HIGHLY EFFICIENT PRICING OF EXOTIC DERIVATIVES UNDER MEAN-REVERSION, JUMPS AND STOCHASTIC VOLATILITY

Chun-Sung Huang

Submitted in fulfillment of the academic
requirements for the degree of
Doctor of Philosophy
in
Computational Finance
in the
School of Management Studies
University of Cape Town

July 17, 2018

Disclaimer

This document describes work undertaken for a PhD programme of study at the University of Cape Town (UCT). All views and opinions expressed therein remain the sole responsibility of the author, and do not necessarily represent those of the institute.

School of Management Studies
University of Cape Town

Declaration 1 - Plagiarism

I, Chun-Sung Huang, declare that

(1) The research reported in this thesis, except where otherwise indicated, is my original research.

(2) This thesis has not been submitted for any degree or examination at any other university.

(3) This thesis does not contain other persons data, pictures, graphs or other information, unless specifically acknowlegded as being sourced from other persons.

(4) This thesis does not contain other persons writing, unless specifically acknowledged as being sourced from other researchers. Where other written sources have been quoted, then

   (a) their words have been re-written but the general information attributed to them has been referenced, or

   (b) where their exact words have been used, then their writing has been placed in italics and referenced.

(5) This thesis does not contain text, graphics or tables copied and pasted from the internet, unless specifically acknowledged, and the source being detailed in the thesis and in the reference sections.

Signed by candidate

Chun-Sung Huang

19 February 2018

Date

Declaration 2 - Publications

The following paper have been published from this thesis:

(1) Huang, C.-S., O'Hara, J.G., Mataramvura, S. (2017). Efficient pricing of discrete arithmetic Asian options under mean reversion and jumps based on Fourier-cosine expansions. *Journal of Computational and Applied Mathematics*, 311, 230-238.

The following papers have been submitted, passed the initial editorial screening, and are now under review in academic journals:

(1) Huang, C.-S., O'Hara, J.G., Mataramvura, S. (2017). Efficient option pricing under the double jump model with stochastic volatility and stochastic interest rate based on Fourier-cosine expansions.

(2) Huang, C.-S., O'Hara, J.G., Mataramvura, S. (2017). Highly efficient option valuation under the double jump framework with stochastic volatility and jump intensity based on Shannon wavelets inverse Fourier technique.

(3) Huang, C.-S., O'Hara, J.G., Mataramvura, S. (2017). Highly efficient Shannon wavelets-based pricing of power options under the double exponential jump framework with stochastic jump intensity and volatility.

Signed by candidate

Chun-Sung Huang

19 February 2018

Date

University of Cape Town

Faculty of Commerce

Co-author declaration for the joint paper in the PhD:

The paper/s to be included as part of the PhD thesis is the sole contribution of the candidate, Mr. Chun-Sung Huang, and that the roles of the other co-authors were moderation as the capacity of supervisors.

Authors: Chun-Sung Huang, John G. O'Hara, Sure Mataramvura

Title: Efficient pricing of discrete arithmetic Asian options under mean reversion and jumps based on Fourier-cosine expansions

Journal: Journal of Computational and Applied Mathematics (311, p230-238)

Huang's independent contribution:

☒ First author ☐ Corresponding author ☐ Other

Contributed to all chapters in the PhD.

Mataramvura's contribution:

☐ First author ☒ Supervisor ☐ Corresponding author

O'Hara's contribution:

☐ First author ☒ External supervisor ☐ Corresponding author

Has the above-mentioned paper been, or will be, part of another Doctoral or Master's thesis?

☐ Yes ☒ No

If yes, please elaborate:

Contribution from other PhD students: None

Contribution from Master's students: None

Do you verify that O'Hara and Mataramvura contributed to the join paper as described above?

☒ Yes ☐ No

If no, please specify:

Co-author's signature: Signed by candidate

Dr. Sure Mataramvura

Signed by candidate

Dr. John G. O'Hara

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my supervisors, Dr. John O'Hara and Dr. Sure Mataramvura. I thank them most sincerely for their guidance, patience and the endless support. I have leant so much on the PhD journey, and it would have been an impossible road to travel without my supervisors. I am honoured to have their mentorship from the get go.

I would also like to thank the School of Management Studies for allowing me to use their facilities. In particular, for providing me a quiet and peaceful space to complete my final write-up, which was nothing short of a blessing. To the Heads of Department, Prof. Anton Schlechter and A. Prof. Suki Goodman, thank you both for assuring me that I have the support of the Department should I needed it. Such gestures are the very spark that one need at times in order to push on.

I am grateful to the National Research Foundation of South Africa (NRF) and the University of Cape Town for their awards of grants, bursaries and other forms of funding that enabled me to carry out this research effectively.

Finally, I would like to take this opportunity to thank my family and friends for their love and encouragement. Most importantly, to my beautiful wife, Chian-Jia, thank you for the unwavering support and the tremendous understanding you have shown me throughout the past few years. During times when I had to be selfish in order to make progress, you have demonstrated that patience is more than just a virtue. You are, and always will be, the beacon of light in my darkest days.

Abstract

The pricing of exotic derivatives continues to attract much attention from academics and practitioners alike. Despite the overwhelming interest, the task of finding a robust methodology that could derive closed-form solutions for exotic derivatives remains a difficult challenge. In addition, the level of sophistication is greatly enhanced when options are priced in a more realistic framework. This includes, but not limited to, utilising jump-diffusion models with mean-reversion, stochastic volatility, and/or stochastic jump intensity. More pertinently, these inclusions allow the resulting asset price process to capture the various empirical features, such as heavy tails and asymmetry, commonly observed in financial data. However, under such a framework, the density function governing the underlying asset price process is generally not available. This leads to a breakdown of the classical risk-neutral option valuation method via the discounted expectation of the final payoff. Furthermore, when an analytical expression for the option pricing formula becomes available, the solution is often complex and in semi closed-form. Hence, a substantial amount of computational time is required to obtain the value of the option, which may not satisfy the efficiency demanded in practice. Such drawbacks may be remedied by utilising numerical integration techniques to price options more efficiently in the Fourier domain instead, since the associated characteristic functions are more readily available.

This thesis is concerned primarily with the efficient and accurate pricing of exotic derivatives under the aforementioned framework. We address the research opportunity by exploring the valuation of exotic options with numerical integration techniques once the associated characteristic functions are developed. In particular, we advocate the use of the novel Fourier-cosine (COS) expansions, and the more recent Shannon wavelets inverse Fourier technique (SWIFT). Once the option prices are obtained, the efficiency of the two techniques are benchmarked against the widely-acclaimed fast Fourier transform (FFT) method. More importantly, we perform extensive numerical experiments and error analyses to show that, under our proposed framework, not only is the COS and SWIFT methods more efficient, but are also highly accurate with exponential rate of error convergence. Finally, we conduct a set of sensitivity analyses to evaluate the

models' consistency and robustness under different market conditions.

Abbreviations

- FFT = fast Fourier transform
- COS = Fourier-cosine
- SWIFT = Shannon wavelets inverse Fourier technique
- CF = closed-form
- MGF = moment generating function
- MAE = mean absolute error
- $\mathbb{R}$ = The set of real numbers
- $\mathbb{Q}$ = Risk-neutral measure
- $\mathcal{F}$ = Filtration generated by the Brownian motion processes
- $\sum'$ = summation with the first term weighted by $\frac{1}{2}$
- $E[...]$ = The expectation of ...
- $Poi(.)$ = Poisson distribution
- $Exp(.)$ = Exponential distribution

Contents

Chapter 1

Introduction

1.1 Background

It is well-known that the celebrated Black-Scholes (BS) option pricing model fails to capture key phenomena, such as excess return rates, mean-reversion, non-constant volatility, stochastic interest rates, extreme price movements, etc., observed in a wide variety of financial data. The inability of the model to account for these empirical features has prompted the emergence of ever more sophisticated asset pricing dynamics, as extensions to the BS, in the attempt to better represent asset price movements, and accurately value the associated derivatives. For instance, one of the most popular extensions is allowing volatility to be governed by a separate stochastic process. Another pertinent, and well-received, improvement is to include jumps in the price process. Although such extensions may produce more relevant stochastic models, the probability density function governing the underlying price process is typically not available due to the additional level of sophistication. Hence, the usual risk-neutral option valuation method, via the discounted expectation of the final payoff, in the original pricing domain is not possible. However, since the associated characteristic function, defined as the Fourier transform of the density function, is often available, we can overcome the above-mentioned drawback by utilising numerical integration methods to price options in the Fourier domain instead.

Numerical integration techniques typically rely on the transformation to the Fourier space, and may be implemented once an analytical expression for the characteristic function has been developed. Notably, the pricing of financial derivatives can also be completed more efficiently in the Fourier domain. However, even state-of-the-art numerical integration techniques exhibit certain restrictions or shortfalls. Hence, once the final option prices have been computed, extensive numerical experiments and error

analyses should be performed to investigate the resulting accuracy and robustness of each method. In addition, conducting a thorough sensitivity analysis to check for pricing consistencies may provide further evidence in support of the method's robustness under different market conditions.

