the residue method described in Section 5.3.1 and those labelled "Spitzer" are from the new method described in Section 5.3.2. Results are presented for error vs. M and CPU time for risk free rates r of 0.02, 0.05 and 0.1. Using a Gaussian process for the underlying asset means the results as $\Delta t \to 0$ can be compared with closed-form calculations for perpetual American options. The convergence of the price for Bermudan options to the continuous case is shown in Tables 5.2 and 5.3 and we discuss the further verification of these results in Section 5.3.8.

Both methods perform well, with polynomial error convergence with grid size M of at least $O(1/M^2)$. Furthermore, as $\Delta t \to 0$, the results approach those for the corresponding perpetual American option. The new Spitzer based method outperformed Green's residue method, achieving an error of 10^{-7} in approximately one tenth of the CPU time required for the residue method. In contrast, it is interesting to note that the convergence of the barrier is slower for the Spitzer based method than Greens residue method at $O(1/N)$ compared to $O(1/N^2)$. Moreover the barriers of both methods converge no faster than the price so we can see that the barrier error has a limited effect on the price error. This can be understood by considering Figure 5.10. In this plot $(K - S_0)$ represents the value of exercising the option, whereas $v(0,0)$ represents the value keeping the option. Close to the optimal exercise barrier D_{opt} this difference is extremely small and so a small error in the calculation of D_{opt} makes a very small difference to the pricing error.

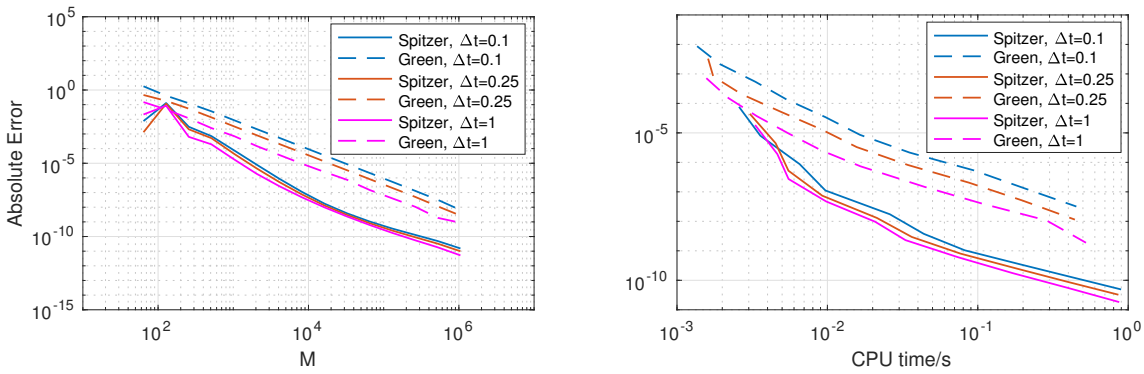

Figure 5.11: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.02$. Notice that the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is better than for the residue method described in Section 5.3.1 labelled "Green".

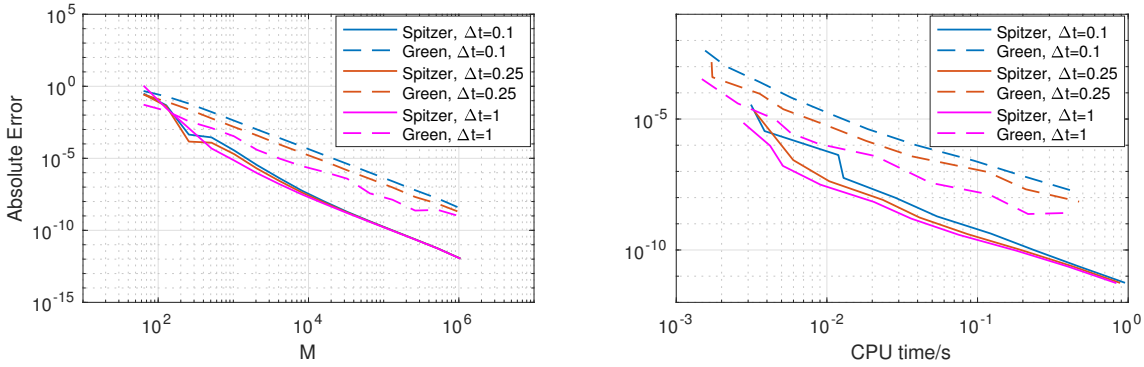

Figure 5.12: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.05$. Notice that the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is better than for the residue method described in Section 5.3.1 labelled "Green".

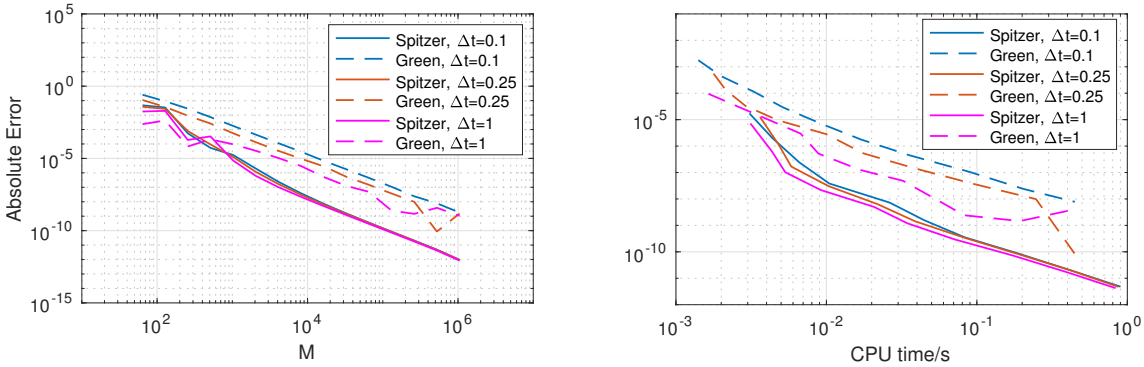

Figure 5.13: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.1$. Notice that the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is better than for the residue method described in Section 5.3.1 labelled "Green".

Δt	$r = 0.1$	$r = 0.05$	$r = 0.2$
1	0.20169919	0.33181098	0.53155442
0.5	0.20737414	0.33522271	0.53328533
0.25	0.21021533	0.33695115	0.53414348
0.1	0.21197983	0.33798846	0.53465453
0.01	0.21305064	0.33861012	0.53495870
American	0.21317038	0.33867902	0.53499224

Table 5.2: Results for the Spitzer based method for perpetual Bermudan options with $M = 2^{20}$ showing the convergence to the price for a perpetual American option.

Δt	$r = 0.1$	$r = 0.05$	$r = 0.2$
1	0.20169919	0.33181098	0.53155362
0.5	0.20737414	0.33522271	0.53328453
0.25	0.21021533	0.33695115	0.53414268
0.1	0.21197984	0.33798846	0.53465373
0.01	0.21305065	0.33861013	0.53495791
American	0.21317038	0.33867902	0.53499224

Table 5.3: Results for Green's residue method for perpetual Bermudan options with $M = 2^{20}$ showing the convergence to the price for a perpetual American option.

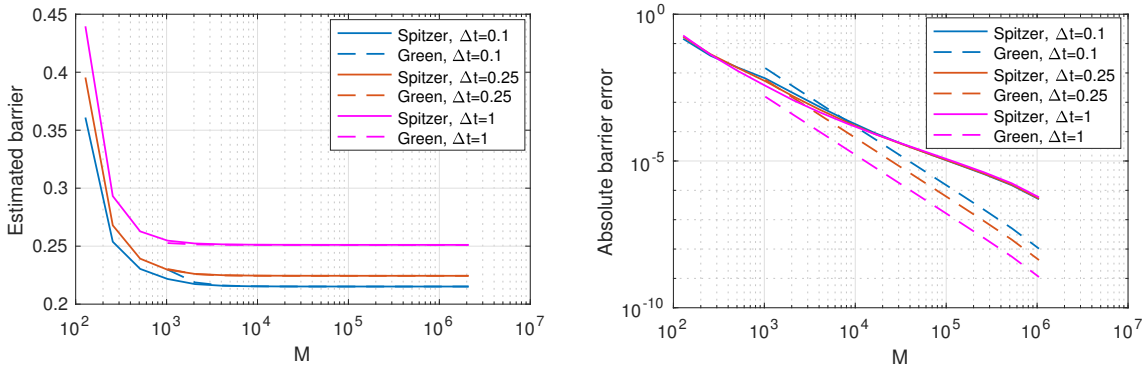

Figure 5.14: Optimal exercise barrier convergence vs. grid size M with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.02$. Notice that, in contrast to the pricing error convergence, the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is worse than for the residue method described in Section 5.3.1 labelled "Green".

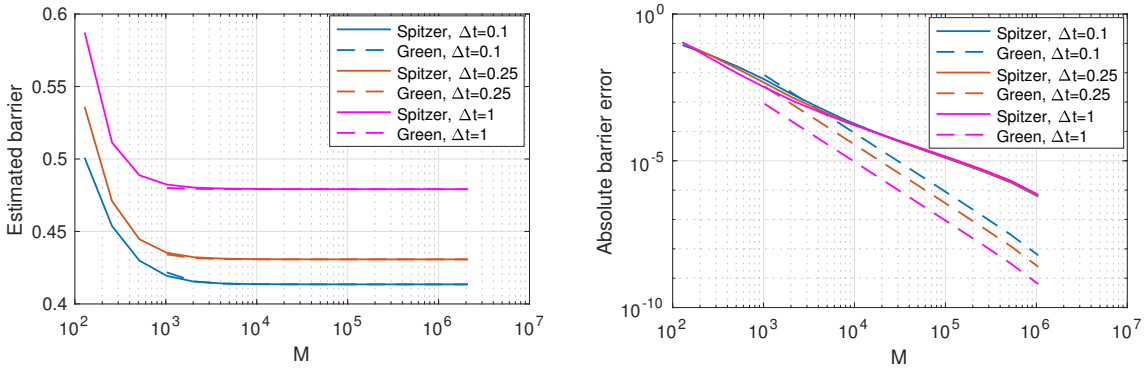

Figure 5.15: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.05$. Notice that, in contrast to the pricing error convergence, the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is worse than for the residue method described in Section 5.3.1 labelled "Green".

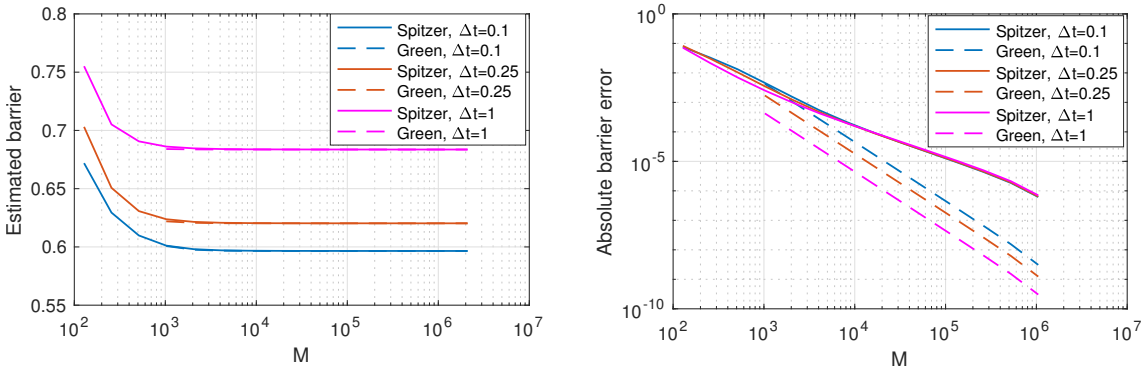

Figure 5.16: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Gaussian process with a risk free rate of $r = 0.1$. Notice that, in contrast to the pricing error convergence, the error convergence for the new method described in Section 5.3.2, labelled "Spitzer", is worse than for the residue method described in Section 5.3.1 labelled "Green".

Δt	$r = 0.1$	$r = 0.05$	$r = 0.02$
1	0.68360194	0.47916587	0.25110359
0.5	0.64678075	0.45056945	0.23520061
0.25	0.62026498	0.43075700	0.22442626
0.1	0.59653531	0.41350275	0.21519240
0.01	0.56851217	0.39363474	0.20472988
American	0.55555556	0.38461538	0.20000000

Table 5.4: Barrier calculated using the new Spitzer based method for perpetual Bermudan options with $M = 2^{20}$ showing the convergence to the barrier for a perpetual American option.

Δt	$r = 0.1$	$r = 0.05$	$r = 0.02$
1	0.68360124	0.47916517	0.25110066
0.5	0.64678008	0.45056880	0.23519670
0.25	0.62026434	0.43075637	0.22442094
0.1	0.59653469	0.41350214	0.21518427
0.01	0.56851159	0.39363417	0.20470522
American	0.55555556	0.38461538	0.20000000

Table 5.5: Barrier calculated using Green's residue method for perpetual Bermudan options with $M = 2^{20}$ showing the convergence to the barrier for a perpetual American option.

5.3.7 Results for perpetual Bermudan options with other Lévy processes

Figures 5.17–5.28 show results for the price and optimal exercise barrier vs. the number of grid points M and CPU time, with the underlying asset driven by the VG and Merton processes. It can be seen that, similarly to the performance with a Gaussian process, the error convergence speed is $O(1/M^2)$. The error is calculated as a precision compared to the result with the maximum number of grid points M and we discuss the verification of these results in Section 5.3.8.

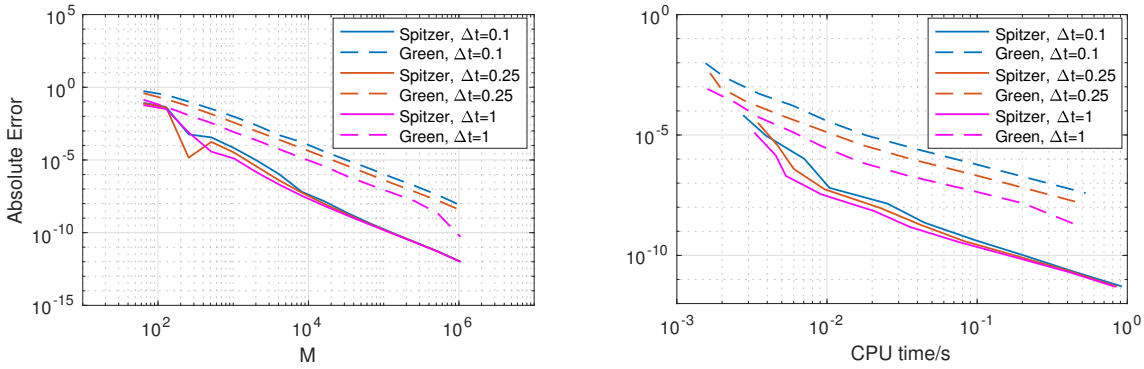

Figure 5.17: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a VG process with a risk free rate of $r = 0.02$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

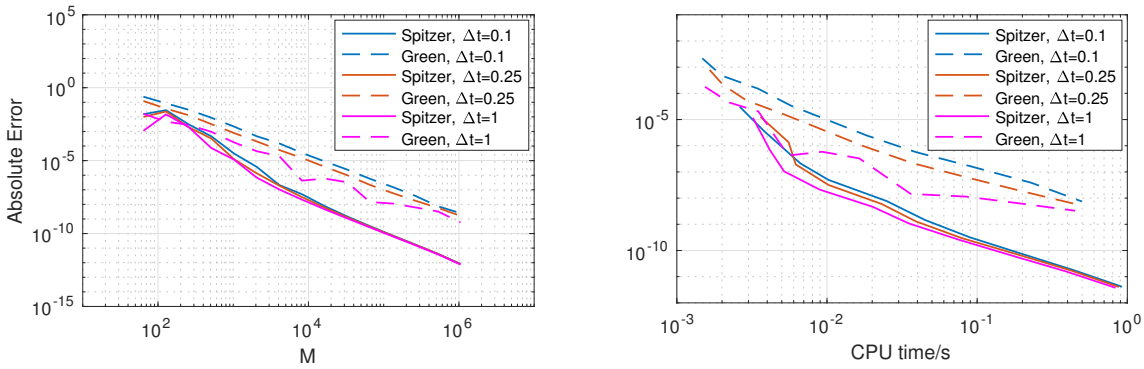

Figure 5.18: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a VG process with a risk free rate of $r = 0.05$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

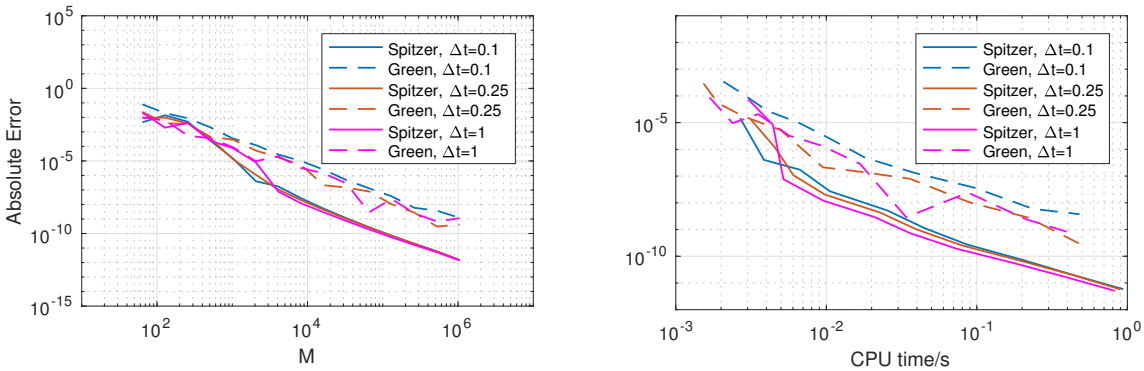

Figure 5.19: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a VG process with a risk free rate of $r = 0.1$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

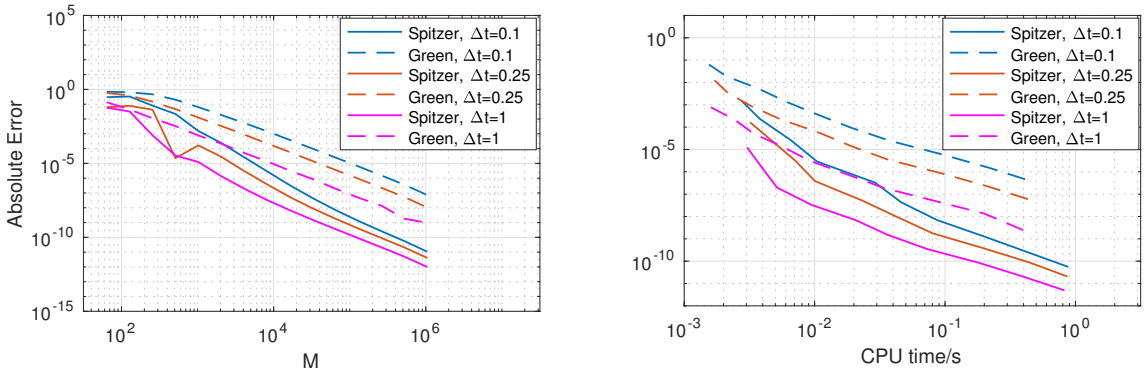

Figure 5.20: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.02$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