One of the most well-known numerical integration techniques in option pricing is the fast Fourier transform (FFT) of Carr & Madan (1999). There is also a strain of literature dealing with the efficient pricing of exotic options with the FFT. Given the associated characteristic function, the FFT method recovers the density function by first applying the inverse continuous Fourier transform. The resulting Fourier integral is then discretised before utilising the FFT algorithm to compute the discrete Fourier transforms. The computational complexity under the FFT is $O(N \log_2 N)$, where N is the number of integration points. This is a considerable reduction from the original requirement of order N^2.

An important focus area in computational finance is to further develop more efficient, yet accurate, option pricing methods. Fang & Oosterlee (2008) proposed the novel Fourier-cosine (COS) method, based on Fourier-cosine expansions, and demonstrated the method's ability to price options more efficiently than the FFT. While the FFT generally require thousands of grid points in the quadrature to achieve an acceptable level of accuracy, which adversely impacts the method efficiency, the COS method does not suffer the same setback. In addition, unlike the FFT, the COS method does not depend on the selection of an arbitrary damping factor for convergence, albeit having restrictions of its own.

Limitations of the COS method include choosing an appropriate integration bound to capture an adequate mass of the density function. This is particularly important given the heavy-tailed nature of density functions that commonly occur in finance. However, there is no existing algorithm to select the most suitable size of such an interval for all asset price processes, nor for the various types of options to be priced. Wider interval choices will require more cosine terms in the approximation to reach a desirable level of accuracy, which comes with the tradeoff of consuming more computational time. Furthermore, increasing the computational interval will also require all series coefficients to be recalculated. Hence, it is not immediately obvious how a satisfactory balance between efficiency and accuracy may be achieved without extensive numerical experiments and error analyses.

More recently, Ortiz-Gracia & Oosterlee (2016) proposed a highly efficient Shannon wavelet inverse Fourier technique (SWIFT) for pricing European options. Wavelet-based methods are generally more flexible and accurate for valuing options with longer expirations. Indeed, the SWIFT method does not require prior decisions on the trunca-

tion of integration range, and can accurately price both long- and short-dated options. This is one of the main improvements over the COS method. In addition, since Shannon wavelets are smooth, we can expect an accurate approximation of heavy-tailed density functions that often emerge in finance. Although both the COS and SWIFT methods may exhibit exponential convergence in its approximations, the latter has been shown empirically by prior studies to be the more efficient alternative in certain cases. The accuracy and error convergence of the SWIFT method, however, relies on the selection of a scale of approximation. We shall explore two separate methods in the sequel to identify the most adequate choice of such a scale.

In this thesis, we are primarily concerned with the efficient pricing of exotic options in a jump-diffusion framework with mean-reversion, stochastic volatility, and/or stochastic jump intensity. To the best of our knowledge, there exists a gap in the current literature that explores the efficient pricing of exotic options in our proposed framework mentioned above. In particular, there is limited research on utilising highly efficient, yet robust, pricing methods beyond the widely-acclaimed FFT technique. We address the research opportunity in three ways. Firstly, we investigate the highly efficient pricing of discretely-monitored arithmetic Asian options through the novel COS method. In particular, we allow for mean reversion and jumps in our underlying price dynamics. We show how the COS pricing methodology may be tailored to price discretely-monitored arithmetic Asian options once the required characteristic function is developed. Our results indicate that the COS method is both accurate and more efficient than the alternative FFT, which requires significantly more terms in the summation in order to obtain the same level of accuracy.

Secondly, while prior studies on the pricing of exotic options in a sophisticated framework generally turn to the result of Monte Carlo simulations as a benchmark for measuring the accuracy of numerical pricing methods, we shall explore a particular scenario within our proposed framework whereby a closed-form solution exists. Such a scenario allows us to compare the resulting COS option prices to that of a closed-form solution, which may yield a better indication on the accuracy and efficiency of the COS method in a more sophisticated framework. Given the level of complexity embedded in the closed-form solution, we show that the COS method is superior in terms of efficiency, and may be the preferred alternative to implement in practice. Moreover, our findings show that the COS method evidently outperforms the FFT in terms of both accuracy and efficiency. In addition to the above, we perform various error analyses to provide further evidence of the COS method's robustness under the proposed framework.

Finally, we explore the highly efficient pricing of exotic options under a double exponential jump framework with stochastic volatility and stochastic jump intensity. In

particular, we focus on the pricing of power options through the novel COS and, the more recent, SWIFT methods. We show how the required characteristic function may be derived and the pricing methodologies adapted to value power options. Our particular interest is on the accuracy, efficiency and robustness of the SWIFT method, compared to the COS and FFT, in a setting whereby randomness is introduced in both the underlying volatility and jump intensity. Once the level of accuracy and efficiency are determined, we conduct various error analyses to illustrate the robustness of the SWIFT method in pricing power options under our proposed framework. In addition, we demonstrate that not only is the SWIFT method accurate, it is also more efficient than both the COS and FFT. Furthermore, we show that the SWIFT method does not suffer the same drawback as the COS method when truncating the required integration range.

Lastly, it is important to emphasise that while the initial error analyses in each of the studies above may already indicate robustness under a particular scenario of our proposed framework, whether the pricing consistency and robustness holds across a wide range of market conditions, as represented by the specific parameters, is not immediately obvious. Hence, as a further contribution, extensive sensitivity analyses are conducted within each study in order to examine and demonstrate the pricing methods' robustness and consistency under different market conditions.

1.2 Organisation and Outline of the Thesis

The remainder of the thesis is organised as follows. We review the existing literature on the importance of jump-diffusion models with mean-reversion, stochastic volatility and/or stochastic jump intensity in Chapter 2. In addition, we explore the literature on the efficient pricing of options through state-of-the-art numerical integration techniques. More precisely, we motivate for the use of both the novel COS and SWIFT methods in pricing exotic options. This is then followed by a brief derivation of their respective pricing methodologies.

In Chapter 3, we present the pricing of arithmetic Asian options under mean-reversion and jumps based on the COS method. Specifically, our efficient pricing method is derived for the discretely monitored versions of the European-style arithmetic Asian options. We first develop the required characteristic function, and show how pricing can be carried out via the COS method. The analytical solutions obtained are then compared to the benchmark FFT-based prices for the examination of its accuracy and computational efficiency. We then conclude the chapter with a thorough sensitivity analysis to provide empirical evidence in support of the COS method's robustness

under different market conditions.

In Chapter 4, we further demonstrate the accuracy and efficiency of the COS method in a more practical double exponential jump framework, with stochastic volatility and stochastic interest rate. Notably, under such a framework, there already exists a semi closed-form solution for the price of an European option. This allows us to evaluate the accuracy of the COS prices to that of a closed-form solution, rather than the usual comparison to that of a Monte Carlo simulation as the benchmark. While prior research has demonstrated that the FFT is more efficient than the closed-form solution, which requires considerable computing power given its complex nature, we recommend the use of the COS method instead as a more efficient and practical alternative. Given the scenario, we are particularly interested in the accuracy of the COS method in comparison to the FFT when benchmarked to an existing closed-form solution under our proposed framework. Indeed, our numerical results demonstrate that not only is the COS method more efficient than the alternative FFT, but it is also more accurate when benchmarked to the existing semi closed-form solution. While the FFT method demands a large number of grid points in the quadrature to achieve a desirable level of accuracy, we show through our numerical analysis that the COS method requires significantly less number of terms in its expansion in order to reach the same level of accuracy. This is not dissimilar to our findings in Chapter 3. Furthermore, we evaluate the robustness of the COS method through various sensitivity analyses. Finally, we show that variability in the jump intensity, as well as the correlation between volatility and the underlying asset, has a significant impact on the resulting option prices across a range of strikes and maturity dates. Such findings advocate the use of stochastic jump intensity in addition to the notion of stochastic volatility.

In Chapter 5, we explore the highly efficient pricing of power options under the double exponential jump framework, with stochastic volatility and jump intensity. In particular, we investigate the highly efficient pricing of power options with both the COS and the novel Shannon wavelet inverse Fourier technique (SWIFT), and analyse the resulting price sensitivities and errors to changes in model parameters. We show, through our numerical analysis, that the SWIFT method is not only more efficient than its close competitors, such as the FFT and COS methods, but is also accurate with exponential error convergence. In addition, results from our sensitivity analyses provide further evidence in support of the SWIFT method's robustness under different market conditions.

Finally, we conclude our thesis in Chapter 6, and provide implications for future work.

Chapter 2

Literature Review and Pricing Methodology

In this chapter, we review the current literature on jump-diffusion models with mean-reversion and stochastic volatility. In addition, we explore prior research on efficient pricing of options with numerical integration techniques.