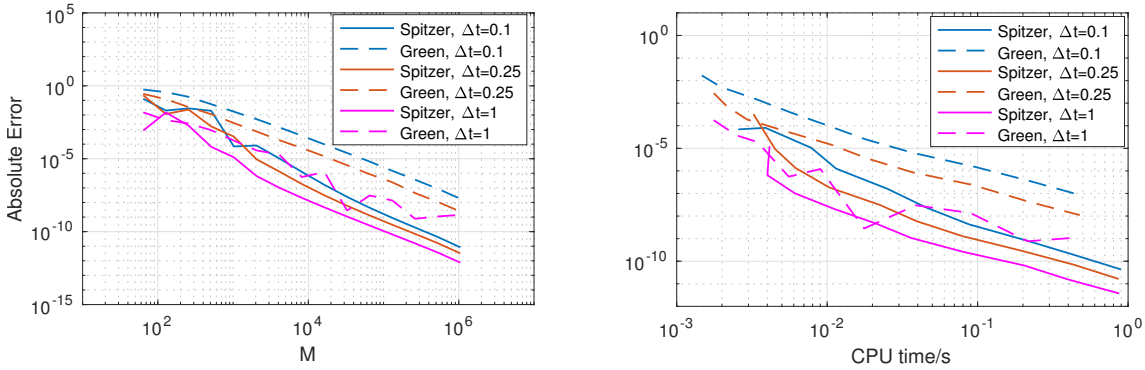

Figure 5.21: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.05$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

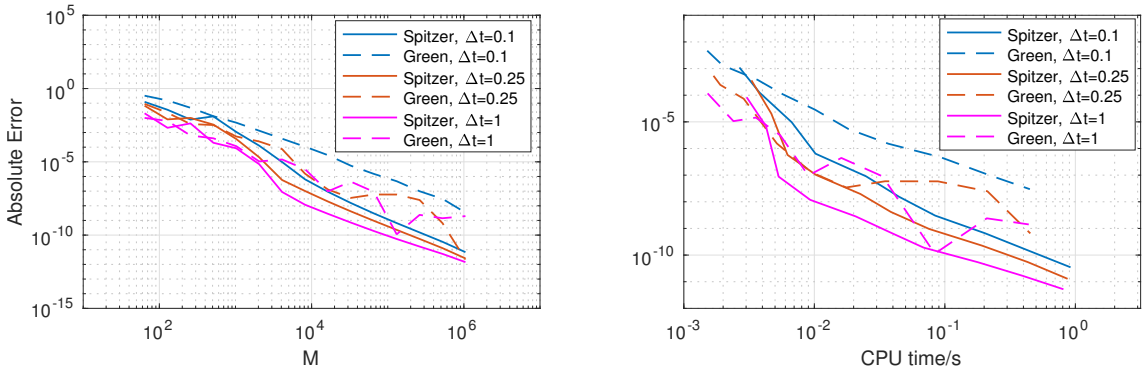

Figure 5.22: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.1$. Notice that, similarly to the technique with a Gaussian process, the error convergence speed is $O(1/M^2)$ and the "Spitzer" method has lower errors.

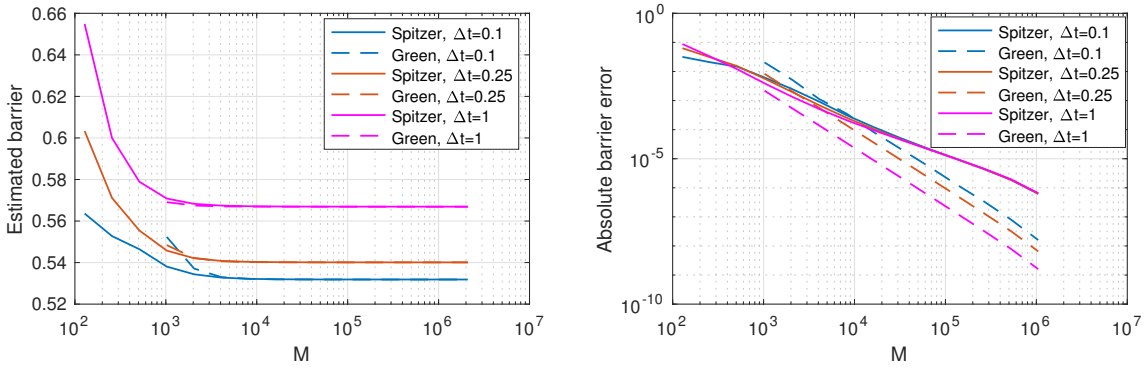

Figure 5.23: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a VG process with a risk free rate of $r = 0.02$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

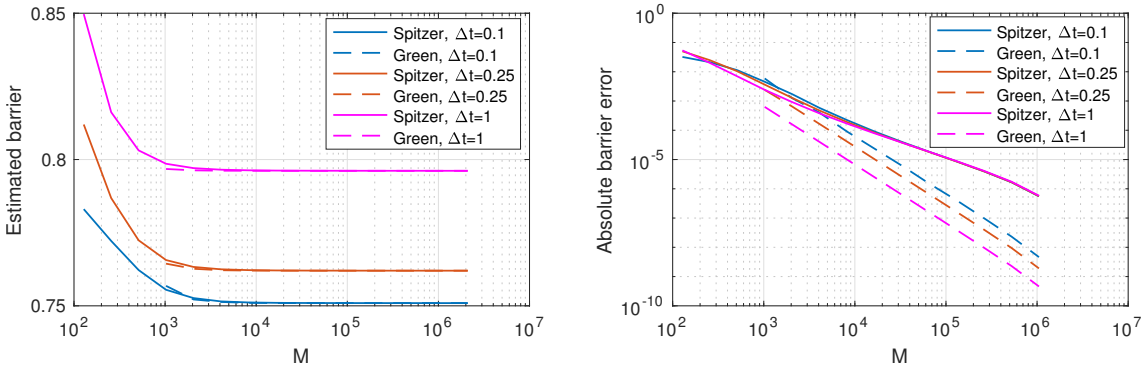

Figure 5.24: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a VG process with a risk free rate of $r = 0.05$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

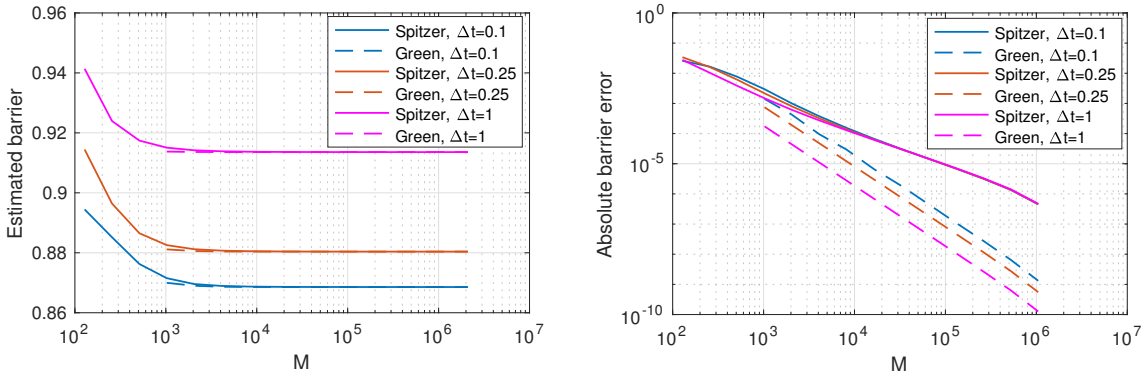

Figure 5.25: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a VG process with a risk free rate of $r = 0.1$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

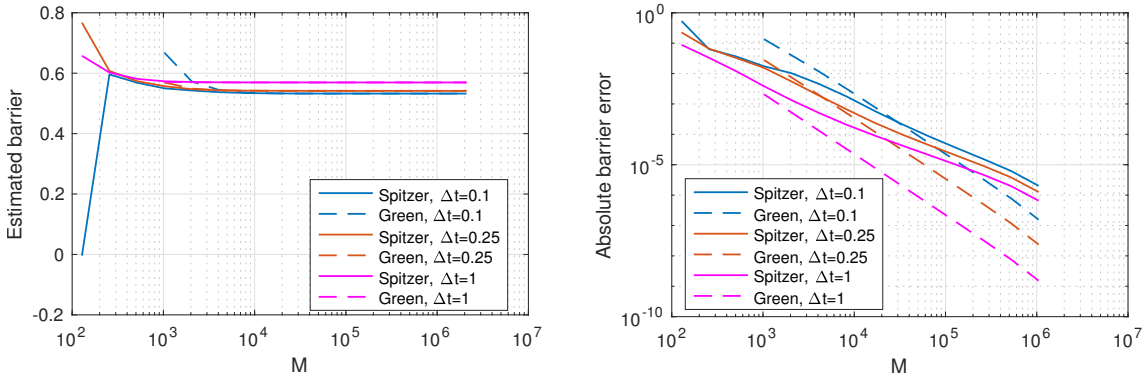

Figure 5.26: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.02$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

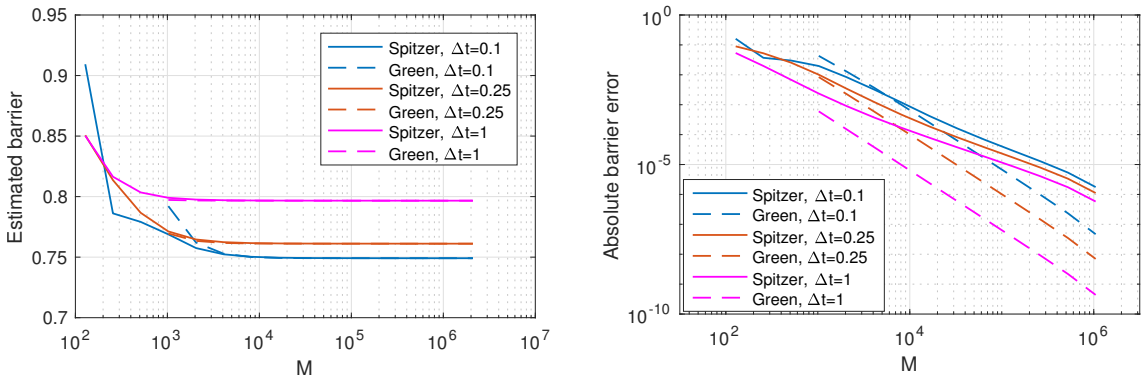

Figure 5.27: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.05$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

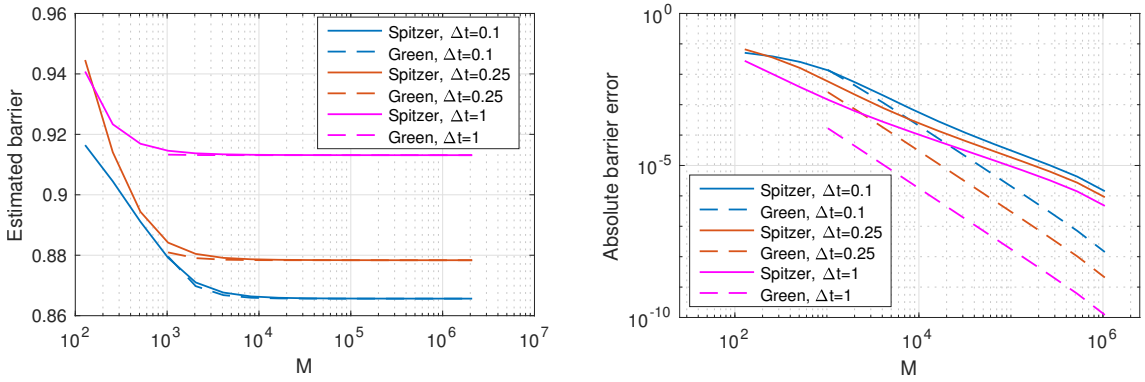

Figure 5.28: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Merton jump-diffusion process with a risk free rate of $r = 0.1$. Notice that, similarly to the technique with a Gaussian process, the "Spitzer" method has slower error convergence.

5.3.8 Comparison of results with Monte Carlo method

In Tables 5.2–5.5 in Section 5.3.6 we showed that that the results with the Gaussian process converge to the closed-form solution by Merton (1973) as $\Delta t \to 0$. However, we do not have a general closed-form solution for other Lévy processes so, similarly to the work on α-quantile options, a Monte Carlo method was used as an approximation.

We wrote Monte Carlo pricing procedures with underlying assets driven by Gaussian, VG and Merton jump-diffusion processes using the same techniques to simulate the paths as described in Section 5.1.3.

Although the discrete nature of perpetual Bermudan options is appropriate for a Monte Carlo simulation, the absence of an expiry date means that a Monte Carlo scheme with a finite number of dates will not represent the contract accurately. However, we truncate the Monte Carlo simulation so far in the future that the effect of disregarding these dates is less than the standard deviation of the Monte Carlo method itself. (Clearly this is more feasible for high discount factors and large time steps as the effect of future dates is discounted away more rapidly.) Finding the optimal exercise barrier uses the same philosophy as the new method described in Section 5.3.2.1, i.e. the price was calculated with $S_0 = D_1$ and $S_0 = D_2$ and we found the intersection between the line through these points and the straight line for the payoff $K - S_0$.

We calculated an approximate 95% confidence interval (corresponding to 2 standard deviations) for the Monte Carlo methods and Table 5.6 shows that the results for our new methods described in Sections 5.3.4 and 5.3.5 are within this range for all cases tested. Furthermore, we can see that the results for Green's method and the new Spitzer based method are the same.

Gaussian							
Parameters		Monte Carlo		Spitzer		Green	
r	Δt	price	2σ	price	difference	price	difference
0.05	1	0.331801	8.41E-05	0.331811	-9.98E-06	0.331811	-9.98E-06
0.05	0.5	0.335196	8.92E-05	0.335223	2.671E-05	0.335223	2.671E-05
0.05	0.25	0.336947	9.00E-05	0.336951	-4.15E-06	0.336951	-4.15E-06
0.05	0.1	0.337911	1.66E-04	0.337988	-7.746E-05	0.337988	-7.746E-05
VG							
Parameters		Monte Carlo		Spitzer		Green	
r	Δt	price	2σ	price	difference	price	difference
0.05	1	0.120225	2.15E-04	0.120237	-1.15E-05	0.120237	-1.15E-05
0.05	0.5	0.123415	2.52E-04	0.123298	1.16E-04	0.123298	1.16E-04
0.05	0.25	0.124996	2.68E-04	0.124919	7.67E-05	0.124919	7.67E-05
0.05	0.1	0.125775	2.54E-04	0.125959	-1.83E-04	0.125959	-1.83E-04
0.02	1	0.247040	3.54E-04	0.247078	-3.83E-05	0.247078	-3.83E-05
0.02	0.1	0.249747	3.56E-04	0.249756	-9.37E-06	0.249756	-9.37E-06
Merton jump-diffusion							
Parameters		Monte Carlo		Spitzer		Green	
r	Δt	price	2σ	price	difference	price	difference
0.05	1	0.119755	2.58E-04	0.119856	-1.01E-04	0.119856	-1.01E-04
0.05	0.5	0.123085	2.39E-04	0.122993	9.21E-05	0.122993	9.21E-05
0.05	0.25	0.124636	2.77E-04	0.124674	-3.82E-05	0.124674	-3.82E-05
0.05	0.1	0.125925	2.39E-04	0.125767	1.58E-04	0.125767	1.58E-04

Table 5.6: Comparison between the value of a perpetual Bermudan option with $M = 2^{20}$ compared with the value for the same contract using a Monte Carlo approximation. Notice that the prices calculated using the new Spitzer based method and Green's method are the same and within two standard deviations of the Monte Carlo price.

5.4 Perpetual American options

Similarly to the work on continuously monitored barrier options in Chapter 4 and quantile options in Section 5.1.5 of this chapter, the Spitzer based pricing method for perpetual Bermudan options can be extended to continuous monitoring, i.e. perpetual American options. Unlike the previous examples for option pricing with continuous monitoring where the application of option pricing was a motivating example for techniques which have relevance for other fields, the continuous (i.e. American) case is commonly used in financial contracts.

Once again, we utilise the conversion between $\Phi(\xi, q)$ in the Fourier-z domain and $\Phi_{\mathrm{c}}(\xi, s)$ in the Fourier-Laplace domain in Eq. (2.83). Here, as $q = e^{-r\Delta t}$ then $s = r$. Both methods described in Sections 5.3.4 and 5.3.5 are converted to continuous monitoring and we compare the results in Section 5.5.

5.4.1 Pricing procedure for perpetual American options

For both methods we adapt the pricing procedure for discretely monitored options to continuous monitoring by replacing Steps 1–2 in the procedures described in Sections 5.3.4 and 5.3.5 with

1. Compute the characteristic exponent $\psi(\xi+i\alpha_{\mathrm{d}})$ of the underlying transition density, where α_{d} is the damping parameter introduced in Section 2.2.3, Eq. (2.91).

2. Use the Plemelj-Sokhotsky relations with the sinc-based Hilbert transform to factorise
$$\Phi_{\mathrm{c}}(\xi,r) := r - \psi(\xi + i\alpha_{\mathrm{d}}) = \Phi_{\mathrm{c}\oplus}(\xi,r)\Phi_{\mathrm{c}\ominus}(\xi,r), \tag{5.65}$$
where r is the risk free rate.

We then continue with the calculations in Section 5.3.4 or 5.3.5 as before but with $\Phi_{\mathrm{c}}(\xi,r)$ in place of $\Phi(\xi,e^{-r\Delta t})$.

5.5 Results for perpetual American options

Figures 5.29–5.34 show results for both methods with the Gaussian, VG and Merton jump-diffusion processes.

For assessing the accuracy with the Gaussian process we have the advantage over Bermudan options that closed-form formulas exist for the calculation of both the barrier and option price and so the results in Figures 5.29 and 5.32 use the closed-form calculation from Merton (1973) to calculate the error. For the other process we do not have a closed-form result, and the continuous nature of American options means that they cannot be accurately represented using Monte Carlo methods which are inherently discrete. Therefore the absolute errors displayed in Figures 5.30 - 5.31 and 5.33 - 5.34 are calculated against the result for the same method with the maximum number of FFT points.

In contrast to the results for perpetual Bermudan options, the performance of the two methods are very different. Our new Spitzer-based method is far superior giving errors of $\approx 10^{-6}$ in $10^{-2}s$ or less. For the Gaussian and Merton processes, Green's method fails to reach an error level of 10^{-6}, and for the VG process it reaches this level about 100 times slower than our new method. The shape of the error convergence of the two methods with M is also interesting. For our new method, we see a $O(1/M^3)$ error convergence for lower values of M which then sharply transitions to a convergence of approximately $O(1/M)$–$O(1/M^{\frac{3}{2}})$ for $M \geq 10^3 - 10^4$. The two distinct shapes are indicative of different error mechanisms dominating the convergence, depending on the value of M. Moreover, the shape of the convergence for higher values of M is consistent with the factorisation error bound calculated for continuously monitored barrier options in Chapter 4. The shape of

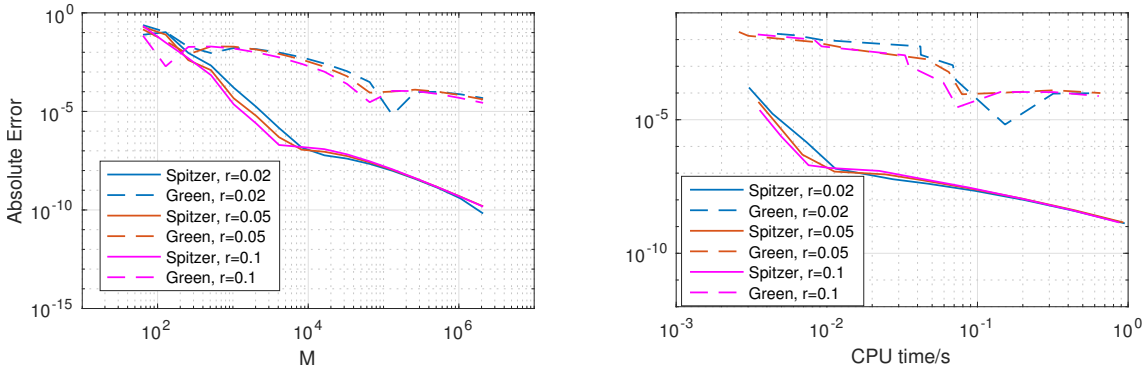

Figure 5.29: Results for the pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Gaussian process. The error is calculated vs. the closed-form expression by Merton, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

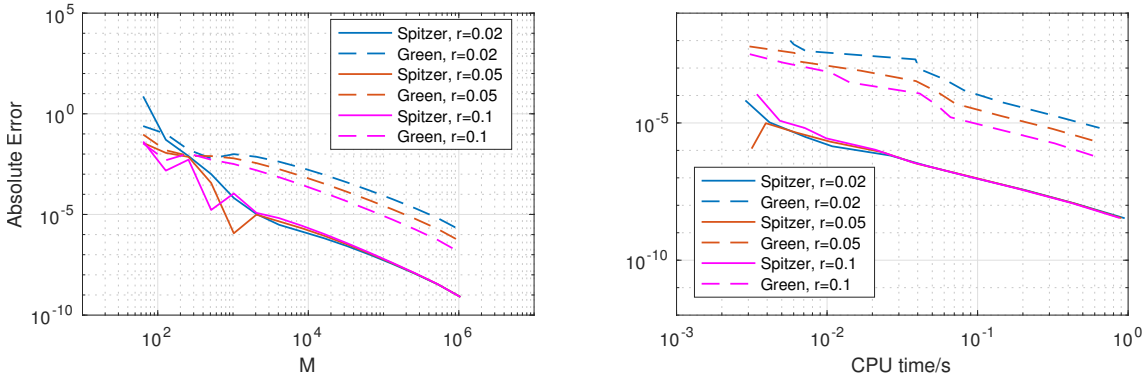

Figure 5.30: Pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a VG process. The error is calculated vs. the numerical result with the maximum grid size, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

the error convergence with M shown by Green's method is also similar to this bound, although less clearly so due to the non-monotonicity of the error results. We discuss this further in Section 5.5.2 with reference to both these methods.

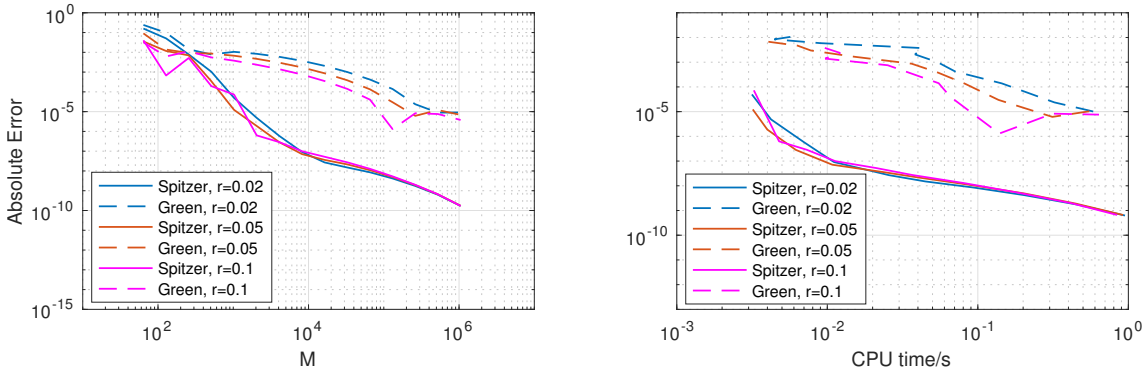

Figure 5.31: Pricing error convergence vs. grid size M and CPU time with an underlying asset driven by a Merton jump-diffusion process. The error is calculated vs. the numerical result with the maximum grid size, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

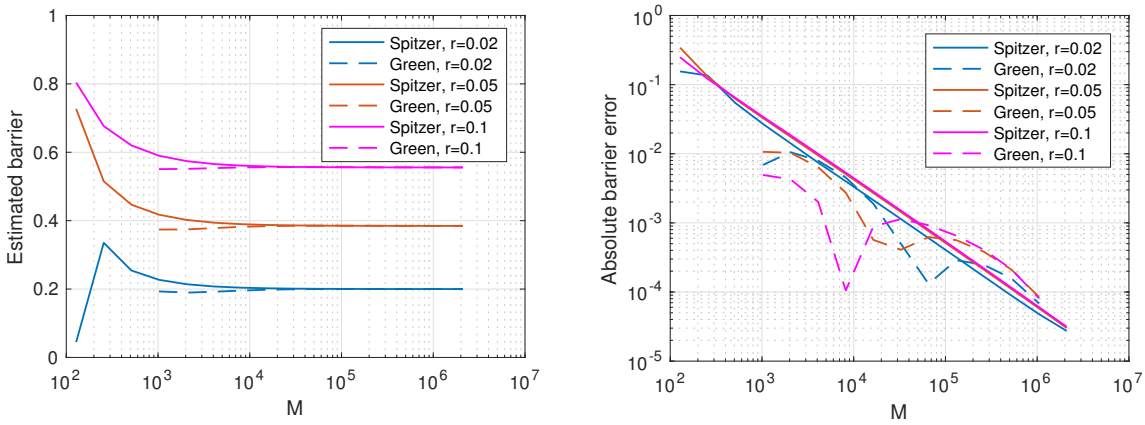

Figure 5.32: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Gaussian process. The error is calculated vs. the closed-form expression by Merton, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

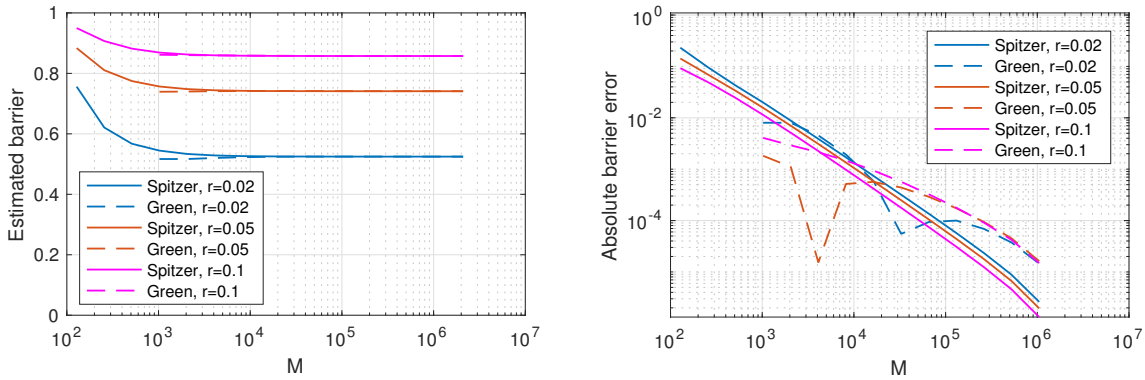

Figure 5.33: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a VG process. The error is calculated vs. the numerical result with the maximum grid size, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

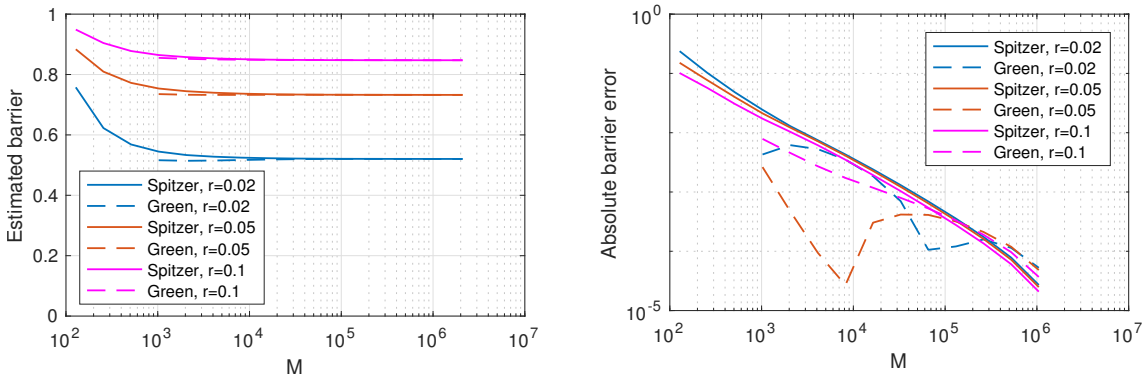

Figure 5.34: Optimal exercise barrier error convergence vs. grid size M with an underlying asset driven by a Merton jump-diffusion process. The error is calculated vs. the numerical result with the maximum grid size, notice that the error convergence for the "Spitzer" method from Section 5.3.2, is superior to the "Green" method from Section 5.3.1.

5.5.1 Comparison between American and Bermudan option prices

The performance for the direct calculation of the price of American options is sufficiently good for practical purposes, with the Spitzer method having an error of 10^{-6} for a CPU time of 10^{-2}s. However, it is of academic interest to study the use of the price for Bermudan options as an approximation to the price for American options, especially as the Bermudan option price converges faster with M for errors $\leq 10^{-8}$.

The left-hand plot in Figure 5.35 shows the price error compared to American options plotted against Δt with an underlying asset driven by a Gaussian process and it is clear that the relationship is linear. By extrapolating this line, we can see that in order to achieve an error of 10^{-8} for $r = 0.02$, a step size of Δt =0.3E-06 is required. The price convergence with this step size is shown in Figure 5.35 and we can see that reducing the step size this low destroys the monotonicity of the convergence and the excellent error performance that we were seeing for more realistic step sizes.

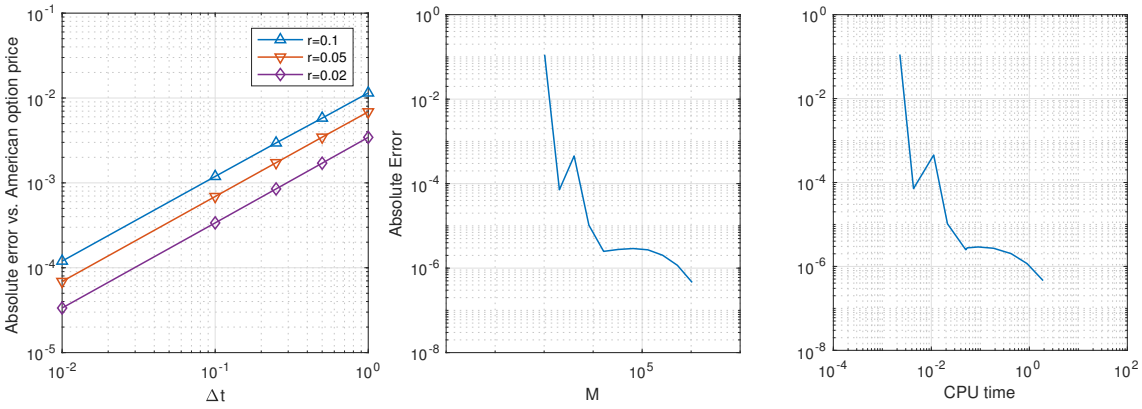

Figure 5.35: Convergence of the price of perpetual Bermudan options to the price of perpetual American options. The left-hand plot shows the convergence as $\Delta t \to 0$ with an underlying asset driven by a Gaussian process. The centre and right-hand side plot shows the error convergence of the price of a perpetual Bermudan option with $r = 0.02$ and Δt =3E-06 with number of grid points M and CPU time.

5.5.2 Factorisation error for continuous monitoring

In Section 4.3 of Chapter 4 we discuss the factorisation error of $\Phi_{\mathrm{c}}(\xi, s)$ in detail and show that we have sub first order polynomial error caused by the increasing nature of $\log \Phi_{\mathrm{c}}(\xi, s)$. However, whilst we repeatedly observe this shape of convergence experimentally for continuously monitored options, the absolute error varies significantly depending on the type of option valued. For barrier options, the level of the error convergence meant that the continuous method performed better than the method with discrete monitoring with very small step sizes. In contrast, the very high level of absolute error for continuously monitored α-quantile options means that it is preferable to use the discretely monitored method with small Δt as an approximation. A further example is provided by the new Spitzer based method method for perpetual American options where the absolute error for the factorisation is so low that it does not dominate the error convergence of the method for errors greater than 10^{-7}. The observation that the error convergence of Green's method for perpetual American options is far inferior to the Spitzer method shows that the absolute error is not dependent on the type of contracts under valuation, but rather the pricing method. The bound for the factorisation error derived in Chapter 4 is

$$\left| \frac{\Phi_{\mathrm{c}\pm}(\xi, s) - \Phi_{M,\mathrm{c}\pm}(\xi, s)}{\Phi_{\mathrm{c}\pm}(\xi, s)} \right| < \left| 1 - (eM'\Delta\xi)^{\frac{\kappa}{M'\Delta\xi}} \right|, \tag{5.66}$$

where $\Phi_{\mathrm{c}\pm}(\xi, s)$ is the true value of the Wiener-Hopf factors and $\Phi_{M,\mathrm{c}\pm}(\xi, s)$ is the value calculated numerically with M FFT grid points, κ is some constant. In the derivation of this bound, many variables are absorbed into κ which depend on the exact form of $\Phi_{\mathrm{c}}(\xi, s) = s - \psi(\xi)$; this in turn depends on the values of s and ξ being used and the function for the characteristic exponent $\psi(\xi)$. Furthermore, we observe that the absolute error is

higher for methods when we make use of $\Phi_{c\oplus}(\xi, 0)$ or $\Phi_{c\ominus}(\xi, 0)$ rather than integrating over the entire range of the function. From Eq. (5.66) we see that the truncation errors from the factorisation are proportional to the value of the function under factorisation and the value of $\Phi_{c\oplus}(\xi, s)$ is at a maximum at $\xi = 0$, we would expect the absolute error to be higher.

In future it may be an interesting theoretical exercise to determine the absolute error of κ based on these. However, it has been shown in this thesis that, in confirmation of the theoretical work by Green et al. (2010) and Fusai et al. (2016), the conversion of a Spitzer-based method from discrete to continuous monitoring is easy to implement. Therefore it may be better to confirm the actual level of the factorisation error experimentally rather than find an expression for κ which is, in any case, likely to require some approximation.

5.6 Conclusion

We implemented new pricing methods based on the Spitzer identities for pricing exotic options. The scheme for discretely monitored α-quantile option was based on the method for pricing lookback options by Fusai et al. (2016) but required the addition of filtering due to the inclusion of $t = 0$ as a monitoring date. We saw very fast error convergence with grid size and, by extension, CPU time. Due to the low errors and the presence of the 10^{-11} error floor in the inverse z-transform, it was not possible to identify whether this error was exponentially or high order polynomially convergent. The extremely low errors and computational time make this a highly attractive method as it stands. However, it would be of academic interest to implement this method using an inverse z-transform with a lower error floor in order to determine the exact order of convergence which we achieve.

Similarly to the work for continuously monitored barrier options, we implemented a numerical method for continuously monitored α-quantile options. The shape of the error convergence was very similar to that for continuously monitored barrier options but the absolute error was much higher. This higher absolute error combined with the excellent performance of the discretely monitored method means that that, unlike barrier options, it is preferable to use the discretely monitored method as an approximation to the continuously monitored method.

We implemented two methods for pricing perpetual Bermudan options. One of these was previously designed by Green (2009) but had not yet been implemented numerically. The other was a new method, also based on the Spitzer identities and the sinc-based fast Hilbert transform. Both methods performed well with high order polynomial error convergence of at least $O(1/M^2)$, but our new method showed significantly lower errors and CPU times, with a computational speed $10\times$ faster for an error of 10^{-7}.

These methods were extended to perpetual American options (i.e. with continuous monitoring) and very different results were observed for the two methods. For Green's

method, the factorisation error was very high and dominated the error convergence. However, for the new Spitzer based method, the factorisation error was much lower and therefore the effect only became dominant for errors below approximately 10^{-7}.

For errors greater than 10^{-7}, the new Spitzer-based method for perpetual American options has errors at least as low as the method for perpetual Bermudan options so we conclude that, for practical purposes, there is no advantage in using discrete monitoring as an approximation for continuous monitoring. Comparing the two as an academic exercise showed that there is as a linear relationship between the size of the time step Δt with discrete monitoring and the error compared to continuous monitoring. However, reducing Δt low enough that the discrete method would be predicted to have a lower error than the continuous method significantly degraded the error convergence such that there was no advantage gained from using the discrete method.

Chapter 6

Numerical solution of the Wiener-Hopf and Fredholm equations

In Chapter 2 we described how the formulation of the Spitzer identities by Green et al. (2010) require the solution of a Wiener-Hopf type equation. The work described in Chapters 3–4 developed an accurate numerical method for solving these equations and produced new bounds for its error performance under different conditions. In this chapter we examine whether the use of these methods can be expanded to the general case and describe the error performance.

6.1 Test cases

As we present a general solution to the Fredholm equation, rather than one limited to a particular application we provide several test cases for solving

$$\lambda f(x) - \int_a^b k(x - x')f(x')dx' = g(x), \quad x \in (a, b), \tag{6.1}$$

for $f(x)$. Solutions to Eq. (6.1) with simple closed-form expressions for $f(x)$, $g(x)$ and $k(x)$ are not readily available. However, if we limit the requirement for simplicity to $f(x)$ and $k(x)$, closed-form expressions for $g(x)$ in Eq. (6.1) can be calculated. Then $g(x)$ and $k(x)$ can be used as inputs to our numerical method and the accuracy of the method can be measured by comparing the result with $f(x)$. We selected $k(x)$ and $f(x)$ to have closed-form expressions for $g(x)$ and also to have Fourier transforms which are easily calculable. Three possible solutions have been derived, $k(x)$ and $f(x)$ both double exponential, $k(x)$ and $f(x)$ both Gaussian, and $k(x)$ and $f(x)$ both Cauchy. The derivation of $g(x)$ is described in the following sections.