2.1 Literature Review

There exist a number of papers in the current literature that proposed the idea of stochastic volatility models. Such a notion seeks to address the volatility smile observed in the prices of traded options, and has become popular since the stock market crash of 1987. Despite the vast amount of research in the current literature that explores the avenue of stochastic volatility, the Hull & White (1987), Stein & Stein (1991) and Heston (1993) models notably received the lion's share of the attention. Among them, it is the work of Heston (1993) in the square root model framework that was hailed as the most prominent. However, it is also well documented that the Heston model does not accommodate for events with extreme price movements, such as a financial crisis. Hence, a more reasonable model would be to incorporate additional jumps to the price process. The notion of jumps has also been widely studied in the literature. For instance, Jorion (1988) examined the prices of stock market indices and exchange rates for discontinuities, while Geman & Roncoroni (2006) and Seifert & Uhrig-Homburg (2007) conducted investigations to provide empirical evidence of jumps in the power market. Further empirical evidence in support of jumps in commodity prices may also be found in the existing literature (see Deng, 2000; Hilliard & Reis, 1999; Schmitz et al., 2014).

Bates (1996) proposed the combination of stochastic volatility with the lognormal jumps of Merton (1976). Under such a framework, there is also a wealth of studies that combine lognormal jumps with stochastic volatility and stochastic interest rate model for option pricing (see, for example, Espinosa & Vives, 2006; Jiang, 2002; Pillay & O'Hara, 2011; Scott, 1997). However, despite the model's ability to explain observed features such as the volatility smile under the specifications of the lognormal jump framework, it contradicts with the asymmetric leptokurtic features in the returns distribution.

Kou (2002) proposed a more fruitful type of jump model, whereby the jump sizes follow an asymmetric double exponential distribution. Such a model has gained wide acceptance for its ability to better fit stock price data due to the high peak and heavy tails in the double exponential distribution (Kou & Wang, 2004). Under such an innovation, together with the notion of stochastic volatility, the model is capable of reproducing both the asymmetric leptokurtic feature embedded in the returns distribution, as well as the volatility smile observed in option prices (Zhang & Wang, 2013). Moreover, the study by Kou & Wang (2004) provided empirical evidence suggesting that the double exponential jump diffusion model outperforms the normal jump diffusion model when fitting stock price data. Part of the model's success may be attributed to its ability to account for overreaction and under-reaction on external information through the use of heavy tails and excess kurtosis, respectively.

Notably, under Bate's framework both the jump and diffusion risks are governed by the same diffusive volatility, and it becomes a difficult task to separate them. However, prior empirical research suggests that the two components of risk may differ significantly over time with a high degree of persistence. For example, the more recent study of Santa-Clara & Yan (2010), where separate stochastic processes were utilised to govern the volatility and jump intensity risks, respectively, provided empirical evidence showing a low level of correlation between the two risks. Such findings contradict the models which allow jump intensity to vary with the underlying diffusive volatility. A more robust model will be to allow the jump intensity to follow a different stochastic process uncorrelated to the volatility risk. Readers can also refer to other related studies, such as Chang et al. (2013) and Huang et al. (2014), in support of the innovation of allowing jump intensity to follow a stochastic process. The latter is an extension of Santa-Clara & Yan (2010) and provides evidence in support of modelling jump intensity as a mean-reversion process to allow for improved correspondence to real market data.

Finally, apart from jumps and stochastic volatility, mean reversion is another significant aspect to consider when modelling financial assets. Indeed, it has been well documented that the prices of certain asset classes, such as commodities, show evidence of jumps and mean reversion. Hence, the pricing of options within these asset classes has also become

an important focus in the field of quantitative finance. For example, due to the impact of relative prices on the supply of both copper and oil, prices tend to fluctuate randomly around some equilibrium level (see Schwartz, 1997). In addition to the above mentioned commodities, Bessembinder et al. (1995) provides strong evidence supporting mean reversion in nine commodity markets, while Casassus & Collin-Dufresne (2005) reveals the existence of such anomaly in the precious metals market. Apart from commodities, however, evident motivating the patterns of mean reversion has also been found in exchange rates and, interestingly, certain stock prices as well (see Jorion & Sweeney, 1996; Chaudhuri & Wu, 2003).

While the ever increasing sophistication of stochastic models provide researchers and practitioners a better tool to effectively capture the various market phenomena, and more accurately value derivatives as a result, finding a closed-form solution for option prices under such framework becomes a mammoth task. In addition, given the significant level of complexity introduced in the underlying stochastic processes, whenever solutions do exist (and usually in semi closed-form) they generally lack analytical tractability. Hence, a substantial amount of computational time is required to obtain the resulting option prices. The importance of such drawback needs to be emphasised as it may not satisfy the calculation speed required in practice.

The aforementioned shortfalls prompt the emergence of ever more efficient numerical integration techniques to price financial options. Firstly, under a highly sophisticated framework, the probability density function of the price process is often not available. The absence of a plausible solution for the density function makes valuing options under the usual discounted expectation of the option payoff, in the original price domain, particularly difficult. However, since the associated characteristic function is generally available, options may be priced by evaluating the equivalent integral under the Fourier domain instead. Numerical integration techniques rely on the transformation to the Fourier domain, and utilises the characteristic function instead of the density function governing the asset price process. Secondly, the excessive computational times required to price options, especially path-dependent exotics such as Asian and Barrier options, when the closed-form (or semi closed-form) solution is too complex, may be circumvented by pricing more efficiently in the Fourier domain through the use of the characteristic functions.

Currently, the most widely-accepted numerical integration technique is the FFT method of Carr & Madan (1999). There is already a strain of literature on the valuation of options using the FFT algorithm (see, for example, Chourdakis, 2005; Dempster & Hong, 2002; Huang et al., 2013, 2014, 2017; Ibrahim et al., 2013, 2014; Lord et al., 2008; Pillay & O'Hara, 2011; Zhang & Wang, 2013, among others). The FFT has proven to

be a successful model, albeit drawbacks that hampers the accuracy and efficiency of the model do exist. Firstly, since the integrands of Fourier transform integrals are highly oscillatory, tiny grids are required for acceptable levels of accuracy when using quadrature-based techniques (such as in the FFT algorithm). However, a choice of finer grids, which may result in thousands of grid points, will significantly increase the computational time required. Secondly, the FFT method relies on an artificial selection of the damping factor for convergence, which may not be ideal in practice.

Fang & Oosterlee (2008) proposed the COS method, based on Fourier-cosine expansions, as an alternative to the FFT. The method recovers the density function through a Fourier series expansion, in terms of a partial sum of orthogonal basis function, with cosines forming the global basis. Thereafter, the characteristic function is used to approximate each coefficient in the series, before its integration with the discounted terminal payoff to obtain the final option price. A number of prior research has shown that the COS method is more efficient than the FFT. While Fang & Oosterlee (2008) provided evidence of COS method's efficiency over the FFT when pricing European options under the Lévy and Heston models, a follow-up paper by Fang & Oosterlee (2009) provided pricing comparisons for early-exercise barrier options. Evidence of efficient pricing of Asian options is provided by Zhang & Oosterlee (2013), Zhang & Oosterlee (2014) and Huang et al. (2017). Studies of COS method pricing with stochastic volatility can be found in Fang & Oosterlee (2010) and Zhang & Geng (2016). It is further noteworthy that, apart from its superior efficiency, the COS method also exhibits exponential error convergence without depending on an arbitrary damping factor. In addition, the COS method is capable of allowing a user-defined vector of strikes to price a basket of options efficiently with just a single computation. While other efficient methods in pricing plain vanilla options already exists, such as the double exponential transform (Mori & Sugihara, 2001; Yamamoto, 2005) and the fast Gauss transform (Broadie & Yamamoto, 2003), the COS method is preferred due to its ability to accommodate for more general asset price dynamics.

More recently, Ortiz-Gracia & Oosterlee (2016) pioneered the novel SWIFT method to efficiently compute the price of European-style options. Such a method, advocated by the use of Shannon wavelets (see Cattani, 2008) to represent the underlying density, overcomes the two main shortfalls embedded in the COS method. Firstly, there exist no prescribed algorithm to determine the most suitable integration range, which is vital to capturing an adequate mass of the underlying density function. This is of particular interest to us, given the heavy-tailed nature in the density functions governing our proposed asset price dynamics. Secondly, due to the periodicity behaviour of the cosine functions, which forms a global basis in Fourier-cosine expansions, attention

needs to be given at the vicinity of the integration boundaries. This is imperative for the case of long-dated options where errors due to rounding-off may accumulate near the boundaries of our integration range. Short-dated options, on the other hand, yields highly peaked density functions, and may require considerable amount of cosine terms for accurate approximations (Ortiz-Gracia & Oosterlee, 2016).