6.1.1 Cauchy

We use $f(x) = k(x) = \frac{1}{\pi(x^2+1)}$. The first step in the calculation of $g(x)$ is to find the solution to the integral in Eq. (6.2), i.e.

$$g_{\text{int}}(x) = \frac{1}{\pi^2}\int_a^b \frac{1}{y^2+1}\frac{1}{(x-y)^2+1}dy. \tag{6.2}$$

Using partial fractions and integrating:

$$\begin{aligned}
g_{\text{int}}(x) =& \frac{1}{\pi^2}\int_a^b \frac{1}{y^2+1}\frac{1}{(x-y)^2+1}dy \\
=& \frac{1}{\pi^2 x(x^2+4)}\int_a^b \frac{2y}{y^2+1} + \frac{x}{y^2+1} - \frac{2(y-x)}{(y-x)^2+1} + \frac{x}{(y-x)^2+1}dy \\
=& \frac{1}{\pi^2 x(x^2+4)}\left[\log(y^2+1) + x\arctan(y) - \log[(x-y)^2+1] + x\arctan(y-x)\right]_a^b \\
=& \frac{1}{\pi^2 x(x^2+4)}\left\{\log\left[\frac{(b^2+1)((a-x)^2+1)}{(a^2+1)((b-x)^2+1)}\right] + \right. \\
& \left. + x\left[\arctan(b) - \arctan(a) + \arctan(b-x) - \arctan(a-x)\right]\right\}.
\end{aligned} \tag{6.3}$$

This gives $g(x)$ in closed-form:

$$\begin{aligned}
g(x) = \frac{1}{\pi(x^2+1)} - & \frac{1}{\pi^2 x(x^2+4)}\left\{\log\left[\frac{(b^2+1)((a-x)^2+1)}{(a^2+1)((b-x)^2+1)}\right] + \right. \\
& \left. + x\left[\arctan(b) - \arctan(a) + \arctan(b-x) - \arctan(a-x)\right]\right\}, \quad x \in [a,b].
\end{aligned} \tag{6.4}$$

6.1.2 Exponential

We use $f(x) = \frac{1}{2}e^{-|x|}$ and $k(x) = \frac{1}{2}e^{-|x|}$. In order to make the calculations of $g(x)$ simpler, the values of a and b have been restricted so that $0 < a < b < \infty$. Then the formula for $g(x)$ in closed-form is

$$\begin{aligned}
g(x) &= \frac{1}{2}e^{-x} - \frac{1}{4}\int_a^b e^{-|x-y|}e^{-y}dy \\
&= \frac{1}{2}e^{-x} - \frac{1}{4}\left[\int_x^b e^{(x-y)}e^{-y}dy + \int_a^x e^{-(x-y)}e^{-y}dy\right] \\
&= \frac{1}{2}e^{-x} - \frac{1}{4}e^{x}\left[\int_x^b e^{-2y}dy + e^{-x}\int_a^x dy\right] \\
&= \frac{1}{2}e^{-x} + \frac{1}{8}e^{x}\left[e^{-2y}\right]_x^b - \frac{1}{4}e^{-x}\left[y\right]_a^x \\
&= \frac{1}{2}e^{-x} + \frac{1}{8}e^{x}\left(e^{-2b} - e^{-2x}\right) - \frac{1}{4}e^{-x}(x-a) \\
&= e^{-x}\left[\frac{3}{8} + \frac{1}{8}e^{-2(b-x)} + \frac{1}{4}(a-x)\right] \quad x \in [a,b].
\end{aligned} \tag{6.5}$$

6.1.3 Gaussian

We use $f(x) = \frac{1}{\sqrt{\pi}} e^{-x^2}$ and $k(x) = \frac{1}{\sqrt{\pi}} e^{-x^2}$. The direct expression for $g(x)$ is then

$$\begin{aligned} g(x) &= \frac{1}{\sqrt{\pi}} e^{-x^2} - \frac{1}{\pi} \int_a^b e^{-(x-y)^2} e^{-y^2} dy \\ &= \frac{1}{\sqrt{\pi}} e^{-x^2} - \frac{1}{\pi} e^{-\frac{x^2}{2}} \int_a^b e^{-2(y-\frac{x}{2})^2} dy \\ &= \frac{1}{\sqrt{\pi}} e^{-x^2} - \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \left\{ \Phi\left[2\left(b - \frac{x}{2}\right)\right] - \Phi\left[2\left(a - \frac{x}{2}\right)\right] \right\}, \quad x \in [a, b], \end{aligned} \tag{6.6}$$

where $\Phi[\cdot]$ is the standard normal CDF.

6.2 Results

The following methods were used to recover $f(x)$ and produce the detailed results shown in this section:

- 4th order quadrature (Press et al., 2007, Eq. (4.1.12)).
- Wiener-Hopf method using sinc-based Hilbert transforms. In order to counteract the oscillations on the recovered function, we used an exponential filter of order 12 on the fixed-point algorithm and one of order 4 on the final stage. No zero padding is used. We discuss the use of the sinc-based fast Hilbert transform and spectral filtering in Section 6.2.1 below
- Wiener-Hopf method using the symmetrical sign function for the Hilbert transform, i.e. with zeros placed at both $\xi = 0$ and $\xi = -\frac{M}{2\Delta\xi}$, similar to the methods described by Rino (1970) and Henery (1974), as referenced by Fusai et al. (2016).

It is common in the literature on numerical methods to refer to the number of grid points as n or N, however to maintain consistency with the option pricing methods described in Chapters 3–5, where N is used for the number of monitoring dates, we use M.

6.2.1 Sinc-based fast Hilbert transform and spectral filtering

In the financial pricing applications described in Chapters 3–5, and indeed in the existing literature by Feng and Linetsky (2008) and Fusai et al. (2016) the sinc-based fast Hilbert transform has shown excellent error convergence, especially when combined with spectral filtering techniques as in Chapter 3. However when we consider its use for this application we must take account of several ways in which the requirements differ from its general use for finding solutions to Wiener-Hopf or Fredholm equations.

Firstly the pricing methods that we implement, as devised by Green et al. (2010), use the analytic continuation of x, i.e. they give results for values of x both inside and outside

the barriers (the integration limits of Fredholm equation). This means that there is no requirement to truncate the functions to the integration limits in the state space. Therefore the inputs to the Hilbert transform are exponentially convergent (or polynomial in the case of the VG process) and so excellent error performance is achieved, especially using spectral filtering to solve the issue with the fixed-point algorithm. By contrast, here we solve the Fredholm and Wiener-Hopf equations as they were originally formulated, i.e. the function is only defined for the range of the integration (a, b) and therefore the functions $k(x)$ and $g(x)$ must be truncated to the ranges of $(a - b, b - a)$ and (a, b) respectively. This truncation will introduce a jump in the functions which means that their Fourier transforms now have first order polynomial decay. Therefore the truncation error from the Hilbert transform will have a first order polynomial convergence unless we can exploit some symmetry between the of the Fourier domain functions for positive and negative ξ as in Chapter 4, in which case we may achieve second order polynomial convergence.

Moreover, there is a second important distinction to be made between the general solution presented here and the work in the previous chapters. For financial applications, the solutions to the Fredholm equation are used to calculate the expectation of a further function, in this case the payoff function. Therefore the exact errors in the function for individual values of x are not particularly important. Rather, the finance literature is concerned with the average error, weighted according to the shape of the payoff function. This also has particular importance when we are considering the use of the sinc-based fast Hilbert transform described in Section 2.3.2 which was instrumental in achieving exponential error convergence with the number of FFT grid points M in Feng and Linetsky (2008); Fusai et al. (2016).

In Figures 6.1 and 6.2, we show results using the sinc-based fast Hilbert transform with no filtering for the Gaussian test case described in Section 6.1.3. It is immediately obvious that, even for high values of M, oscillations are visible in the numerical solution.

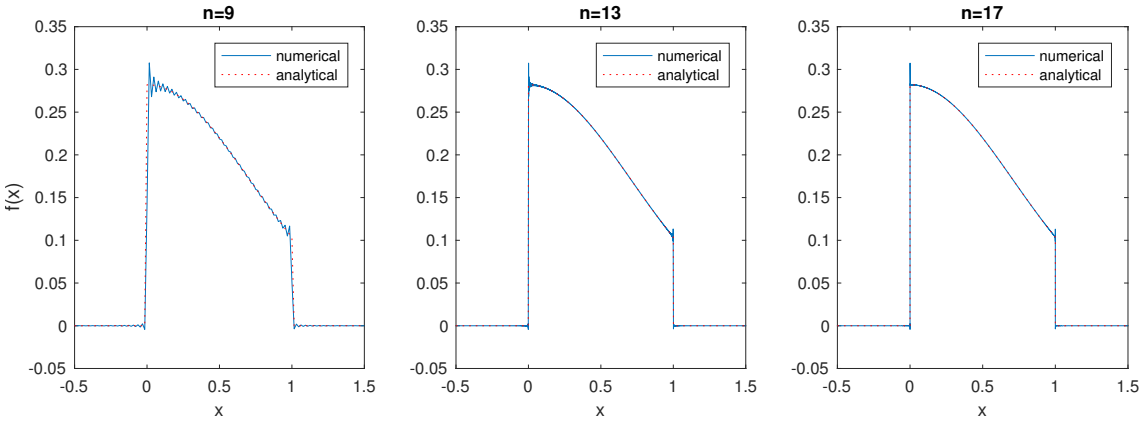

Figure 6.1: Numerical and analytical $f(x)$ using the sinc-based Hilbert transform with no filtering. Notice that oscillations are visible in the numerical solution even for high values of M and that the maximum size of the overshoot is not reduced by increasing this value.

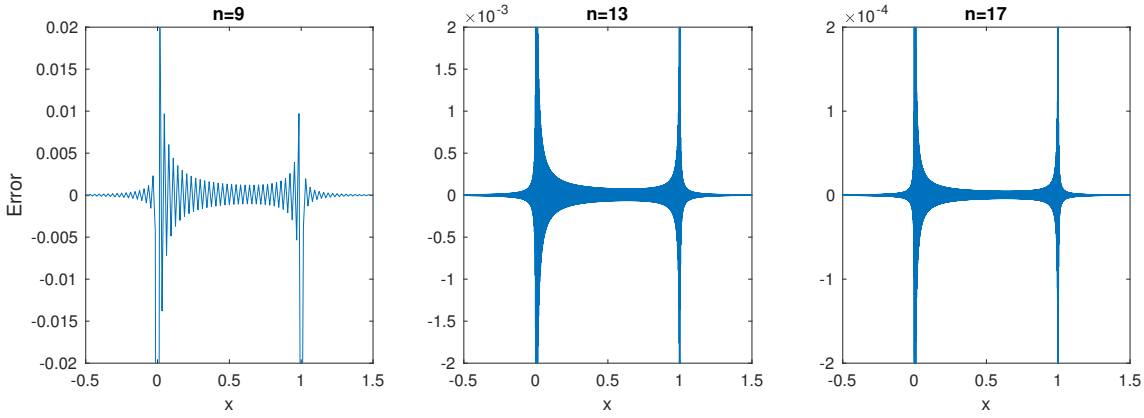

Figure 6.2: Error in the numerical calculation of $f(x)$ using the sinc-based Hilbert transform with no filtering showing the oscillations in more detail. Notice that they increase in frequency as M increases.

We can use spectral filtering to overcome the oscillations, however this can have a negative effect on the accuracy of the numerical method, especially close to the discontinuities in the state space; this is illustrated in Figures 6.3–6.5. Figure 6.4 shows that the lower order filter gives a shallower slope at the discontinuity, but has a stronger effect on the oscillations. However, we can see from Figure 6.4 that, regardless of the order of the filtering, the overshoot at the discontinuity remains approximately the same. Figure 6.5 shows that spectral filtering removes the oscillations away from the discontinuity and that the best results are achieved with a filter of order 8. Although the behaviour of the numerical method using the sinc-based fast Hilbert transform is not appropriate for a general solution to the Fredholm equation due to the high errors at function discontinuities, it remains the case that for applications where we are solely interested in a function value away from any jumps this may be an appropriate method to use.

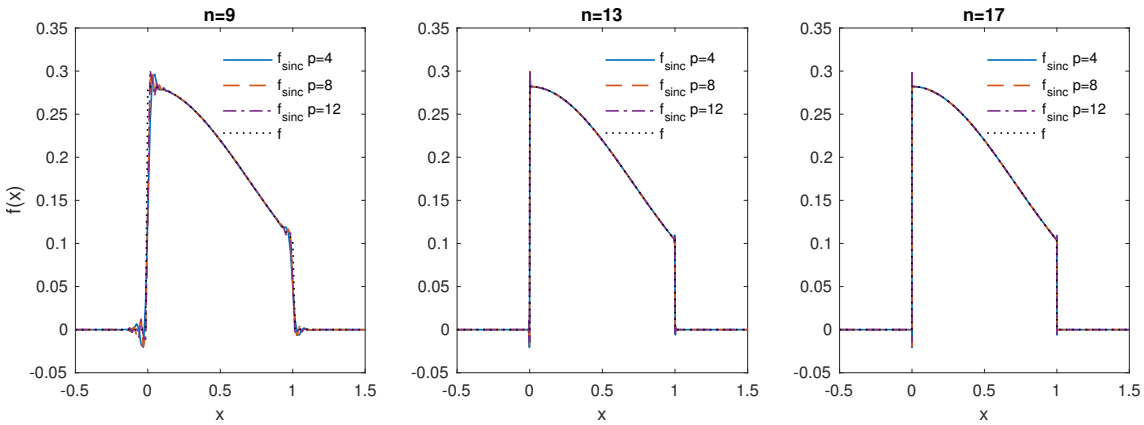

Figure 6.3: Numerical and analytical $f(x)$ using the sinc-based Hilbert transform with exponential filtering. The parameter p describes the order of the filter.

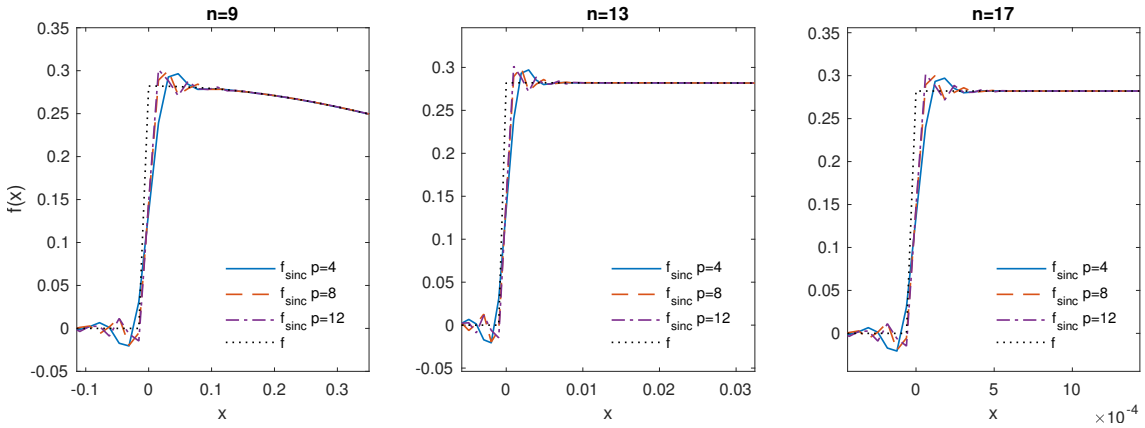

Figure 6.4: Numerical and analytical $f(x)$ using the sinc-based Hilbert transform with exponential filtering, focusing on the discontinuity at $x = 0$. The parameter p describes the order of the filter.

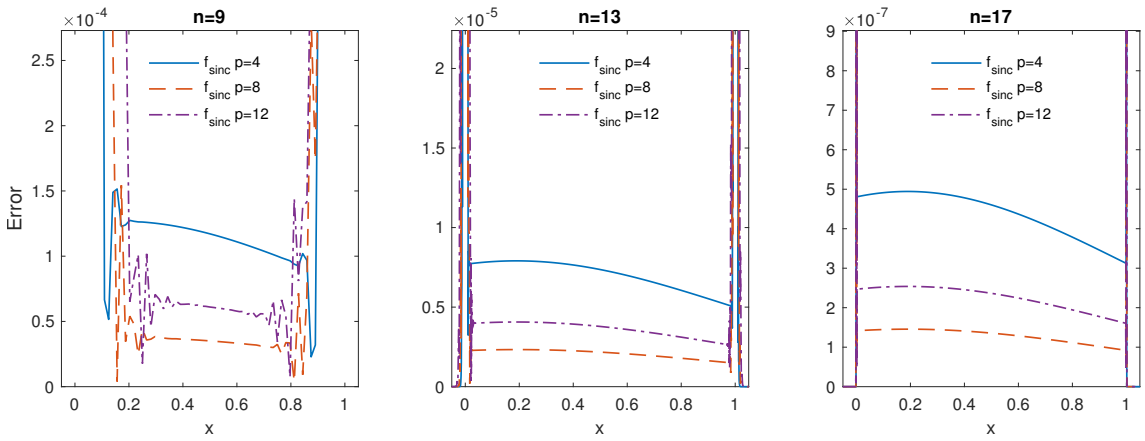

Figure 6.5: Error between the numerical and analytical calculation of $f(x)$ using the sinc-based Hilbert transform with exponential filtering. The parameter p describes the order of the filter. The scale has been chosen to display the error away from the discontinuities of $f(x)$. Notice that the oscillations in the centre of the range are removed but the large error at the discontinuities remain.

6.2.2 FFT method based on the sign function

As an alternative to the sinc-based fast Hilbert transform, we examine the method used by Henery (1974) and Rino (1970) which was also discussed in Fusai et al. (2016). Simply, we realise the relationship between the Hilbert transform and the Fourier transform described in Section 2.1.3, i.e.

$$\mathcal{H}[\widehat{f}(\xi)] = -i\mathcal{F}_{\xi\to x}[\mathrm{sgn}(x)\mathcal{F}^{-1}_{\xi\to x}\widehat{f}(\xi)], \tag{6.7}$$

where

$$\operatorname{sgn}(x) = \begin{cases} 1 & x > 0, \\ 0 & x = 0, \\ -1 & x < 0. \end{cases} \tag{6.8}$$

Notice that Eq. (6.8) is 0 at $x = 0$ and this is is reflected in our numerical implementation.

The results for the Gaussian test case are shown in Figures 6.6–6.8, f is the analytic solution, f_{sin0} is the output from the numerical calculations using the Wiener-Hopf iterative method with the Hilbert transform implemented with the sign function. It is immediately apparent from Figure 6.6 that we no longer suffer from the overshoot that was seen using the sinc-based Hilbert transform. However, looking at the discontinuity more closely in Figure 6.7, we can see that we will have a peak error at a single state-space grid point as the numerical solution increases to the final value of $f(x)$ more slowly than the analytic function. However, unlike the sinc-based function, where the extent of the oscillations depends not only on the filtering but also the shape of the function used, we can state here that as long as the value of x is at least one grid step away from the discontinuity, the answer will be unaffected by the peak error. It is also interesting to note that the error is symmetrical around the discontinuity when the iterative Wiener-Hopf method is used with the sign-based Hilbert transform.

Figure 6.8 displays the error results away from the discontinuity and we can see that, although there is some variation in the error across x the results are all of the same order of magnitude and far superior to the error results seen for the sinc-based Hilbert transform.

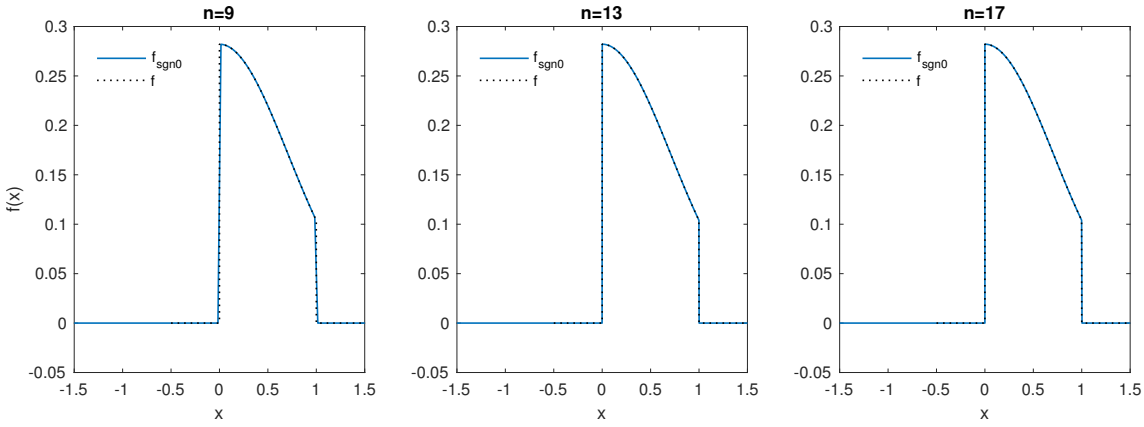

Figure 6.6: Numerical and analytical $f(x)$ using the FFT based method with a symmetrical sign function. Notice that there are no oscillations or overshoot.

Although it is important to observe the functions which are calculated numerically, when assessing the performance of the numerical methods, the error convergence with

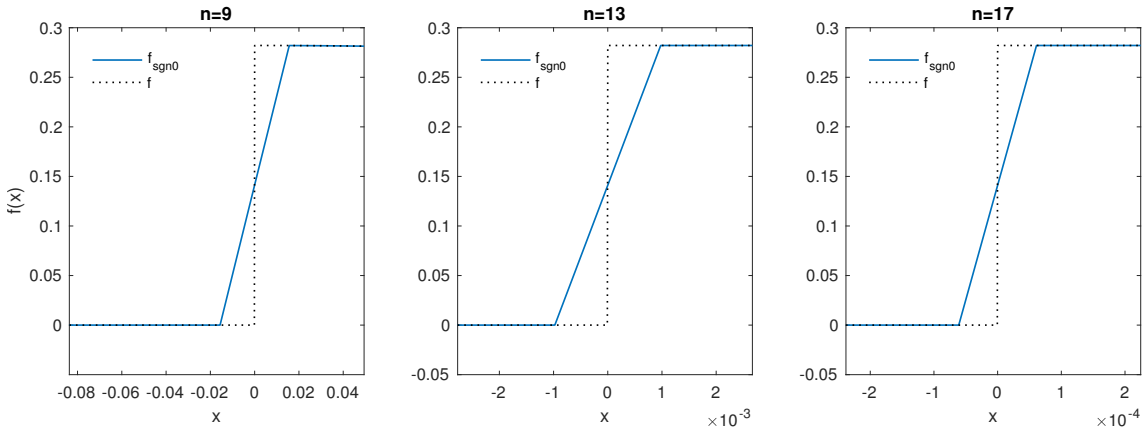

Figure 6.7: Numerical and analytical $f(x)$ using the FFT based method with a symmetrical sign function, focusing on the discontinuity at $x = 0$. Notice that there is an error before and after the discontinuity due to the size of the grid step size.

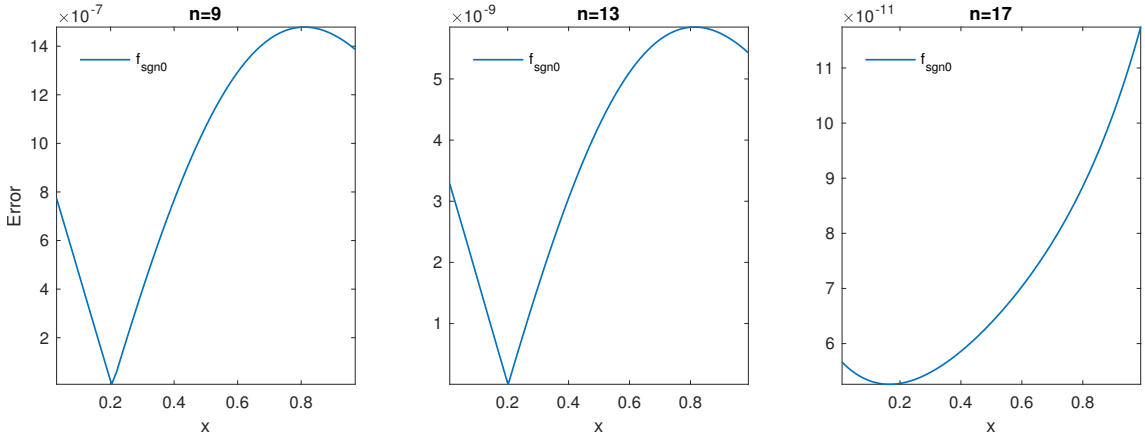

Figure 6.8: Error between the numerical and analytical calculation of $f(x)$ using the FFT based method with a symmetrical sign function. The scale has been chosen to display the error away from the discontinuities of $f(x)$. Notice that the magnitude of the error is far lower than with the sinc-based Hilbert transform.

CPU time and number of grid points M is also important. We measured this at 10%, 50% and 90% of the range between a and b; results for the Gaussian test case are shown in Figures 6.9–6.10. We compare the results for the test cases in Sections 6.1.1–6.1.3 for the iterative Wiener-Hopf method with both the sinc-based and sign-based Hilbert transform methods, in the figures these are labelled f_{sinc} and f_{sin0} respectively. An 8^{th} order exponential function was used with the sinc-based Hilbert transform to counteract the oscillations, as described in Section 6.2.1. We also include results for the existing quadrature based method (labelled f_q), as described in (Press et al., 2007, Eq. (4.1.12)) and which was the previous state of the art, as a benchmark.

The fastest converging method is our new Wiener-Hopf iterative method using the sign-based Hilbert transform, achieving a rate of $O(1/M^2)$. The other methods exhibited

$O(1/M)$ convergence, with the method using the sinc-based Hilbert transform with spectral filtering achieving better absolute error performance vs. M but converging with CPU time almost identically to the quadrature based method. The $O(1/M)$ convergence for the sinc-based Hilbert transform is consistent with the error bound described by Stenger (1993) for a function with a first order discontinuity and the $O(1/M^2)$ convergence seen for the sign-based method is consistent with that reported by Fusai et al. (2016).

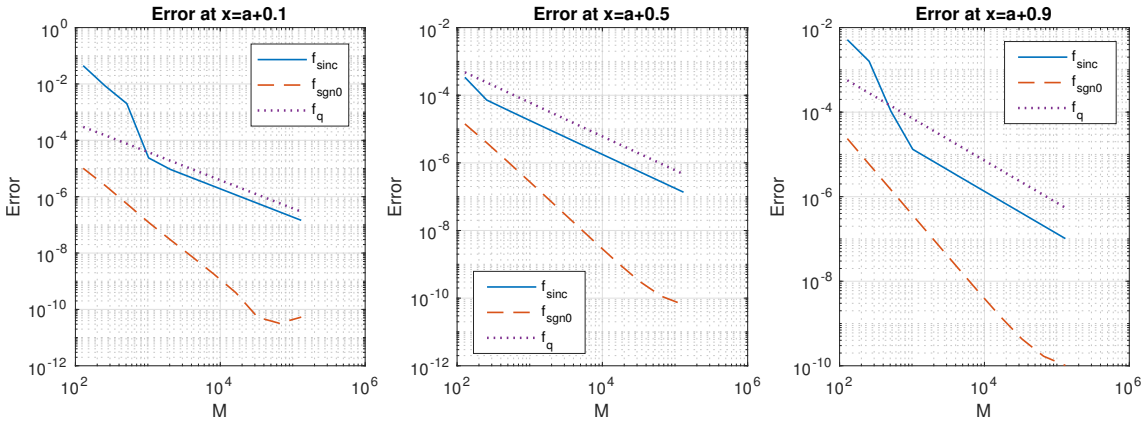

Figure 6.9: Error convergence of the numerical methods vs. M with the Gaussian test case. Notice that our new FFT based method with a symmetrical sign function converges twice as quickly as the existing quadrature based method or our new method with the sinc-based Hilbert transform.

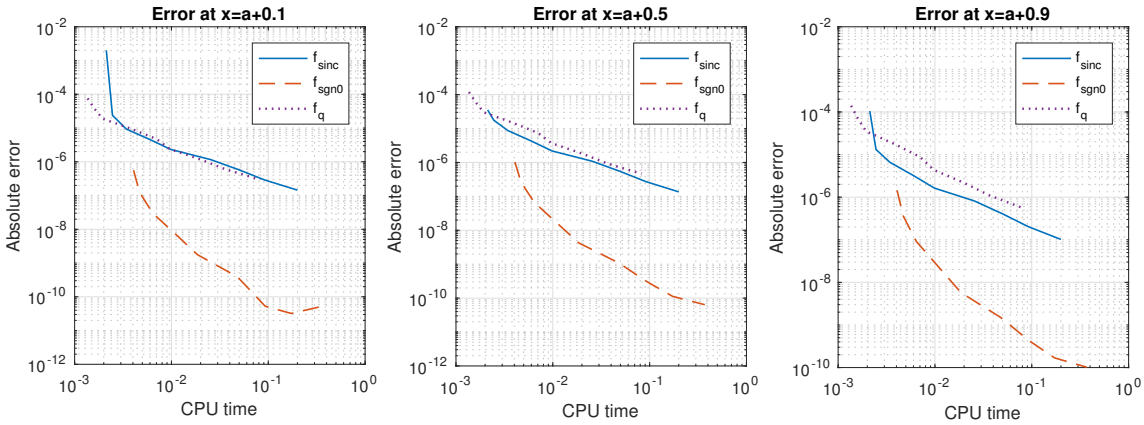

Figure 6.10: Error convergence of the numerical methods vs. CPU time with the Gaussian test case. Notice that the new method with a symmetrical sign function has the best performance.

6.2.3 Results for Cauchy and exponential test cases

Figures 6.11–6.14 show that the results for the Cauchy and exponential test cases are consistent with those for the Gaussian test case; the use of the sinc-based Hilbert transform results in an overshoot at the function discontinuities and the symmetrical sign function

results in a spot error at function discontinuities. We also notice that the quadrature method has a spot error at the discontinuity but this effects a smaller range of x than for our new numerical methods. The reason for this smaller range is that the state space for Fourier-based methods need to be truncated at $\pm \geq 4(b-a)$ in order to avoid wrap-round effects. In contrast, the range of x for the quadrature method need only be truncated at the integration limits a and b.

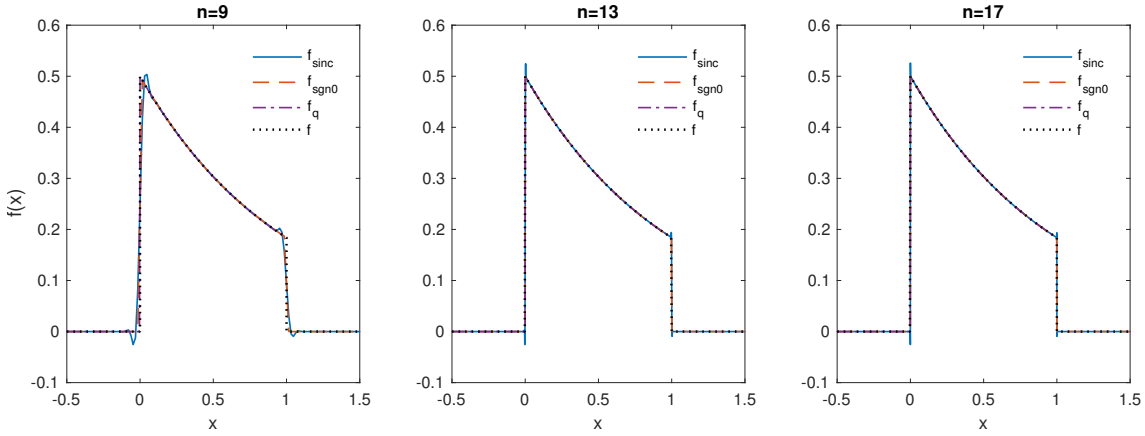

Figure 6.11: Numerical and analytical $f(x)$ with the exponential test case. Notice that the results are consistent with the Gaussian test case.

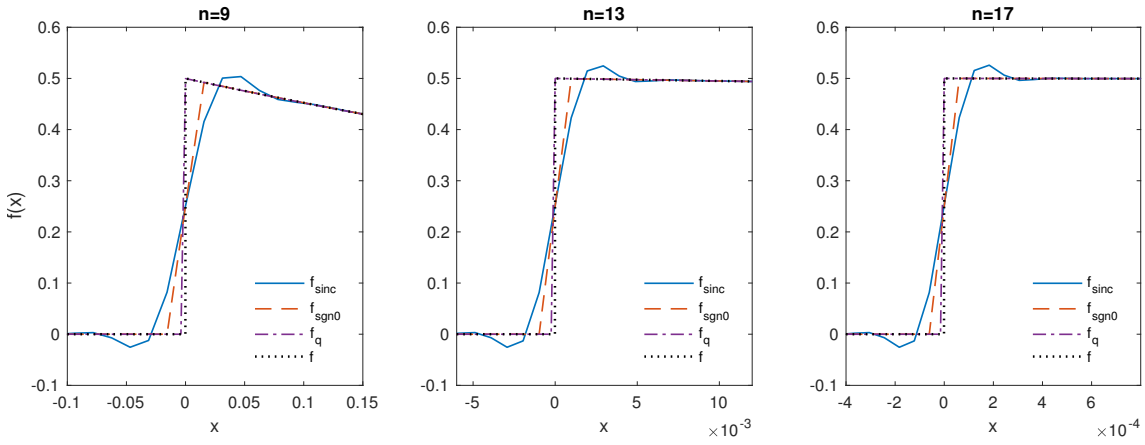

Figure 6.12: Numerical and analytical $f(x)$ with the exponential test case focusing on the first discontinuity. Notice that the results are consistent with the Gaussian test case.

We also measured the error convergence with the Cauchy and exponential test cases and the results are shown in Figures 6.15–6.18. These confirm the findings with the Gaussian test case in Section 6.1.3 which showed that the best performing method is the new iterative solution to the Wiener-Hopf equation with the Hilbert transform implemented using the FFT with the symmetrical sign function.

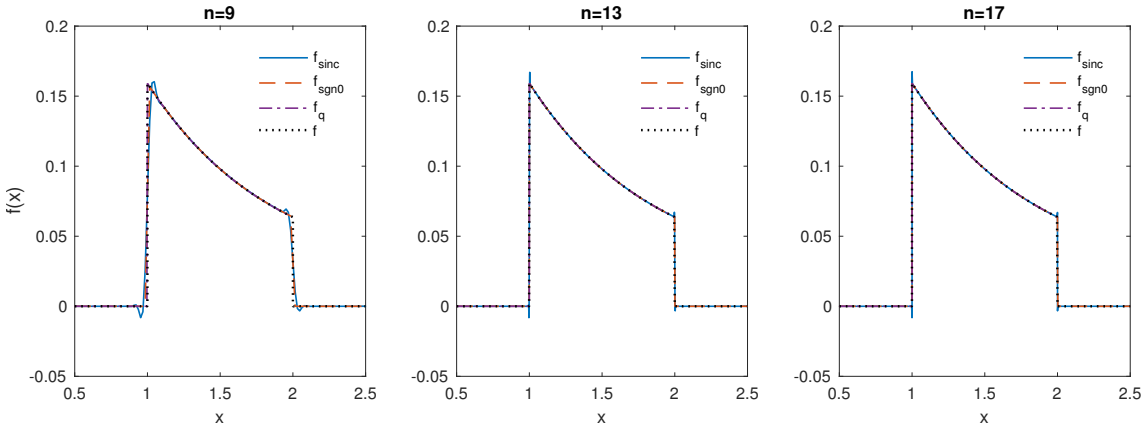

Figure 6.13: Numerical and analytical $f(x)$ with the Cauchy test case. Notice that the results are consistent with the Gaussian and exponential test cases.

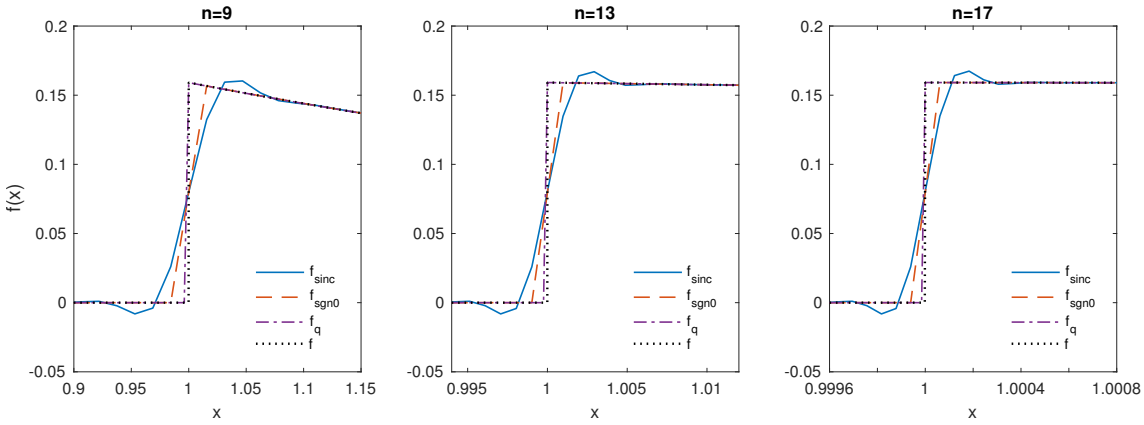

Figure 6.14: Numerical and analytical $f(x)$ with the Cauchy test case focusing on the first discontinuity. Notice that the results are consistent with the Gaussian and exponential test cases.