Shannon wavelets are smooth wavelets generated by the cardinal sine functions, and benefit from local wavelet basis. Since the wavelets provide local approximations, it is possible to resolve the errors occurring near the boundaries of the integration domain. One of the main advantages of the SWIFT method is that it does not rely on prior decisions for the truncation of the required integration range. Instead, we can define an error tolerance level, and allow the local wavelet basis to indicate whether the amount of density mass captured is in accordance such a level of tolerance. It is further noteworthy that the number of cardinal sine terms required for our approximation is defined automatically according to the integration interval identified (see Ortiz-Gracia & Oosterlee, 2016).

2.2 Pricing Methodology

In this Section, we briefly derive the COS and SWIFT pricing methods in line with Fang & Oosterlee (2008) and Ortiz-Gracia & Oosterlee (2016), respectively.

2.2.1 Fourier-Cosine Expansions

The pricing of options under the COS method, as with all numerical integration techniques, follows from the discounted expected payoff approach under the risk-neutral measure $\mathbb{Q}$:

$$\begin{aligned} v(x,t) &= e^{-r(T-t)}\mathbb{E}[v(y,T)|x] \\ &= e^{-r(T-t)}\int_{\mathbb{R}} v(y,T)g(y|x)dy, \end{aligned} \tag{2.1}$$

where $v(x,t)$ denotes the option value at time t, and r is the interest rate. In addition, x and y are state variables at time t and expiration date T, respectively. Typically, the option's payoff function, v, is known, but its transitional density $g(y|x)$ is not. Fang & Oosterlee (2008) proposed an approximation of the transition probability, based on (2.1), with a truncated domain $[a,b]$ by a truncated Fourier-cosine series expansion with N_C terms, based on the conditional characteristic function, i.e.:

$$g(y|x) \approx \frac{2}{b-a}\sum_{h=0}^{N_C-1}{}' \operatorname{Re}\left[\psi\left(\frac{h\pi}{b-a};x\right)e^{-ih\pi\frac{a}{b-a}}\right]\cos\left(h\pi\frac{y-a}{b-a}\right), \tag{2.2}$$

where $\psi(\nu; x)$ is the conditional characteristic function of $g(y|x)$, and a, b denotes the integration range in the original domain.[1] $\sum'$ indicates that the first term of the summation is multiplied by a weight of one-half. Re[$\cdot$] denotes the real part of the argument.

Approximation (2.2) is possible since, for all real functions g defined on a finite range $[a, b]$, its corresponding cosine series expansion is given by the expression:

$$g(y) = \sum_{h=0}^{\infty}{}' A_h \cdot \cos\left(h\pi \frac{y-a}{b-a}\right) \quad \text{where} \quad A_h = \frac{2}{b-a}\int_a^b g(y)\cos\left(h\pi\frac{y-a}{b-a}\right)dy.$$

The conditions for the existence of the characteristic function (a Fourier transform) ensures that the integrand of the characteristic function, $\psi(\nu) = \int_{\mathbb{R}} e^{iy\nu}g(y)dy$, decays to 0 at $\pm\infty$. Hence, for appropriately chosen bounds $[a, b]$, we can approximate:

$$\psi(\nu) := \int_{\mathbb{R}} e^{iy\nu}g(y)dy \approx \int_a^b e^{iy\nu}g(y)dy. \tag{2.3}$$

without significant loss of accuracy. Finally, substituting approximation (2.3) into A_h, and identifying that cosines represent the real part on the Euler formula, equation (2.2) follows after truncating the infinite series. The overall error introduced in the COS method approximations, as shown above, converges exponentially for adequately selected truncation range $[a, b]$ (Fang & Oosterlee, 2008, 2009).[2]

The appropriate choice of integration range $[a, b]$ may be determined by making use of the cumulants, such that the error of the approximation is within some tolerance level Tol, i.e. $|\int_{\mathbb{R}} g(y|x)dy - \int_a^b g(y|x)dy| < Tol$ (see Fang & Oosterlee, 2008). Denoting μ_j the j^{th} cumulant of the log-asset price, and are determined by evaluating the j^{th} derivative of the associated moment generating function (MGF) at 0, the truncation range can be given by:

$$[a, b] = \left[\mu_1 - L\sqrt{\mu_2 + \sqrt{\mu_4}},\ \ \mu_1 + L\sqrt{\mu_2 + \sqrt{\mu_4}}\right], \tag{2.4}$$

where $L = 10 \sim 12$, in accordance with Fang & Oosterlee (2008, 2009). The number of cumulants to include in (2.4) may depend on the nature of the density function governing the asset price process. For instance, we can include the 4th cumulants to accommodate the sharp peaks and fat tails in the density function, whereas μ_6 may be required when dealing with extreme short maturities (see Fang & Oosterlee, 2008).

Finally, replacing the conditional density function in (2.1) with its approximation (2.2), and interchanging the summation and integration, we obtain the COS formula to price

[1] The integration does not suffer a significant loss of accuracy from the truncation for adequately chosen a and b values, as the density functions decays to zero at $\pm\infty$

[2] The method converges exponentially for $g(y|x) \in \mathbb{C}^{\infty}[a, b] \in \mathbb{R}$, and algebraic otherwise.

an option with payoff $v(x,t)$:

$$\hat{v}(x,t) = e^{-r(T-t)} \sum_{h=0}^{N_C-1}{}' \text{Re}\left[\psi\left(\frac{h\pi}{b-a}; x\right) e^{-ih\pi\frac{a}{b-a}}\right] V_h, \tag{2.5}$$

where $\hat{v}(x,t)$ is the approximation of the option value at time t, Re[.] denotes the real part of the argument, and

$$V_h := \frac{2}{b-a}\int_a^b v(y,T)\cos\left(h\pi\frac{y-a}{b-a}\right)dy \tag{2.6}$$

are the Fourier-cosine coefficients of payoff $v(y,T)$. We solve for V_h in the respective Chapters for the different options to price. Readers are also referred to Fang & Oosterlee (2008) for detailed derivations of the above.

It is worthwhile emphasising at this point the potential pricing error due to the choice of parameter L in (2.4). This is particularly the case with call options, where the payoff function is generally unbounded and grows exponentially with respect to the log-asset price. Consequently, significant approximation errors may emerge when truncating the integration domain of the risk-neutral formula with large values of L. However, since put options are bounded by the strike price K, they do not suffer the same setback. Hence, one can also utilise the put-call parity to obtain call option values once the put option prices are computed. This may provide slight improvements to the pricing accuracy (Fang & Oosterlee, 2008).

2.2.2 Shannon Wavelet Inverse Fourier Technique

The SWIFT method, as with other efficient numerical integration techniques, departs from the risk-neutral option pricing formula:

$$\begin{aligned} v(x,t) &= e^{-r(T-t)}\mathbb{E}[v(y,T)|x] \\ &= e^{-r(T-t)}\int_{\mathbb{R}} v(y,T)g(y|x)dy, \end{aligned} \tag{2.7}$$

where x and y, as with the COS method, are state variables at times t and T, respectively, $\mathbb{E}$ is the expectation operator, and v denotes the value of the option. The expectation is taken under the risk-neutral measure $\mathbb{Q}$, with $g(y|x)$ the conditional density under such a measure. Finally, r is the constant risk-free rate of interest.

While the transition density of sophisticated stochastic processes is typically not known, its corresponding characteristic function is generally available in the Fourier space. The SWIFT framework proposed by Ortiz-Gracia & Oosterlee (2016) consists of approximating density g by utilising a finite combination of Shannon scaling functions.

Thereafter, the density coefficients of the approximation from its Fourier transform may be recovered. Following Ortiz-Gracia & Oosterlee (2016), we consider the transition density g in $L^2(\mathbb{R})$ for some random variable X, and its Fourier transform

$$\psi^*(\nu) = \int_{\mathbb{R}} e^{-i\nu x} g(x) dx.$$

In wavelet theory, a multiresolution analysis (MRA) consists of a family of closed subspaces, $\cdots \subset V_{-2} \subset V_{-1} \subset V_0 \subset V_1 \subset V_2 \cdots$ in $L^2(\mathbb{R})$, such that $\overline{\bigcup V_j} = L^2(\mathbb{R})$ and $\bigcap V_j = \{0\}$, for $j \in \mathbb{Z}$. In addition, $g(x) \in V_j \iff g(2x) \in V_{j+1}$. Under such conditions, a function $\phi \in V_0$ exists, such that the set $\{\phi_{j,k}(x) := 2^{j/2}\phi(2^j x - k)\}_{k\in\mathbb{Z}}$ generates an orthonormal basis for V_j.[3] The function ϕ is known as the father wavelet or scaling function, and generates an orthonormal basis for every V_j subspace (which are scaled versions of V_0).

If we define W_j as the space of functions in V_{j+1} and not in V_j, i.e. $V_{j+1} = V_j \bigoplus W_j$, then there exists a function $\upsilon \in W_0$, known as the mother wavelet or wavelet functions, such that the set $\{\upsilon_{j,k}(x) := 2^{j/2}\upsilon(2^j x - k)\}_{k\in\mathbb{Z}}$ forms the orthonormal basis of W_j.