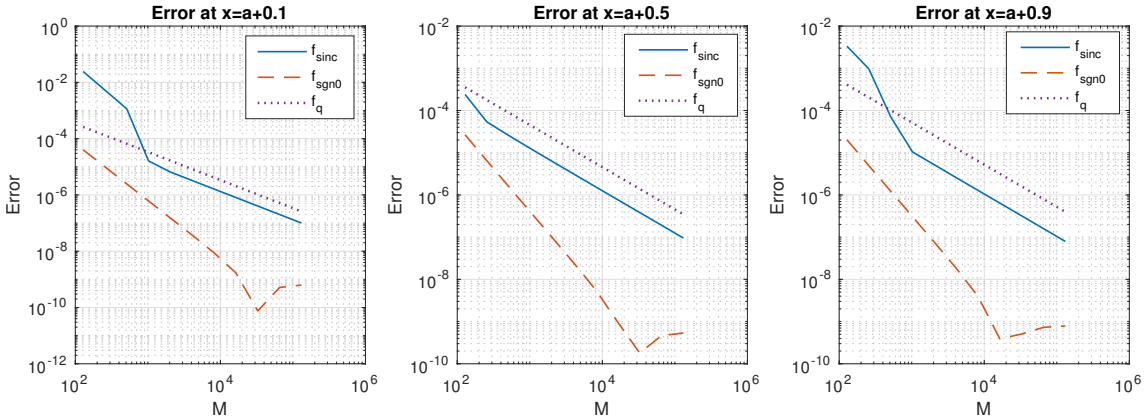

Figure 6.15: Error convergence of the numerical methods vs. M with the Cauchy test case. Notice that the results are consistent with the Gaussian test case.

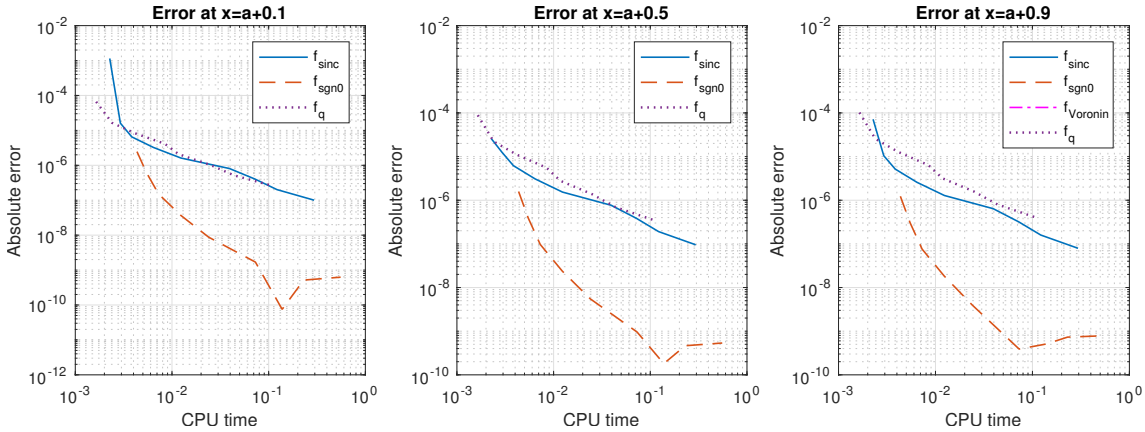

Figure 6.16: Error convergence of the numerical methods vs. CPU time with the Cauchy test case. Notice that the results are consistent with the Gaussian test case.

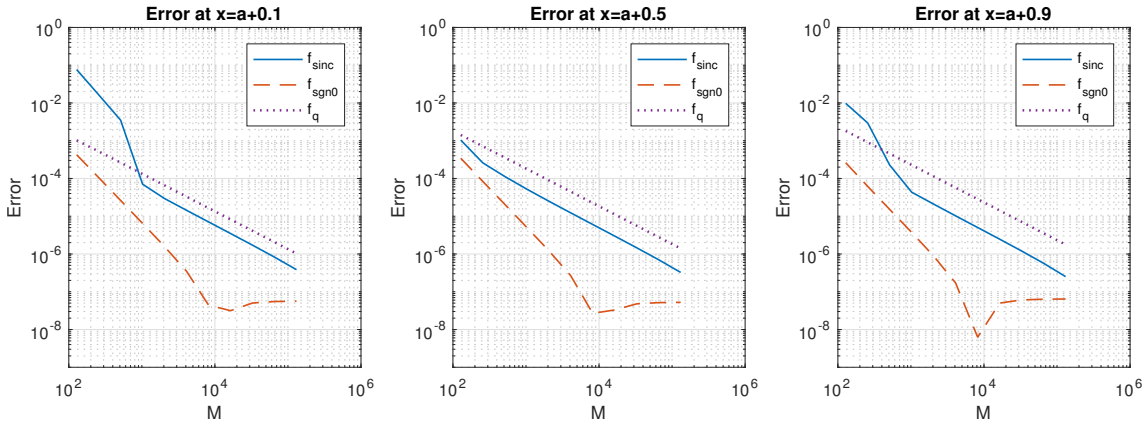

Figure 6.17: Error convergence of the numerical methods vs. M with the exponential test case. Notice that the results are consistent with the Gaussian and Cauchy test cases.

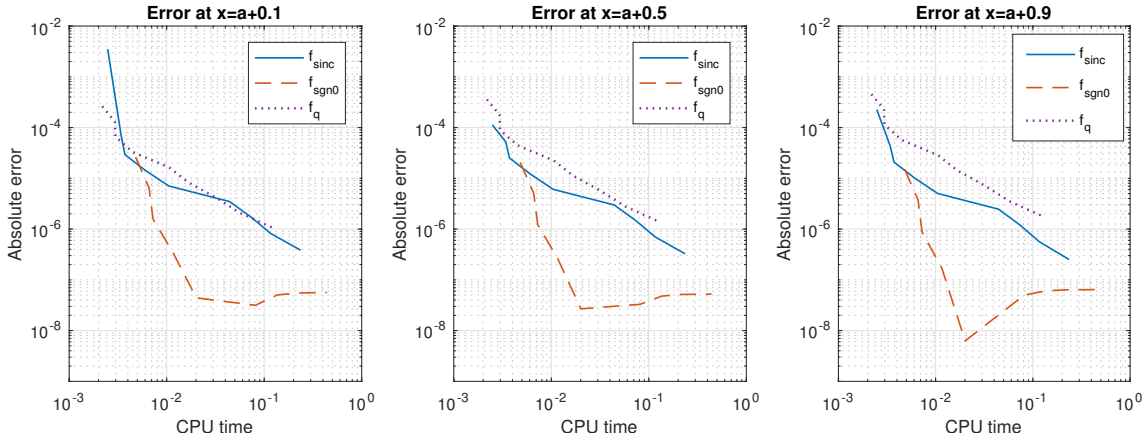

Figure 6.18: Error convergence of the numerical methods vs. CPU time with the exponential test case. Notice that the results are consistent with the Gaussian and Cauchy test cases.

6.3 Conclusion

We implemented three methods for solving general Fredholm equations of the second kind and assessed their performance for three test cases with analytical solutions. These methods comprised of the 4th order quadrature method from (Press et al., 2007, Eq. (4.1.12)) and two iterative methods based on the Wiener-Hopf method with a fixed-point algorithm to solve the coupled equations, as used to obtain the Spitzer identities in Chapters 3–5. The latter two methods differed only in their implementation of the Hilbert transform, with the first using the sinc-based fast Hilbert transform as used so successfully in the application of option pricing and the second using the combination of the FFT and the sign function as described by Rino (1970) and Henery (1974).

The difference in requirements between the application of option pricing and the general solution of the Fredholm equation mean that the method using the sinc-based fast Hilbert transform is not particularly optimal, having $O(1/M)$ convergence and high errors close to the function discontinuity. However, the iterative Wiener-Hopf method using the sign-based Hilbert transform performed much better than both the method with the sinc-based Hilbert transform and the existing quadrature method, having $O(1/M^2)$ convergence and therefore represents a new state of the art in terms of error convergence.

The only other aspect to consider with the Wiener-Hopf method using the sign-FFT Hilbert transform is the width of the peak error at a discontinuity of $f(x)$ as shown in Figures 6.12 and 6.14. The peak error at the discontinuity is wider for the Wiener-Hopf method than for the quadrature method due to the requirement for a wider range of x to overcome the wrap-round effects which we encounter with Fourier-based methods. Therefore if we require an accurate answer very close to a discontinuity the quadrature method could be considered. However, due to the excellent CPU time vs. error performance shown in Figures 6.10, 6.16 and 6.18 we would instead recommend using the Wiener-Hopf method with the sign-based Hilbert transform with a larger grid size to take account of the required accuracy close to the discontinuity

Chapter 7

Conclusions

The work in this thesis concerns the problem of pricing path-dependent options with general Lévy processes. To achieve this, we address outstanding issues for the calculation of fluctuation identities using a numerical implementation of the Wiener-Hopf technique to calculate the Spitzer identities. Despite using the problem of option pricing as a motivating example, the work has a much wider application beyond this. The Spitzer identities have applications in physical, biological, social, actuarial and other sciences and there is also the wider application of the general numerical method for solving the Wiener-Hopf and Fredholm equations.

There is a large body of existing literature concerning the numerical pricing of exotic options, but most particularly we build on the work by Feng and Linetsky (2008), Feng and Linetsky (2008), Green (2009), Green et al. (2010), Fusai et al. (2006) and Fusai et al. (2016). Specifically, the work described in Chapters 3 and 4 addressed outstanding issues for the numerical implementation of the Wiener-Hopf technique to calculate the Spitzer identities using the problem of barrier option pricing as a motivating example. The work in Chapter 5 extended the pricing techniques to other exotic options, namely α-quantile and perpetual Bermudan and American options and for Chapter 6 we looked at applying the numerical techniques to the general problem of solving the Wiener-Hopf and Fredholm equations. The original contributions to the literature are described in more detail in the following paragraphs.

Chapter 3. We investigated the polynomial error performance of the methods for pricing discrete barrier options described by Fusai et al. (2016), and produced new insights into the use of the fast Hilbert transform within an iterative fixed-point algorithm. This, in turn, lead to a rigorous examination of the error convergence of the scheme by Fusai et al. (2016) and the design of a new scheme using spectral filters with exponential error convergence for general Lévy processes. This is the first exponentially convergence pricing scheme for double-barrier options with general Lévy processes which has a CPU time

independent of the number of monitoring dates. Moreover, although spectral filters have been used with general Fourier transform based pricing methods starting with Ruijter et al. (2015), to our knowledge this is the first time spectral filters have been used with the sinc-based fast Hilbert transform from Stenger (1993).

Chapter 4. We adapted the method for discretely monitored barrier options to create, to the best of our knowledge, the first method for the valuation of continuously monitored barrier options which can be used for general Lévy process and is based neither on approximating the characteristic exponent nor approximating the price with continuous monitoring with the price for discrete monitoring. The error performance was investigated and we developed highly accurate error bounds for the pricing procedure which were well matched by empirical results for the error convergence.

Chapter 5. We extended the pricing methods for discrete and continuously monitored barrier options to α-quantile hindsight options and perpetual Bermudan and American options. For discretely monitored α-quantile options, we combined the Spitzer identities and spectral filtering to present the first pricing method with very fast (either exponential or very high order polynomial) convergence and a CPU time independent of the number of monitoring dates. The work on perpetual Bermudan and American options predominantly contributes to the literature in two ways. Firstly we provide an implementation of the residue method described in Green (2009). We also present a completely new method for pricing perpetual options based on the Spitzer identities which also included a novel and efficient way to calculate the optimal exercise barrier. The success of the implementation of Green's method for perpetual Bermudan options is a useful addition to the literature validating, as it does, the elegant mathematics of Green's method. However of more value in practice is the new method we developed based on pricing a put option which not only showed excellent error convergence for perpetual Bermudan options but also retained very good error convergence when the technique was converted to the continuous monitoring case, i.e. perpetual American options.

Chapter 6. The final contribution to the literature is the use of the numerical technique that was developed for the Spitzer identities to solve general Wiener-Hopf and Fredholm equations. We found closed-form solutions to allow the testing of the method and compared the technique under various conditions. Although the results show that general technique does not converge as quickly as the pricing techniques, for reasons which are discussed in more detail in Chapter 7, the new technique is still state of the art, converging with M twice as fast as the existing quadrature method.

The Wiener-Hopf methods which were developed by Green (2009); Green et al. (2010) and Fusai et al. (2016) have been successfully extended to other types of contract and

refined with the use of spectral filtering. Further developments may be possible with local stochastic volatility (LSV) models. The pricing scheme with LSV models would require either the factorisation of a $N \times N$ Wiener-Hopf matrix, where N is the number of points in the volatility grid, or the extension of the fixed-point algorithm. As the current techniques are designed to work with general Lévy processes the techniques can be extended to products with underlying assets characterised by diverse probability distributions. Specifically, it would be useful to apply these techniques to options on crypto-currencies whose probability distributions are very different from the ones used to characterise traditional financial markets. The novelty of the crypto-currency market, especially in terms of options contracts, means that there is a requirement to develop robust pricing methods for these products.

In addition to the extension of the techniques, there are areas where the methods could benefit from small refinements. The Abate and Whitt inverse z-transform works well down to an error floor of 10^{-12} which is, of course, more than sufficient for practical use. However, using an inverse z-transform with a lower noise floor would be useful for the theoretical understanding of the techniques, especially for α-quantile options where it is currently impossible to tell if the convergence is exponential or very high order polynomial. For the continuously monitored methods, we have comprehensively shown that the noise floor of the inverse Laplace transform does not limit the error of the numerical methods which are dominated by the convergence of the factorisation step. However it may be possible to use a Laplace transform which requires fewer calculations and thus reduces the computation time for a given error.

Bibliography

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.

Abrahams, I. D. 1997. On the solution of Wiener–Hopf problems involving noncommutative matrix kernel decompositions. *SIAM Journal on Applied Mathematics* **57**(2) 541–567. doi:10.1137/S0036139995287673.

Akahori, J. 1995. Some formulae for a new type of path-dependent option. *The Annals of Applied Probability* **5**(2) 383–388. doi:10.1214/aoap/1177004769.

Andricopoulos, A. D., M. Widdicks, P. W. Duck, D. P. Newton. 2003. Universal option valuation using quadrature methods. *Journal of Financial Economics* **67**(3) 447 – 471. doi:https://doi.org/10.1016/S0304-405X(02)00257-X.

Askari, H., N. Krichene. 2008. Oil price dynamics (2002–2006). *Energy Economics* **30**(5) 2134–2153. doi:10.1016/j.eneco.2007.12.004.

Atkinson, C., G. Fusai. 2007. Discrete extrema of Brownian motion and pricing of exotic options. *Journal of Computational Finance* **10**(3) 1. doi:10.21314/JCF.2007.174.

Bakshi, G. S., Z. Chen. 1997. An alternative valuation model for contingent claims. *Journal of Financial Economics* **44**(1) 123–165. doi:10.1016/S0304-405X(96)00009-8.

Ballotta, L. 2002. *α-Quantile Option in a Jump-Diffusion Economy*. Springer US, Boston, MA, 75–87. doi:10.1007/978-1-4757-5226-7_5.

Ballotta, L., A. E. Kyprianou. 2001. A note on the α-quantile option. *Applied Mathematical Finance* **8**(3) 137–144. doi:10.1080/13504860210122375.

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

Baxter, G. 1961. An analytic approach to finite fluctuation problems in probability. *Journal d'Analyse Mathématique* **9**(1) 31–70. doi:10.1007/BF02795339.

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. doi:10.2307/1992962.

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.

Beaglehole, D. R., P. H Dybvig, G. Zhou. 1997. Going to extremes: Correcting simulation bias in exotic option valuation. *Financial Analysts Journal* **53**(1) 62–68. doi:10.2469/faj.v53.n1.2057.

Beerends, R. J., H. G. ter Morsche, J. C. van den Berg, E. M. van de Vrie. 2003. *Fourier and Laplace transforms*. Cambridge University Press, Cambridge. doi:10.1017/CBO9780511806834.

Björk, T. 2009. *Arbitrage Theory in Continuous Time*. 3rd ed. Oxford University Press, Oxford.

Black, F., M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* **81**(3) 637–654. doi:10.1086/260062.

Boyarchenko, S. I., S. Z. Levendorskii. 2002a. Perpetual American options under Lévy processes. *SIAM Journal on Control and Optimization* **40**(6) 1663–1696. doi:10.1137/S0363012900373987.

Boyarchenko, S. I., S. Z. Levendorskii. 2002b. Pricing of perpetual Bermudan options. *Quantitative Finance* **2**(6) 432–442. doi:10.1080/14697688.2002.0000010.

Boyd, J. P. 2001. *Chebyshev and Fourier Spectral Methods*. Springer, Heidelberg. doi:10.1002/zamm.19910710715.

Boyle, P. P., Y. Tian. 1998. An explicit finite difference approach to the pricing of barrier options. *Applied Mathematical Finance* **5**(1) 17–43. doi:10.1080/135048698334718.

Bray, A. J., S. N. Majumdar, G. Schehr. 2013. Persistence and first-passage properties in nonequilibrium systems. *Advances in Physics* **62**(3) 225–361. doi:10.1080/00018732.2013.803819.

Brennan, M. J., E. S. Schwartz. 1977. The valuation of American put options. *The Journal of Finance* **32**(2) 449–462. doi:10.1111/j.1540-6261.1977.tb03284.x.

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.

Cai, N., N. Chen, X. Wan. 2010. Occupation times of jump-diffusion processes with double exponential jumps and the pricing of options. *Mathematics of Operations Research* **35**(2) 412–437. doi:10.1287/moor.1100.0447.

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. doi:10.21314/JCF.1999.043.

Chang, C., C.-D. Fuh, S.-K. Lin. 2013. A tale of two regimes: Theory and empirical evidence for a Markov-modulated jump diffusion model of equity returns and derivative pricing implications. *Journal of Banking and Finance* **37**(8) 3204–3217. doi:10.1016/j.jbankfin.2013.03.009.

Chen, L. 1996. *A Three-Factor Model of the Term Structure of Interest Rates*. Springer Berlin Heidelberg, Berlin, 1–36. doi:10.1007/978-3-642-46825-4_1.

Chen, R.-R., L. Scott. 1992. Pricing interest rate options in a two-factor Cox-Ingersoll-Ross model of the term structure. *The Review of Financial Studies* **5**(4) 613–636. doi:10.1093/rfs/5.4.613.

Chen, S.-N., M.-H. Chiang, P.-P. Hsu, C.-Y. Li. 2014. Valuation of quanto options in a Markovian regime-switching market: A Markov-modulated Gaussian HJM model. *Finance Research Letters* **11**(2) 161–172. doi:10.1016/j.frl.2013.09.002.

Chernoff, H. 1961. Sequential tests for the mean of a normal distribution. Jerzy Neyman, ed., *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, Berkeley, CA, 79–91. URL `https://projecteuclid.org/euclid.bsmsp/1200512160`. Volume 1: Contributions to the Theory of Statistics.

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.

Choi, J., D. Margetis, T. M. Squires, M. Z. Bazant. 2005. Steady advection–diffusion around finite absorbers in two-dimensional potential flows. *Journal of Fluid Mechanics* **536** 155–184. doi:10.1017/S0022112005005008.

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 Publishing Company, Amsterdam.

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

Cont, R., P. Takov. 2004. Non-parametric calibration of jump?diffusion option pricing models. *Journal of Computational Finance* **7**(3) 1–49. doi:10.21314/JCF.2004.123.

Cont, R., P. Tankov. 2004. *Financial Modeling with Jump Processes*. Chapman & Hall/CRC, Boca Raton.

Cont, R., P. Tankov. 2006. Retrieving Lévy processes from option prices: Regularization of an ill-posed inverse problem. *SIAM Journal on Control and Optimization* **45**(1) 1–25. doi:10.1137/040616267.

Cox, J. 1975. Notes on option pricing I: Constant elasticity of diffusions.

Cox, J. C., S. A. Ross, M. Rubinstein. 1979. Option pricing: A simplified approach. *Journal of Financial Economics* **7**(3) 229–263. doi:https://doi.org/10.1016/0304-405X(79)90015-1.

Cui, Z., J. L. Kirkby, D. Nguyen. 2017. Equity-linked annuity pricing with cliquet-style guarantees in regime-switching and stochastic volatility models with jumps. *Insurance: Mathematics and Economics* **74** 46–62. doi:10.1016/j.insmatheco.2017.02.010.

Dadachanji, Z. 2015. *FX Barrier Options – A Comprehensive Guide for Industry Quants*. Applied Quantitative Finance, Palgrave Macmillan, London. doi:10.1057/9781137462756.

Daniele, V. G., G. Lombardi. 2006. Wiener-Hopf solution for impenetrable wedges at skew incidence. *IEEE Transactions on Antennas and Propagation* **54**(9) 2472–2485. doi:10.1109/TAP.2006.880723.

Daniele, V. G., R. S. Zich. 2014. *The Wiener-Hopf Method in Electromagnetics*. SciTech Publishing, Stevenage.

Darling, D. A., T. Liggett, H. M. Taylor. 1972. Optimal stopping for partial sums. *The Annals of Mathematical Statistics* **43**(4) 1363–1368. doi:10.1214/aoms/1177692491.

Dassios, A. 1995. The distribution of the quantile of a Brownian motion with drift and the pricing of related path-dependent options. *The Annals of Applied Probability* **5**(2) 389–398. doi:10.1214/aoap/1177004770.

Dassios, A. 2006. Quantiles of Lévy processes and applications in finance URL `http://http://stats.lse.ac.uk/angelos/docs/reviewpaper.pdf`. Working paper.

Dingeç, K. D., W. Hörmann. 2012. A general control variate method for option pricing under Lévy processes. *European Journal of Operational Research* **221**(2) 368–377. doi:10.1016/j.ejor.2012.03.046.

Domínguez, A. 2016. Highlights in the history of the Fourier transform. *IEEE Pulse* **7**(1) 53–61. doi:10.1109/MPUL.2015.2498500.

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.

Fang, F., C. W. Oosterlee. 2011. A Fourier-based valuation method for Bermudan and barrier options under Heston's model. *SIAM Journal on Financial Mathematics* **2**(1) 439–463. doi:10.1137/100794158.

Farnoosh, R., A. Sobhani, M. H. Beheshti. 2017. Efficient and fast numerical method for pricing discrete double barrier option by projection method. *Computers and Mathematics with Applications* **73**(7) 1539–1545. doi:10.1016/j.camwa.2017.01.019.

Feng, L., X. Lin. 2013. Pricing Bermudan options in Lévy process models. *SIAM Journal on Financial Mathematics* **4**(1) 474–493. doi:10.1137/120881063.

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.

Fredholm, I. 1903. Sur une classe d'équations fonctionelles. *Acta mathematica* **27** 365–390. doi:10.1007/BF02421317.

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, 1998.*, 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. doi:10.1007/s00780-005-0170-y.

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:10.1016/j.ejor.2015.11.027.

Fusai, G., D. Marazzina, M. Marena. 2011. Pricing discretely monitored Asian options by maturity randomization. *SIAM Journal on Financial Mathematics* **2**(1) 383–403. doi:10.1137/09076115X.

Fusai, G., D. Marazzina, M. Marena, M. Ng. 2012. z-transform and preconditioning techniques for option pricing. *Quantitative Finance* **12**(9) 1381–1394. doi:10.1080/14697688.2010.538074.

Fusai, G., M. C. Recchioni. 2007. Analysis of quadrature methods for pricing discrete barrier options. *Journal of Economic Dynamics and Control* **31**(3) 826–860. doi:10.1016/j.jedc.2006.03.002.

Fusai, G., A. Tagliani. 2001. Pricing of occupation time derivatives: continuous and discrete monitoring. *Journal of Computational Finance* **5**(1) 1–38. doi:10.21314/JCF.2001.059.

Germano, G., C. E. Phelan, D. Marazzina, G. Fusai. 2018. Solution of Weiner-Hopf and Fredholm integral equations by fast Hilbert and Fourier transforms. Working paper.

Gibbs, J. W. 1898. Fourier's series. *Nature* **59** 200.

Gibbs, J. W. 1899. Fourier's series. *Nature* **59** 606.

Golbabai, A., L. V. Ballestra, D. Ahmadian. 2014. A highly accurate finite element method to price discrete double barrier options. *Computational Economics* **44**(2) 153–173. doi:10.1007/s10614-013-9388-5.

Goldman, M. B., H. B. Sosin, M. A. Gatto. 1979. Path dependent options: "buy at the low, sell at the high". *The Journal of Finance* **34**(5) 1111–1127. doi:10.2307/2327238.

Gottlieb, D., C.-W. Shu. 1995a. On the Gibbs phenomenon IV. recovering exponential accuracy in a subinterval from a Gegenbauer partial sum of a piecewise analytic function. *Mathematics of Computation* **64**(211) 1081–1095. doi:10.1090/S0025-5718-1995-1284667-0.

Gottlieb, D., C.-W. Shu. 1995b. On the Gibbs phenomenon V:recovering exponential accuracy from collocation point values of a piecewise analytic function. *Numerische Mathematik* **71**(4) 511–526. doi:10.1007/s002110050155.

Gottlieb, D., C.-W. Shu. 1996. On the Gibbs phenomenon III: Recovering exponential accuracy in a sub-interval from a spectral partial sum of a piecewise analytic function. *SIAM Journal on Numerical Analysis* **33**(1) 280–290. doi:10.1137/0733015.

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

Gottlieb, D., C.-W. Shu, A. Solomonoff, H. Vandeven. 1992. On the Gibbs phenomenon I: recovering exponential accuracy from the Fourier partial sum of a non-periodic analytic function. *Journal of Computational and Applied Mathematics* **43**(1) 81–98. doi:https://doi.org/10.1016/0377-0427(92)90260-5.

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. doi:10.1016/S0927-0507(05)80169-0.

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. 2009. Application of the Wiener-Hopf technique to derivative pricing. Ph.D. thesis, University of Manchester.

Green, R., D. I. Abrahams, G. Fusai. 2007. Pricing financial claims contingent upon an underlying asset monitored at discrete times. *Journal of Engineering Mathematics* **59**(4) 373–384. doi:10.1007/s10665-007-9176-0.

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.

Hagan, P. S., Deep D. Kumar, A. S. Lesniewski, D. E. Woodward. 2002. Managing smile risk. *Wilmott Magazine* 84–108.

Henery, R. J. 1974. Solution of Wiener-Hopf integral equations using the fast Fourier transform. *Journal of the Institute of Mathematics and its Applications* **13** 89–96. doi:10.1093/imamat/13.1.89.

Henery, R. J. 1977. Solution of Fredholm integral equations with symmetric difference kernels. *Journal of the Institute of Mathematics and its Applications* **19** 29–37. doi:10.1093/imamat/19.1.29.

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.

Hewitt, E., R. E. Hewitt. 1979. The Gibbs-Wilbraham phenomenon: an episode in Fourier analysis. *Archive for History of Exact Sciences* **21**(2) 129–160. doi:10.1007/BF00330404.

Hull, J. C. 2017. *Options, Futures and Other Derivitives*. 10th ed. Pearson, New Jersey.

Ifeachor, E., B. Jervis. 1993. *Digital Signal Processing — A Practical Approach*. Addison Wesley, Wokingham.

Jones, D. S. 1984. Factorization of a Wiener-Hopf matrix. *IMA Journal of Applied Mathematics* **32** 211–220. doi:10.1093/imamat/32.1-3.211.

Jones, D. S. 1991. Wiener-Hopf splitting of a 2×2 matrix. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **434** 419–433. doi:10.1098/rspa.1991.0101.

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. doi:10.1214/aoms/1177703855.

King, F. W. 2009. *Hilbert Transforms*. Cambridge University Press, Cambridge.

Kisil, A. V. 2015. The relationship between a strip Wiener–Hopf problem and a line Riemann–Hilbert problem. *IMA Journal of Applied Mathematics* **80**(5) 1569–1581. doi:10.1093/imamat/hxv007.

Klüppelberg, C., A. E. Kyprianou, R. A. Maller, et al. 2004. Ruin probabilities and overshoots for general Lévy insurance risk processes. *Annals of Applied Probability* **14**(4) 1766–1801. doi:10.1214/105051604000000927.

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.

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.

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

Kunitomo, N., M. Ikeda. 1992. Pricing options with curved boundaries. *Mathematical Finance* **2**(4) 275–298. doi:10.1111/j.1467-9965.1992.tb00033.x.

Lawrie, J. B., I. D. Abrahams. 2007. A brief historical perspective of the Wiener-Hopf technique. *Journal of Engineering Mathematics* **59**(4) 351–358. doi:10.1007/s10665-007-9195-x.

Lewis, A. 2001. A simple option formula for general jump-diffusion and other exponential Lévy processes. doi:10.2139/ssrn.282110. SSRN 282110.

Lian, G., S. Zhu, R. J. Elliott, Z. Cui. 2017. Semi-analytical valuation for discrete barrier options under time-dependent Lévy processes. *Journal of Banking and Finance* **75**(Supplement C) 167–183. doi:10.1016/j.jbankfin.2016.11.012.

Lord, R., F. Fang, F. Bervoets, C. W. Oosterlee. 2008. A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes. *SIAM Journal on Scientific Computing* **30** 1678–1705. doi:10.1137/070683878.

Mack, C. A. 2011. Fifty years of Moore's law. *IEEE Transactions on Semiconductor Manufacturing* **24**(2) 202–207. doi:10.1109/TSM.2010.2096437.

Madan, D. B., E. Seneta. 1990. The variance gamma (V.G.) model for share market returns. *Journal of Business* **63**(4) 511–524. doi:10.1086/296519.

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. doi:10.1007/978-1-4614-3433-7. Springer Proceedings in Mathematics & Statistics, vol. 19.

Martin, J. D. 1991. *Signals and Processes: A Foundation Course*. Pitman, London.

McKechan, D. J. A., C. Robinson, B. S. Sathyaprakash. 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. doi:10.1088/0264-9381/27/8/084020.

Merton, R. C. 1973. Theory of rational option pricing. *The Bell Journal of Economics and Management Science* **4**(1) 141–183. doi:10.2307/3003143.

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.

Mishuris, G., S. Rogosin. 2016. Factorization of a class of matrix-functions with stable partial indices. *Mathematical Methods in the Applied Sciences* **39**(13) 3791–3807. doi:10.1002/mma.3825.

Miura, R. 1992. A note on look-back options based on order statistics. *Hitotsubashi Journal of Commerce and Management* **27**(1) 15–28. doi:10.15057/5748.

Moore, G. E. 1998. Cramming more components onto integrated circuits, reprinted from Electronics, volume 38, number 8, April 19, 1965, pp. 114 ff. *Proceedings of the IEEE* **86**(1) 82–85. doi:10.1109/JPROC.1998.658762.

Mordecki, E. 2002. Optimal stopping and perpetual options for Lévy processes. *Finance and Stochastics* **6**(4) 473–493. doi:10.1007/s007800200070.

Noble, B. 1958. *Methods Based on the Wiener-Hopf Technique for the Solution of Partial Differential Equations*. Pergamon Press, London. Reprinted New York: Chelsea, 1988.

Nolan, J. P. 2018. *Stable Distributions — Models for Heavy Tailed Data*. Birkhäuser, Boston. URL `http://fs2.american.edu/jpnolan/www/stable/stable.html`. In progress, Chapter 1 online.

O'Cinneide, C. A. 1997. Euler summation for Fourier series and Laplace transform inversion. *Stochastic Models* **13** 315–337.

Paley, R. E. A. C., N. Wiener. 1933. Notes on the theory and application of Fourier transforms. iii, iv, v, vi, vii. *Transactions of the American Mathematical Society* (33) 761–791. doi:10.1090/S0002-9947-1933-1501716-9.

Phelan, C., D. Marazzina, G. Fusai, G. Germano. 2018. Hilbert transform, spectral filters and option pricing. *Annals of Operational Research* doi:10.1007/s10479-018-2881-4.

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

Port, S. C. 1963. An elementary probability approach to fluctuation theory. *Journal of Mathematical Analysis and Applications* **6**(1) 109–151. doi:10.1016/0022-247X(63)90097-0.

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.

Press, W. H., S. A. Teukolsky, W. T. Vetterling, B. P. Flannery. 2007. *Numerical Recipes in C++*. 3rd ed. Cambridge University Press, Cambridge.

Ramezani, C. A., Y. Zeng. 2007. Maximum likelihood estimation of the double exponential jump-diffusion process. *Annals of Finance* **3**(4) 487–507. doi:10.1007/s10436-006-0062-y.

Rino, C. L. 1970. Factorization of spectra by discrete Fourier transforms. *IEEE Transactions on Information Theory* **16** 484–485. doi:10.1109/TIT.1970.1054502.

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. doi:10.1214/aoap/1177005065.

Rogers, L. C. G. 2002. Monte Carlo valuation of American options. *Mathematical Finance* **12**(3) 271–286. doi:10.1111/1467-9965.02010.

Ruijter, M. J., 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. doi:10.21314/JCF.2015.306.

Schoutens, W. 2003. *Lévy Processes in Finance*. Wiley, New York.

Scott, L. O. 1997. Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Applications of Fourier inversion methods. *Mathematical Finance* **7**(4) 413–426. doi:10.1111/1467-9965.00039.

Shreve, S. 2004. *Stochastic Calculus for Finance II*. Springer, New York.

Sobhani, A., M. Milev. 2018. A numerical method for pricing discrete double barrier option by Legendre multiwavelet. *Journal of Computational and Applied Mathematics* **328**(Supplement C) 355–364. doi:10.1016/j.cam.2017.07.033.

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.

Tadmor, E. 2007. Filters, mollifiers and the computation of the Gibbs phenomenon. *Acta Numerica* **16** 305–378. doi:10.1017/S0962492906320016.

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

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

Veitch, B. H., I. D. Abrahams. 2007. On the commutative factorization of $n \times n$ matrix Wiener–Hopf kernels with distinct eigenvalues. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 463. 613–639. doi:10.1098/rspa.2006.1780.

Voronin, A. F. 2004. A complete generalization of the Wiener-Hopf method to convolution integral equations with integrable kernel on a finite interval. *Differential Equations* **40**(9) 1259–1267. doi:10.1007/s10625-005-0004-x.

Wang, R. 2014. Conformal mapping between s-plane to z-plane. `http://fourier.eng.hmc.edu/e102/lectures/Z_Transform/node2.html`. URL `http://fourier.eng.hmc.edu/e102/lectures/Z_Transform/node2.html`. Accessed: 09/10/2015.

Wendel, J. G. 1960. Order statistics of partial sums. *The Annals of Mathematical Statistics* **31**(4) 1034–1044. doi:10.1214/aoms/1177705676.

Widder, D. V. 1945. What is the Laplace transform? *The American Mathematical Monthly* **52**(8) 419–425. doi:10.2307/2305640.

Wiener, N., E. Hopf. 1931. Über eine Klasse singulärer Integralgleichungen. *Sitzungsberichte der Preußischen Akademie der Wissenschaften, Mathematisch-Physikalische Klasse* **31** 696–706.

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

Yor, M. 1995. The distribution of Brownian quantiles. *Journal of Applied Probability* **32**(2) 405–416. doi:10.2307/3215296.

Appendix A

Process and contract parameters

We list the process and contract parameters for the numerical experiments carried out for this thesis.

A.1 Discretely monitored barrier options

Table A.1 contains all the parameters of the numerical experiments for discretely monitored barrier options, presented in Chapter 3.

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

Table A.1: 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.

A.2 Continuously monitored barrier options

Table A.2 contains all the parameters of the numerical experiments for continuously monitored barrier options, presented in Chapter 4.

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

Table A.2: 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.

A.3 Alpha quantile and early exercise options

Table A.3 contains the process parameters used for the α-quantile and early exercise options described in Chapter 5.

Model	$\Psi(\xi, t)$	Symbol	Value
Gaussian	$e^{-t\left(i\mu\xi - \frac{\sigma^2\xi^2}{2}\right)}$	μ	0
		σ	0.4
VG	$(1 - i\nu\xi\theta + \nu\sigma^2\xi^2/2)^{-t/\nu}$	θ	$\frac{1}{9}$
		σ	$\frac{1}{3\sqrt{3}}$
		ν	0.25
Merton jump-diffusion	$e^{-t\left(\frac{1}{2}\sigma^2\xi^2 + \lambda(e^{i\alpha_m\xi - \frac{1}{2}\delta^2\xi^2} - 1)\right)}$	σ	0.1
		λ	3
		α_m	-0.05
		δ	0.086

Table A.3: 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.

Appendix B

Results with varying parameters

We provide additional results with varying values of the damping factor α_{d} and the process parameters to demonstrate that the performance of our pricing methods is robust with respect to changes in these values.