Shannon wavelets analysis allows for the approximation of a density function g at a chosen level of resolution m. Hence, we have

$$g(x) \approx \mathcal{P}_m g(x) = \sum_{j=-\infty}^{m-1} \sum_{k=-\infty}^{k=\infty} \Big[\int_{-\infty}^{\infty} g(x)\upsilon_{j,k}(x)dx \Big] \upsilon_{j,k}(x) = \sum_{k\in\mathbb{Z}} c_{m,k}\phi_{m,k}(x), \quad (2.8)$$

where the projection $\mathcal{P}_m g(x)$, such that $P_m : L^2(\mathbb{R}) \to V_m$, converges to $g(x)$ in $L^2(\mathbb{R})$ when $m \to \infty$. Hence, higher values of m improves the approximation of the density function g.[4] The integral in (2.8) represents the wavelet coefficients, and $c_{m,k} = \int_{\mathbb{R}} g(x)\phi_{m,k}(x)dx$ are the density, or scaling, coefficients. Finally, $\phi_{m,k}$ are the scaling functions.

In this work, we consider the set of Shannon scaling (or *sinc*) functions in the subspace V_m, defined by

$$\begin{aligned} \phi_{m,k}(x) &= 2^{m/2}\phi(2^m x - k) \\ &= 2^{m/2}\frac{\sin(\pi(2^m x - k))}{\pi(2^m x - k)}, \quad \forall k \in \mathbb{Z}, \end{aligned}$$

as the father wavelets, such that when $m = k = 0$ we have the basic scaling function $\phi(x) = \text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}$. The mother wavelet functions in subspace W_m, on the other hand, are given by

$$\upsilon_{m,k}(x) = 2^{\frac{m}{2}} \frac{\sin(\pi(2^m x - k - 1/2)) - \sin(2\pi(2^m x - k - 1/2))}{\pi(2^m x - k - 1/2)}, \quad \forall k \in \mathbb{Z},$$

[3] a different function ϕ will result in a different MRA.

[4] The entire wavelet summation is obtained if we allow j to go up to infinity.

such that $\upsilon(x) = \text{sinc}(x - 1/2) - 2\,\text{sinc}(2x - 1)$, when $m = k = 0$. It follows that the Fourier transform of the father and mother wavelets are given by

$$\begin{aligned}\hat{\phi}_{m,k}(\nu) &= \int_{\mathbb{R}} e^{-i\nu x}\phi_{m,k}dx \\ &= \frac{e^{-i\frac{k}{2^m}\nu}}{2^{m/2}} \cdot \text{rect}\left(\frac{\nu}{2^{m+1}\pi}\right),\end{aligned}$$

and

$$\hat{\upsilon}_{m,k}(\nu) = \frac{e^{-i\frac{k+1/2}{2^m}\nu}}{2^{m/2}}\left[\text{rect}\left(\frac{\nu}{2^m\pi} - \frac{3}{2}\right) + \text{rect}\left(-\frac{\nu}{2^m\pi} - \frac{3}{2}\right)\right],$$

where the *rectangular* function, $\text{rect}(\cdot)$, is given by

$$\text{rect}(x) = \begin{cases} 1 & \text{if } |x| < 1/2, \\ 1/2 & \text{if } |x| = 1/2, \\ 0 & \text{if } |x| > 1/2. \end{cases}$$

Due to its simplicity, we consider the father wavelets in the time domain instead of the mother wavelets (illustrated in Figure 2.1). In addition, it is worthwhile underlining that although Shannon wavelet functions have a slow rate of decay in the time domain, they exhibit very sharp decay in the Fourier domain.

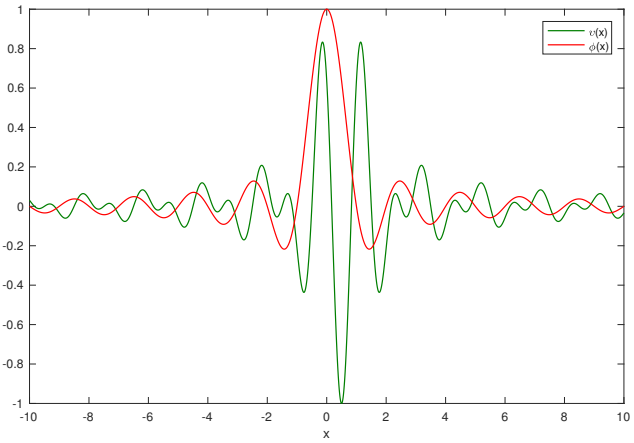

Figure 2.1: Shannon father $\phi(x)$ and mother $\upsilon(x)$ functions

If we assume that the associated density function tends towards zero at either tails of the distribution, as with most density functions in finance, then the infinite summation in (2.8) can be accurately approximated by:

$$g_m(x) := \sum_{k=k_1}^{k_2} c_{m,k}\phi_{m,k}(x), \tag{2.9}$$

for appropriately chosen k_1 and k_2 values (see Lemma 1 of Ortiz-Gracia & Oosterlee, 2016).

Through Vieta's formula, the cardinal sine functions can be represented by an infinite product of cosine terms, i.e. $\text{sinc}(x) = \prod_{j=1}^{\infty} \cos(\pi x/2^j)$ (see Gearhart & Schultz, 1990). If we truncate the infinite product to a domain with J factors only, the cosine product-to-sum identity allows us to approximate the cardinal sine function as (Quine & Abrarov, 2013):

$$\text{sinc}(x) \approx \text{sinc}^*(x) := \frac{1}{2^{J-1}} \sum_{j=1}^{2^{J-1}} \cos\left(\frac{2j-1}{2^J}\pi x\right). \tag{2.10}$$

Hence, by replacing the cardinal sine function in $c_{m,k}$, we can show that

$$\begin{aligned} c_{m,k} &= 2^{m/2} \int_{\mathbb{R}} g(x)\,\text{sinc}(2^m x - k)dx \\ &\approx \frac{2^{m/2}}{2^{J-1}} \sum_{j=1}^{2^{J-1}} \int_{\mathbb{R}} g(x) \cos\left(\frac{2j-1}{2^J}\pi(2^m x - k)\right)dx \\ &=: \tilde{c}_{m,k}. \end{aligned}$$

Finally, since $\int_{\mathbb{R}} g(x)\cos(\nu x)dx = \text{Re}[\psi^*(\nu)]$, where Re[.] denotes the real part of the argument, we can obtain the following expression for the density coefficients in (2.9) above:

$$c_{m,k} \approx \tilde{c}_{m,k} = \frac{2^{m/2}}{2^{J-1}} \sum_{j=1}^{2^{J-1}} \text{Re}\left[\psi^*\left(\frac{(2j-1)\pi 2^m}{2^J}\right) e^{\frac{ik\pi(2j-1)}{2^J}}\right]. \tag{2.11}$$

Detailed derivations can be found in Ortiz-Gracia & Oosterlee (2016).

While the most computationally demanding part of (2.11) above is the evaluation of the characteristic function $\psi^*(.)$ at each grid points, the process may be simplified by involving the FFT. This is possible by first choosing a constant J, instead of a different J for each k, defined by $\bar{J} := \lceil \log_2(\pi\Gamma_m) \rceil$, where $\Gamma_m := \max_{k_1<k<k_2} \Gamma_{m,k}$, with k_1 and k_2 fixed indices, and $\lceil x \rceil$ denoting the smallest integer greater than x. We can define $\Gamma_{m,k} := \max(|2^m a - k|, |2^m a + k|)$ for some constant a, such that $G(-a) + 1 - G(a) < \epsilon$ for $\epsilon > 0$, where $G(\cdot)$ represents the distribution function (see Theorem 1 of Ortiz-Gracia & Oosterlee, 2016).

It follows that each evaluation of $\psi^*(.)$ requires only a single calculation and used by the FFT algorithm. Furthermore, if we assume $\psi^*(\frac{(2j+1)\pi 2^m}{2^{\bar{J}}}) = 0$ for $j = 2^{\bar{J}-1}, ..., 2^{\bar{J}} - 1$, then expression (2.11) becomes:

$$\tilde{c}_{m,k} = \frac{2^{m/2}}{2^{\bar{J}-1}} \text{Re}\left[\sum_{j=0}^{2^{\bar{J}-1}} \psi^*\left(\frac{(2j+1)\pi 2^m}{2^{\bar{J}}}\right) e^{\frac{2\pi i k j}{2^{\bar{J}}}} e^{\frac{ik\pi}{2^{\bar{J}}}}\right]. \tag{2.12}$$

Hence, the density coefficients may be obtained more efficiently through the FFT.[5]

We highlight at this point the importance of an appropriate choice for the wavelet scale, m, upfront, since such a scale of approximation cannot be adjusted at a later stage without recalculating all approximation terms. Hence, along the lines of Maree et al. (2017), we utilise the Fourier transform, $\psi^*(\nu)$, of the density function to determine the appropriate choice of wavelet scale, m, in analytical form. This is possible since our stochastic process of interest (5.1) satisfies the following condition:

$$\left|\psi^*(\nu)\right| = \left|e^{\mathcal{E}(\nu)\tau}\right| \leq \mathcal{C}e^{-\tilde{d}\tau|\nu|^{\iota}}, \tag{2.13}$$

where $\tau := T - t$, $\iota = 2$ and $\mathcal{C}, \tilde{d} > 0$.[6] From (2.13), it can be shown through integration that the mass of the tails of the characteristic function, $\mathcal{A}(2^m\pi)$, is bounded by

$$\mathcal{A}(2^m\pi) \leq \frac{\mathcal{C}}{\pi\iota(\tilde{d}\tau)^{1/\iota}}\Gamma\Big(\frac{1}{\iota}, \tilde{d}\tau(2^m\pi)^{\iota}\Big) =: \epsilon_m, \tag{2.14}$$

where Γ represents the incomplete gamma function, and

$$\epsilon_m \sim \frac{\mathcal{C}(2^m\pi)^{1-\iota}}{\pi\iota\tilde{d}\tau}e^{-\tilde{d}\tau(2^m\pi)^{\iota}}, \tag{2.15}$$

for large values of $2^m\pi$ (see Maree et al., 2017). Note that (2.15) holds uniformly for ν. Notably, from (2.15), it can be observed that the ϵ_m converges exponentially with respect to m. However, since parameters $\mathcal{C}$ and $\tilde{d}$ are not available in most cases, we overcome the drawback by allowing the substitution of $\mathcal{C}e^{-\tilde{d}\tau|\nu|^{\iota}} \sim |\psi^*(\pm 2^m\pi)|$ into (2.15). Hence, ϵ_m can we approximated by,

$$\epsilon_m \sim \frac{(2^m\pi)^{1-\iota}}{2\pi\iota\tau}\Big(\left|\psi^*(-2^m\pi)\right| + \left|\psi^*(2^m\pi)\right|\Big). \tag{2.16}$$

Using (2.16) above, we can efficiently compute the approximation error for a particular choice of wavelet scale, m, without a significant amount of CPU time. We can, therefore, carry out a straight forward iterative procedure, through choosing $m = 0, 1, 2, ...$, until ϵ_m is less than a tolerance level, TOL, determined a priori.

Finally, it is also worthwhile mentioning that one major advantage of utilising Shannon scaling functions is the convenience of determining the area under the approximated density function with minimal efforts once the coefficients $c_{m,k}$ are calculated. The approximated area may then be determined by (Ortiz-Gracia & Oosterlee, 2016):

$$\tilde{H} = \frac{1}{2^{m/2}}\left(\frac{c_{m,k_1}}{2} + \sum_{k_1<k<k_2} c_{m,k} + \frac{c_{m,k_2}}{2}\right), \tag{2.17}$$

[5] it is worthwhile noting that the computational complexity is reduced from $\mathcal{O}(2^{\bar{J}-1}\cdot(k_2-k_1+1))$ in a direct calculation in (2.11) to $\mathcal{O}(\log(2)\cdot\bar{J}\cdot 2^{\bar{J}+1})$ with (2.12).

[6] Note that ι may take on different values depending on the expression of $\psi^*(\nu)$.

where $\tilde{H} \approx 1$. Adequate determination of indices k_1 and k_2 with the cumulants, as with the COS method, are presented in our numerical experiments in Chapter 5. The closer the estimated area is to one, the more accurate our approximation of the target density function.

Chapter 3

Efficient pricing of discrete arithmetic Asian options under mean reversion and jumps based on Fourier-cosine expansions

3.1 Introduction

A topic of ongoing interest is the long standing hard problem of pricing arithmetic Asian options. The payoffs of these path-dependent exotics are based on the arithmetic average of the underlying prices monitored at fixed dates prior to maturity. The monitoring dates used to measure the arithmetic averages may also be taken at different frequencies, such as daily, weekly or monthly. Unlike its closely related geometric type, the prices of the more commonly traded arithmetic Asian options must be approximated numerically. This is mainly due to the absence of an analytically tractable solution for the distribution of the sum of log normally distributed random variables.

Asian options, introduced in 1987, are now widely traded in the commodities market as a hedging tool. For instance, various delivery companies in the gas market utilise Asian options to their advantage under risk management (see Eydeland & Wolyniec, 2003). The popularity of Asian options arises mainly from its averaging effect, which is able to reduce possible risk of market manipulation in the price of the underlying at maturity. In addition, since averages move in a more stable way in comparison to individual prices, the volatility inherent in the underlying price is reduced as a result. Further information on Asian options with its history and evolution may also be found

in Boyle & Boyle (2001) and Marena et al. (2014).

We propose an efficient pricing method for discrete arithmetic Asian options under a pricing dynamic which exhibits both jumps and mean reversion. Essentially, the model is a jump-diffusion extension of the one utilised by Fusai et al. (2008) (as proposed by Chung & Wong, 2014). Apart from accuracy in the pricing, computational efficiency is also of equal importance, if not more, particularly, for high frequency traders. Such notion brings about the non-trivial problem of finding a reasonable tradeoff between accuracy and efficiency in the pricing methods. As a result, efficient pricing methods of exotic options have also gained much interest from both practitioners and academics alike.

In option pricing, the valuation of complex contracts requires efficient numerical methods. The conditional expectation of the option payoff under the risk-neutral measure can be bridged with the solution of a partial differential equation through the well-known Feynman-Kac theorem. It then follows that various numerical pricing techniques, including numerical integration, can be developed. These numerical integration techniques rely on the transformation into the Fourier domain, which is particularly useful especially since the density function of many relevant underlying price process, required for the integration in the original domain, is not known. However, its Fourier transform, the characteristic function, often is. It then follows that the fast Fourier transform (FFT) method, introduced by Carr & Madan (1999) and Dempster & Hong (2002), may be applied to calculate the option price efficiently. However, Fang & Oosterlee (2008) proposed a novel pricing method, the Fourier-cosine expansions (COS method), as an alternative to the FFT. Such method could further improve the speed in the pricing.

In this chapter, we propose to price discrete arithmetic Asian options under the assumption of mean reversion and jumps with the COS method. We show through numerical examples that the COS method is indeed more efficient than the benchmark FFT, used by Chung & Wong (2014). It was also shown in Chung & Wong (2014) that the FFT is superior to the commonly implemented Monte Carlo simulation.

3.2 Price Process with Mean-Reversion and Jumps

3.2.1 Model Specification

Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space on which a Brownian motion process W_t and a Poisson process N_t, with intensity $\lambda > 0$, is defined for $0 \leq t \leq T$. Furthermore, we assume independence between the Brownian motion and Poisson process. Suppose $\mathbb{Q}$

is the risk neutral measure under which the price process is governed by the following dynamics:

$$dS_t = \kappa\Big(\theta - \frac{\mu\lambda}{\kappa} - S_{t_-}\Big)dt + \sigma\sqrt{S_{t_-}}dW_t + JdN_t, \tag{3.1}$$

where $J \sim Exp(\mu)$ and $N_t \sim Poi(\lambda t)$.

The model proposed here is an extension of the Fusai et al. (2008) model, whereby a jump component has been added to the original spot price process, which is defined as a square root process driven by a Brownian motion. The jump size J and its arrival rate N_t are independent, and are modelled with an exponential distribution and a Poisson process, respectively. More specifically, the proposed price process is a CIR model with an exponential jump extension. Further justifications for the specific choice of jump dynamics can be found in Hoepfner (2009) and Beliaeva & Nawalkha (2012), with the latter suggesting the non-existence of an analytical solution under lognormal jumps.

The use of the CIR as a base model gives rise to two main advantages in terms of Asian option pricing. Firstly, since we are interested in the average price of the underlying, the existence of the characteristic function for $\int_0^T r_t dt$ in the CIR model, used widely in the modelling of interest rates, helps simplify the problem at hand. Secondly, instead of a log price, by choosing suitable parameters according to the Feller condition, we can model the stock price directly under the CIR model while maintaining its positivity. Such positivity is consistent even after jumps are added, as the jump sizes are modelled using an exponential distribution, which is always positive.

Our aim is to price an Asian option at initial time 0 that matures at terminal time T. The underlying price will be recorded at some regular time interval to allow for the discretely monitored Asian option in question. We split the pricing interval $[0, T]$ into $n+1$ sets of $\triangle$-spaced monitoring dates. Hence, we have dates $0, \triangle, ..., n\triangle = T$. Such a setup allows for the computation of an analytical price for the Asian option with payoff depending on arithmetic average $A_n = \sum_{j=0}^{n} \omega_j S_{j\triangle}$ and the terminal price $S_{n\triangle}$, where ω_j is the weight assigned to price $S_{j\triangle}$ and $\sum_{j=0}^{n} \omega_j = 1$. It is worthwhile mentioning that the weights assigned to the underlying price at different time intervals need not be equal. Table 3.1 summarises the various type of options that may be priced under our model assumption.