B.1 Discretely monitored Barrier options with filtering

These additional results relate to the pricing method for discretely monitored barrier options described in Chapter 4

B.1.1 Variation of the damping parameter α_{d}

The method of selecting the damping parameter α_{d} originates in the paper by Feng and Linetsky (2008), who specify an acceptable range for α_{d} based on the process parameters. In Chapter 3, and in line with Fusai et al. (2016), we selected a value of α_{d} at the centre of this range. For these additional tests we selected a number of values within the calculated accepted range for α_{d}. We can see from Figures B.1 and B.2 that the performance of the method is very robust against the variation of this parameter. It is important to note that the selection of α_{d} is done automatically based on the calculations by Feng and Linetsky (2008) and does not need to be manually tuned in order to obtain a good performance.

B.1.2 Variation of the process paramters

We tested the method with the NIG and Kou processes using four different parameter sets shown in Table B.1. We carried out all tests with 52 monitoring dates as the effect of varying the number of dates is covered in the results in Chaper 3. The results are shown in Figures B.3 and B.4 and demonstrate that the error convergence varies very little between the different parameter sets.

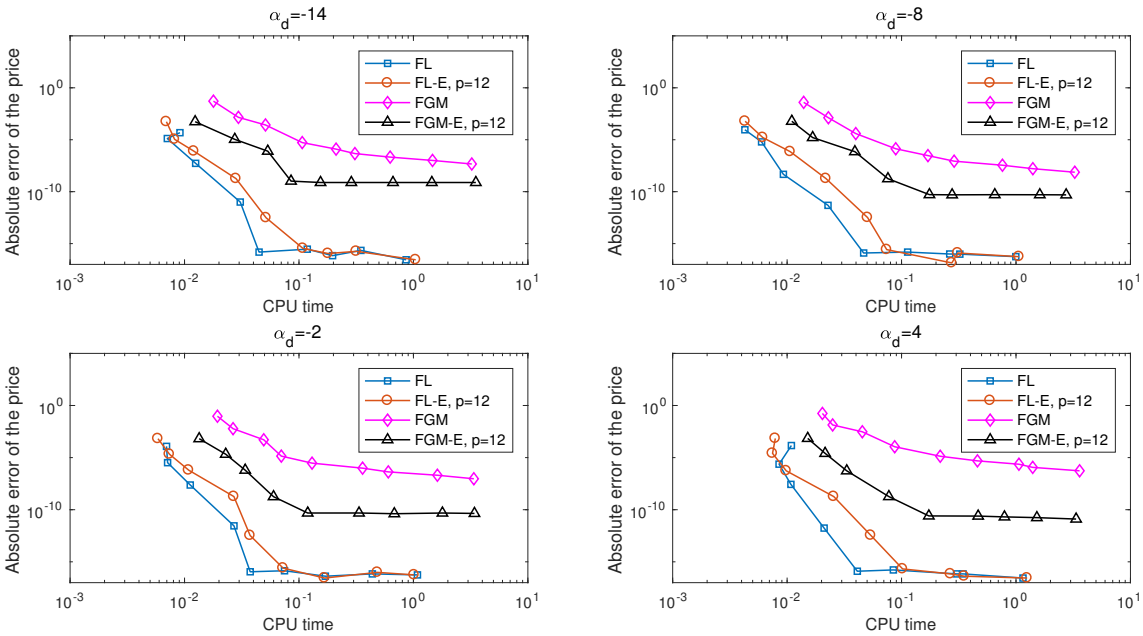

Figure B.1: Results for the pricing error convergence with CPU time for the NIG process with varying values of the damping parameter α_d.

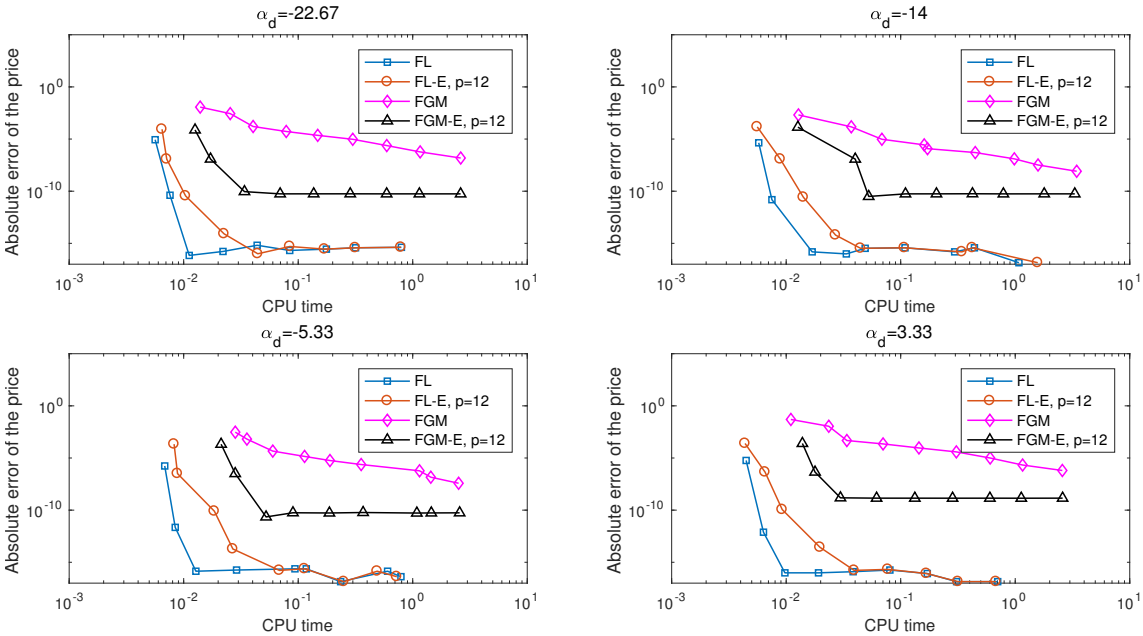

Figure B.2: Results for the pricing error convergence with CPU time for the Kou process with varying values of the damping parameter α_d.

Model	$\Psi(\xi,t)$	Symbol	Set 1	Set 2	Set 3	Set 4
NIG	$e^{-t\delta\left(\sqrt{\alpha^2-(\beta+i\xi)^2}+\sqrt{\alpha^2-\beta^2}\right)}$	α	15	20	18	15
		β	-5	-8	-12	-8
		δ	0.5	0.3	0.4	0.3
Kou	$e^{-t\left(\frac{\sigma^2\xi^2}{2}-\lambda\left(\frac{(1-p)\eta_2}{\eta_2+i\xi}+\frac{p\eta_1}{\eta_1-i\xi}-1\right)\right)}$	p	0.5	0.3	0.3	0.5
		λ	3	3	7	7
		σ	0.2	0.15	0.2	0.15
		η_1	50	25	50	50
		η_2	50	25	50	50

Table B.1: Parameter sets for the underlying processes; $\Psi(\xi,t)$ is the characteristic function of the process that models the log return of the underlying asset.

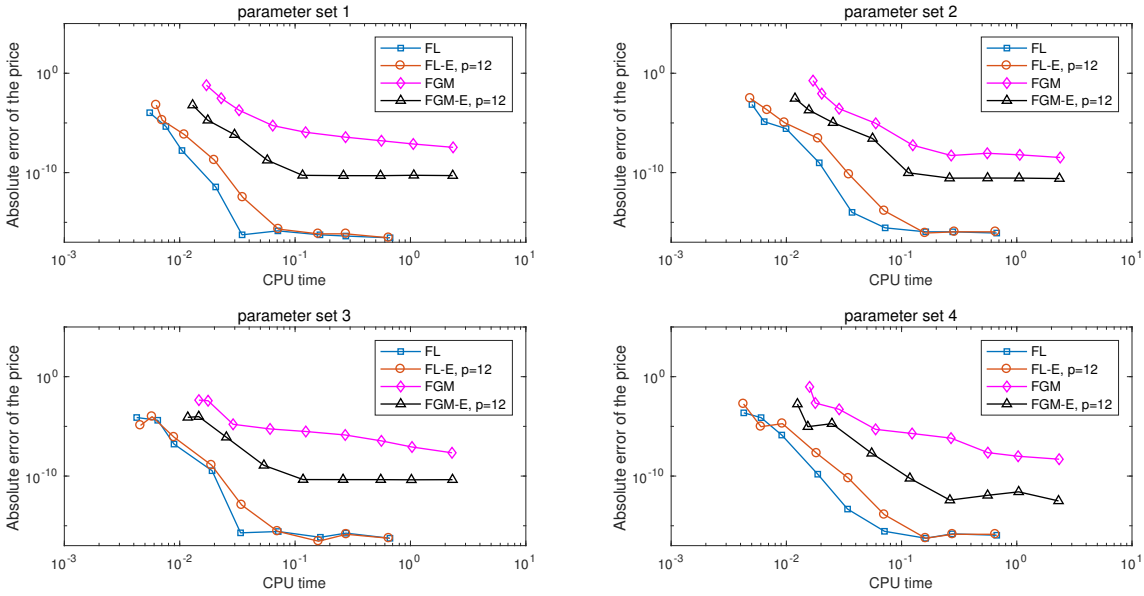

Figure B.3: Results for the pricing error convergence with CPU time for the NIG process with the parameter sets described in Table B.1.

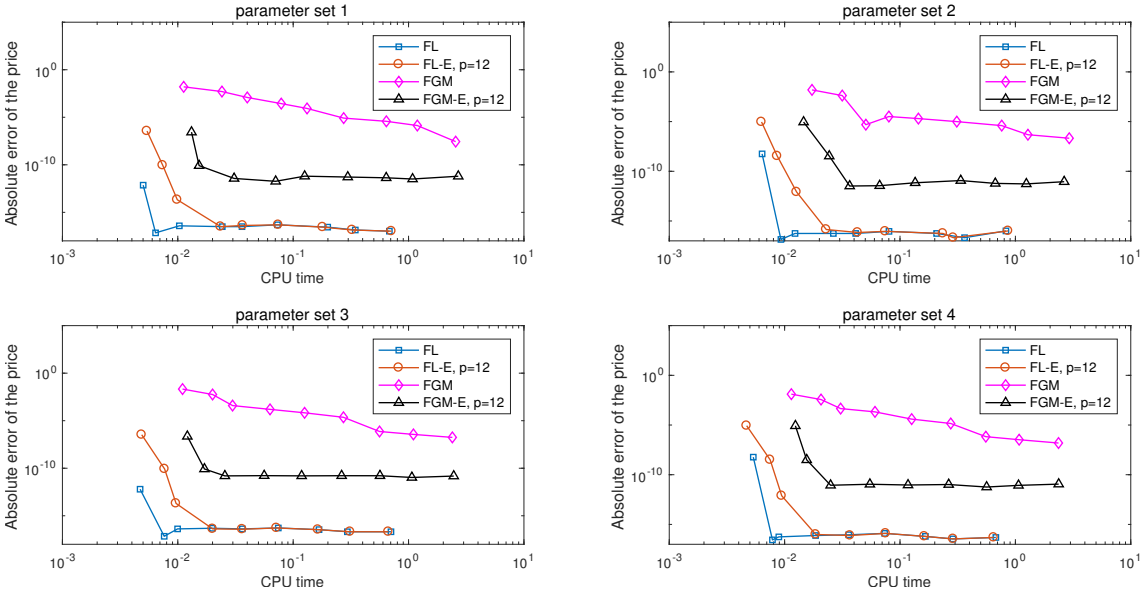

Figure B.4: Results for the pricing error convergence with CPU time for the Kou process with the parameter sets described in Table B.1.

B.2 Continuously monitored Barrier options

These additional results relate to the pricing method for continuously monitored barrier options described in Chapter 4

B.2.1 Variation of the damping parameter α_d on the results for double-barrier options

As for discretely monitored barrier options, we set α_d as the centre of the range specified by Feng and Linetsky (2008). These additional results show that our method is robust with respect to the choice of the value of α_d within a wide range. This can be seen from Figures B.5–B.7, which should be compared with Figure 4.10 in Chapter 4.

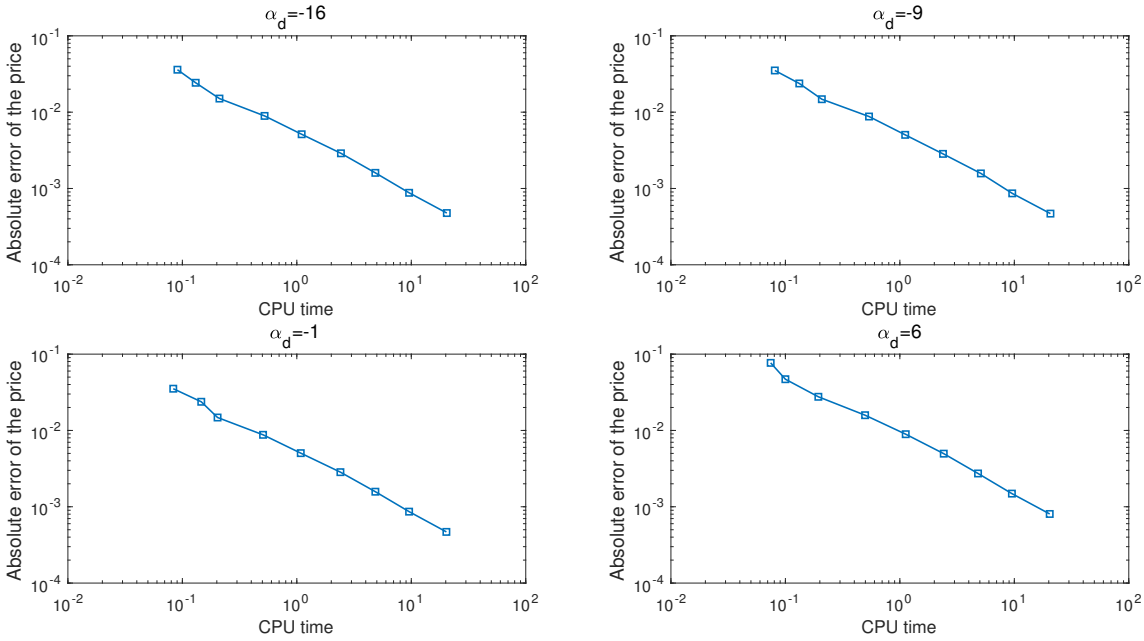

Figure B.5: Results for the error convergence versus CPU time with double barriers and the NIG process varying the value of the damping parameter α_{d}.

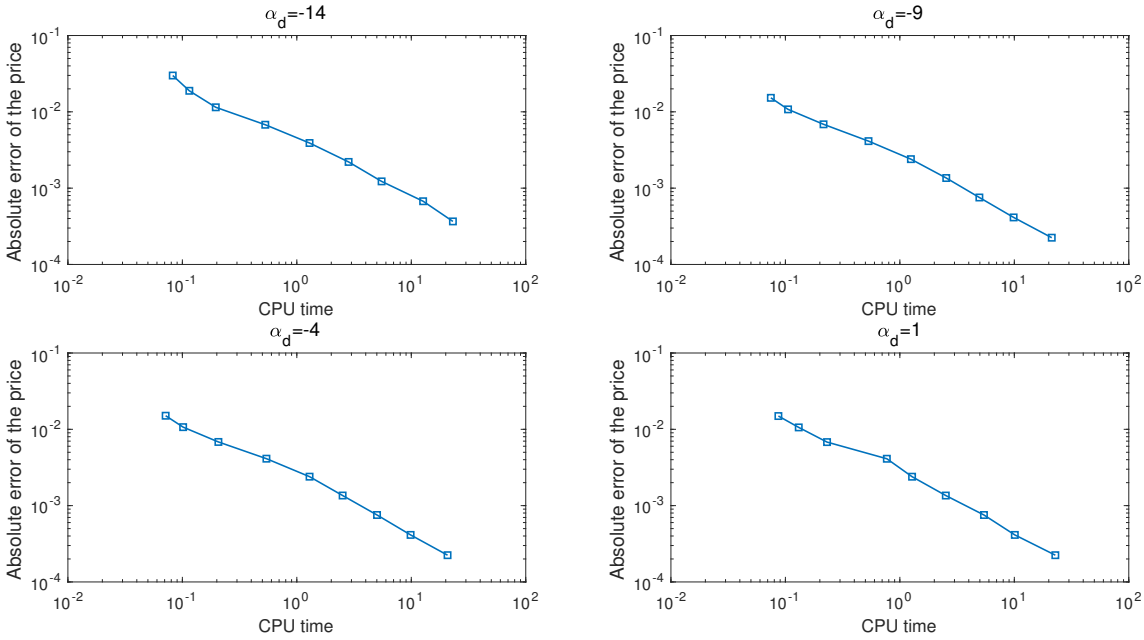

Figure B.6: Results for the error convergence versus CPU time with double barriers and the VG process varying the value of the damping parameter α_{d}.

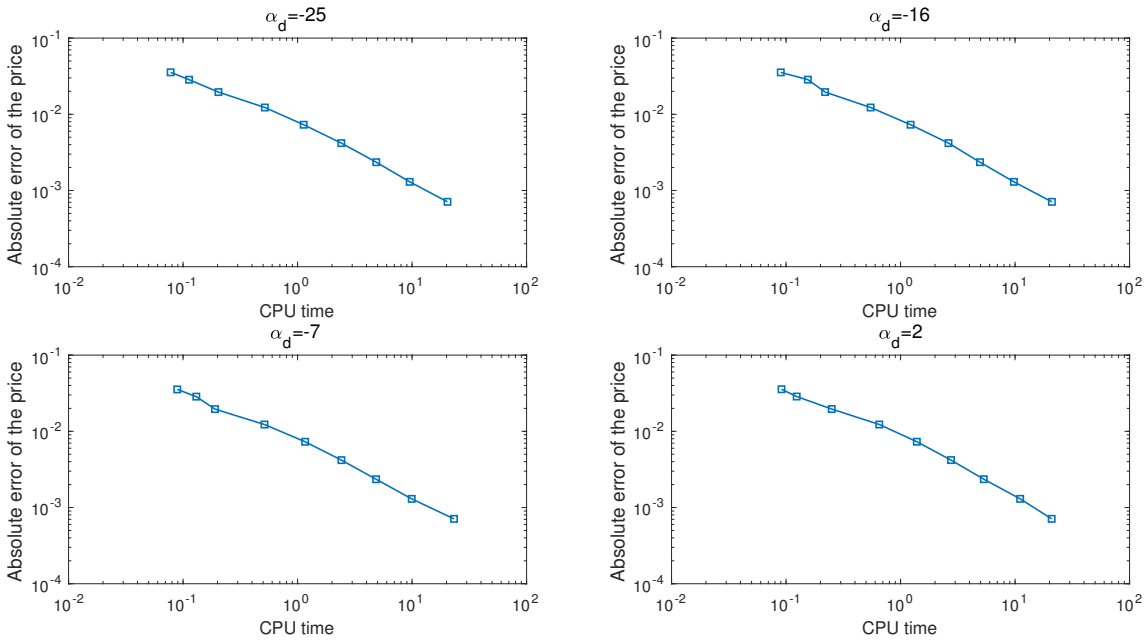

Figure B.7: Results for the error convergence versus CPU time with double barriers and the Kou process varying the value of the damping parameter α_{d}.