3.2.2 Derivation of Joint Characteristic Function

The joint characteristic function between $S_{n\triangle}$ and A_n is required for us to price the Asian options analytically. We first determine the characteristic function of $S_{t+\triangle}$ and proceed to derive the joint characteristic function of the pair $S_{n\triangle}$ and A_n.

Table 3.1: Payoff functions of various options

Option type	Payoff function
Fixed strike Asian call	$\max\{A_n - K, 0\}$
Fixed strike Asian put	$\max\{K - A_n, 0\}$
Floating strike Asian call	$\max\{S_{n\triangle} - A_n - K, 0\}$
Floating strike Asian put	$\max\{K + A_n - S_{n\triangle}, 0\}$
European call	$\max\{S_{n\triangle} - K, 0\}$
European put	$\max\{K - S_{n\triangle}, 0\}$

The characteristic function of $S_{t+\triangle}$ can be defined as $\psi^{\varphi}(t, S_t) \equiv \mathbb{E}_t^{\mathbb{Q}}[e^{i\upsilon S_{t+\triangle}}]$ with parameter set $\varphi = \{\kappa; \mu; \lambda; \theta; \sigma\}$.[1] The generalised Feynman-Kac theorem (see Duffie et al. (2000) and Cont & Tankov (1975)) implies that ψ^{φ} solves the following partial integro-differential equation (PIDE):

$$\frac{\partial \psi^{\varphi}}{\partial \tau} + \kappa\Big(\theta - \frac{\mu\lambda}{\kappa} - S_u\Big)\frac{\partial \psi^{\varphi}}{\partial S_u} + \frac{1}{2}\sigma^2 S_u \frac{\partial^2 \psi^{\varphi}}{\partial S_u^2} + \lambda \int_{-\infty}^{\infty} [\psi^{\varphi}(u, S_u + J) - \psi^{\varphi}(u, S_u) q(J)] dJ = 0, \tag{3.2}$$

with boundary condition $\psi^{\varphi}(t + \triangle, S_{t+\triangle}) = e^{i\upsilon S_{t+\triangle}}$, where $u \in [t, t + \triangle]$ and $\tau = t + \triangle - u$, and $q(J)$ is the distribution of J. Coefficients, $\kappa(\theta - \frac{\mu\lambda}{\kappa} - S_u)$ and σ, of the mean reverting asset price process (3.1) are both affine in nature. It follows that the solution to (3.2) is of exponential affine form $\psi^{\varphi}(u, S_u) = e^{-\alpha^{\varphi}(\tau;\upsilon)S_u - \beta^{\varphi}(\tau;\upsilon)}$. Substituting into (3.2) above, and matching the characteristic function of the exponential distribution governing the jumps, we obtain $\mathbb{E}(e^{-\alpha^{\varphi}(\tau;\upsilon)}) = \frac{1}{1+\mu\alpha^{\varphi}(\tau;\upsilon)}$. Further simplification will allow us to obtain the following ODE (with differentiations taken with respect to τ):

$$\alpha^{\varphi\prime}(\tau; \upsilon) + \kappa\alpha^{\varphi}(\tau; \upsilon) + \frac{1}{2}\sigma^2[\alpha^{\varphi}(\tau; \upsilon)]^2 = 0, \tag{3.3}$$

$$\beta^{\varphi\prime}(\tau; \upsilon) - \kappa\Big(\theta - \frac{\mu\lambda}{\kappa}\Big)\alpha^{\varphi}(\tau; \upsilon) + \lambda\Big(\frac{1}{1 + \mu\alpha^{\varphi}(\tau; \upsilon)} - 1\Big) = 0, \tag{3.4}$$

with initial conditions $\alpha^{\varphi}(0; \upsilon) = -i\upsilon$ and $\beta^{\varphi}(0; \upsilon) = 0$.

Solving for $\alpha^{\varphi}(\triangle; \upsilon)$ from the Bernoulli equation, and $\beta^{\upsilon}(\triangle; \upsilon)$ through integration, we obtain the following:

$$\Psi^{\varphi}(\upsilon) = \mathbb{E}_t(e^{i\upsilon S_{t+\triangle}}) = e^{-\alpha^{\varphi}(\triangle;\upsilon)S_t - \beta^{\varphi}(\triangle;\upsilon)}, \tag{3.5}$$

where

$$\alpha^{\varphi}(\triangle; \upsilon) = \frac{\frac{\sigma^2}{2\kappa}(e^{\kappa\triangle} - 1) - \frac{e^{\kappa\triangle}}{\upsilon} i}{\frac{\sigma^4}{4\kappa^2}(e^{\kappa\triangle} - 1)^2 + \frac{e^{2\kappa\triangle}}{\upsilon^2}}, \tag{3.6}$$

$$\beta^{\varphi}(\triangle; \upsilon) = \kappa\Big(\theta - \frac{\mu\lambda}{\kappa}\Big)\int_0^{\triangle} \alpha^{\varphi}(\tau; \upsilon) d\tau - \lambda \int_0^{\triangle} \Big(\frac{1}{1 + \mu\alpha^{\varphi}(\tau; \upsilon)} - 1\Big) d\tau. \tag{3.7}$$

[1] From here, we drop the $\mathbb{Q}$ for notational convenience

The joint characteristic function between $S_{n\triangle}$ and A_n can be derived by utilising (3.5) and repeating the law of iterated expectation. Hence, following the methodology as outlined in Chung & Wong (2014), we have the joint characteristic function between $S_{n\triangle}$ and $\sum_{j=0}^{n}\omega_j S_{j\triangle}$ under price dynamics (3.1):

$$\begin{aligned}\Psi_{A_n}^{\varphi}(\phi;\gamma) &= \mathbb{E}_0\Big(e^{i\phi S_{n\triangle}+i\gamma\sum_{j=0}^{n}\omega_j S_{j\triangle}}\Big)\\ &= e^{i\Gamma_0^{\varphi}(\triangle;\phi,\gamma)S_0-\sum_{j=0}^{n-1}\beta^{\varphi}(\triangle;\Gamma_{j+1}^{\varphi}(\triangle;\phi,\gamma))},\end{aligned} \tag{3.8}$$

where $\Gamma_j^{\varphi}(\triangle;\phi,\gamma)$ satisfies the following recursive equation:

$$\Gamma_j^{\varphi}(\triangle;\phi,\gamma) = i\Gamma^{\varphi}(\triangle;\Gamma_{j+1}^{\varphi}(\triangle;\phi,\gamma)) + \gamma\omega_j, \tag{3.9}$$

for $j = n-1, n-2, ..., 0$, and starting value $\Gamma_n^{\varphi}(\triangle;\phi,\gamma) = \phi + \gamma\omega_n$.

3.3 Asian Option Pricing with Fourier-Cosine Expansions

The pricing of options under the COS method, as with all numerical integration techniques, follows from the discounted expected payoff approach under the risk-neutral measure $\mathbb{Q}$:

$$v(x,t) = e^{-r(T-t)}\mathbb{E}[v(y,T)|x] = e^{-r(T-t)}\int_{\mathbb{R}} v(y,T)g(y|x)dy, \tag{3.10}$$

where $v(x,t)$ denotes the option value at time t, and r is the interest rate. In addition, x and y are state variables at time t and expiration date T, respectively. Typically, the option's payoff function, v, is known, but its transitional density function $g(y|x)$ is not. Fang & Oosterlee (2008) proposed an approximation of the transition probability, based on (3.1), with a truncated domain $[a,b]$ by a truncated Fourier-cosine series expansion with N_C terms, based on the conditional characteristic function, i.e.:

$$g(y|x) \approx \frac{2}{b-a}\sum_{h=0}^{N_C-1}{}'\,\mathrm{Re}\Big[\psi\Big(\frac{h\pi}{b-a};x\Big)e^{-ih\pi\frac{a}{b-a}}\Big]\cos\Big(h\pi\frac{y-a}{b-a}\Big), \tag{3.11}$$

where $\psi(\nu;x)$ is the conditional characteristic function of $g(y|x)$, and a, b denotes the integration range in the original domain. The integration range $[a,b]$ may also be determined by making use of the cumulants, such that the error of the approximation is within some tolerance level Tol, i.e. $|\int_{\mathbb{R}} g(y|x)dy - \int_a^b g(y|x)dy| < Tol$ (see Fang & Oosterlee, 2008). Finally, replacing the conditional density function in (3.10) with its approximation (3.11), and interchanging the summation and integration, we obtain the COS formula to price an option with payoff $v(x,t)$:

$$\hat{v}(x,t) = e^{-r(T-t)}\sum_{h=0}^{N_C-1}{}'\,\mathrm{Re}\Big[\psi\Big(\frac{h\pi}{b-a};x\Big)e^{-ih\pi\frac{a}{b-a}}\Big]V_h, \tag{3.12}$$

where $\hat{v}(x,t)$ is the approximation of the option value at time t, and

$$V_h := \frac{2}{b-a}\int_a^b v(y,T)\cos\Big(h\pi\frac{y-a}{b-a}\Big)dy \tag{3.13}$$

are the Fourier-cosine coefficients of payoff $v(y,T)$.

Having derived the joint characteristic function between $S_{n\triangle}$ and A_n, we use the COS method to price Asian options. First, consider the contingent claim $v(y,T)=\max(\rho y-k,0)$ at time T, where $k=\rho K$, $\rho=+1$ for calls and $\rho=-1$ for puts. This setup allows for both fixed strike Asian options $\Big(y=\sum_{j=0}^n \omega_j S_{j\triangle}\Big)$ and floating strike Asian options $\Big(y=S_{n\triangle}-\sum_{j=0}^n \omega_j S_{\omega j}\Big)$, as well as for a plain vanilla option $\Big(y=S_{n\triangle}\Big)$. Under the assumption of risk neutrality, with $g_{\rho y}(u)$ as the density of ρy, the arbitrage free price of our option at initial time 0 is:

$$v(y,T;k;\rho)=e^{-r(T-t)}\int_{-\infty}^{\infty}\max(u-k,0)g_{\rho y}(u)du. \tag{3.14}$$

Through expression (3.14), and substituting the joint characteristic function derived in (3.8) into (3.12), we arrive at the COS method for pricing the various arithmetic Asian options mentioned in Table 3.1, where:

$$V_h=\begin{cases}\frac{2}{b-a}\Big(\Pi_{1,h}(K,b)-K\Pi_{2,h}(K,b)\Big) & \text{for calls,}\\ \frac{2}{b-a}\Big(K\Pi_{2,h}(a,K)-\Pi_{1,h}(a,K)\Big) & \text{for puts,}\end{cases} \tag{3.15}$$

where $\Pi_{1,h}$ and $\Pi_{2,h}$ are from the mathematical results below.

Proposition 1. The cosine series coefficients, $\Pi_{1,h}$, of a function $\mathcal{H}(y)=y$ on $[x_1,x_2]\subset[a,b]$ given by,

$$\Pi_{1,h}(x_1,x_2):=\int_{x_1}^{x_2} y\cos\Big(h\pi\frac{y-a}{b-a}\Big)dy, \tag{3.16}$$

and the cosine series coefficients, $\Pi_{2,h}$, of another function $\mathcal{H}(y)=1$ on $[x_1,x_2]\subset[a,b]$ given by,

$$\Pi_{2,h}(x_1,x_2):=\int_{x_1}^{x_2}\cos\Big(h\pi\frac{y-a}{b-a}\Big)dy, \tag{3.17}$$

are both known analytically.

Proof. A straightforward calculation shows that

$$\Pi_{1,k}(x_1,x_2)=\frac{1}{\Big(\frac{h\pi}{b-a}\Big)^2}\Bigg[\cos\Big(h\pi\frac{x_2-a}{b-a}\Big)-\cos\Big(h\pi\frac{x_1-a}{b-a}\Big) +\frac{h\pi}{b-a}\sin\Big(h\pi\frac{x_2-a}{b-a}\Big)x_2-\frac{h\pi}{b-a}\sin\Big(h\pi\frac{x_1-a}{b-a}\Big)x_1\Bigg],$$

and

$$\Pi_{2,k}(x_1,x_2)=\begin{cases}\Big[\sin\Big(h\pi\frac{x_2-a}{b-a}\Big)-\sin\Big(h\pi\frac{x_1-a}{b-a}\Big)\Big]\frac{b-a}{h\pi}, & \text{if } h\neq 0,\\ (x_1-x_2), & \text{otherwise.}\end{cases}$$

3.4 Numerical Results

In this section, a variety of numerical analyses are performed to test the performance of the COS method against its alternative competitor, the FFT method, as a benchmark. In addition to Chung & Wong (2014), which concluded that the FFT outperforms Monte Carlo simulations in terms of both pricing accuracy and efficiency, we further show that the COS method is more efficient, and does not compromise the pricing accuracy.

3.4.1 Truncation Range for COS Method

The error analysis of the COS method, presented in Fang & Oosterlee (2008), has shown that over a well-specified truncation range for the integration in (3.10), the overall error converges either exponentially or algebraically, depending on whether the density function belongs to $\mathbb{C}^{\infty}([a,b] \subset \mathbb{R})$ or has a discontinuity in one of its derivatives, respectively. Such a truncation range, $[a,b]$, may be determined by making use of the n-th cumulant, c_n, of $y = \sum_{j=0}^{n} \omega_j S_{j\triangle}$ (for fixed strike) or $y = S_{n\triangle} - \sum_{j=0}^{n} \omega_j S_{\omega j}$ (for floating strike), as proposed in Fang & Oosterlee (2008):

$$[a,b] := \left[c_1 - L\sqrt{c_2 + \sqrt{c_4}}, \quad c_1 + L\sqrt{c_2 + \sqrt{c_4}} \right] \quad \text{with } L = 10. \tag{3.18}$$

Readers are referred to Fang & Oosterlee (2008) and Fang & Oosterlee (2009) for detailed discussions on the choices of c_n and L.

3.4.2 Comparison of COS Method Against FFT

Our main focus is the performance comparison between the proposed COS method and that of the FFT. Constant parameters are utilised in our models to ease the demonstration. In particular, we make use of the same constant parameters as specified in Chung & Wong (2014). These parameter values are summarised in Table 3.2. In addition, our comparison will be performed on both fixed and floating strike Asian options, for the different frequencies of monitoring dates n, where $n = 4, 12, 26, 52$ and 252. These dates correspond to quarterly, monthly, biweekly, weekly and daily monitoring setups. The resulting relative price differences between the COS method and FFT are shown in Figure 3.1.

Our numerical result shows a relative pricing difference (or error) in the order between 0.006% and 0.02% for the fixed strike Asian options, and between 0.02% and 0.2% for the floating strike, indicating a negligible difference between the two approaches. These results, together with that of Chung & Wong (2014), suggest a high pricing accuracy

Table 3.2: Parameter values for the numerical analyses

$S_0 = 1$	$\kappa = 0.3$
$\theta = 1.05$	$\sigma = 0.7$
$\lambda = 5$	$\mu = 0.1$
$T = 1$	$r = 0.04$
$\omega_j = \frac{1}{n+1}$	

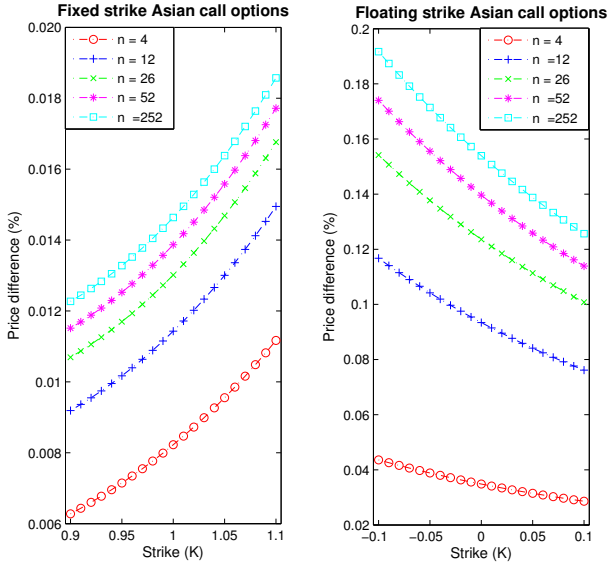

Figure 3.1: Price difference between COS and FFT methods for Fixed and Floating strike Asian calls

for both the FFT and COS method in pricing arithmetic Asian options. In terms of pricing efficiency, the COS method dominates that of the FFT. Using a computer equipped with a 3.5GHz quad-core Intel Core i7-4850HQ processor, the COS method takes only between 0.01-1s (ranging between $\frac{1}{15}$ and $\frac{1}{2}$ of the time required by the FFT method) to obtain the option prices, depending on the choice of monitoring dates and integration grid sizes, N. It should be highlighted that, while N denotes the grid sizes used for the FFT method, we allow the number of terms in the Fourier-cosine series expansion to match such a grid size (i.e. $N_C = N$) for a direct comparison between the two methods.

Table 3.3: cpu time differences and relative error between COS method and FFT

n = 52	N	128	256	512	1024
COS	sec	0.0244	0.0431	0.0524	0.1207
FFT	sec	0.2458	0.2551	0.2712	0.2954
relative error		1.6754e-05	1.6840e-05	1.6838e-05	1.6838e-05

In Table 3.3, the cpu time and relative error information, comparing the COS and the FFT method, are presented for the pricing of Asian options. For this particular example, we price for fixed strike arithmetic Asian options, with weekly monitoring dates ($n = 52$), and grid sizes ranging from $N = 128$ to $N = 1024$. The COS method uses significantly less cpu time to obtain the option prices, while at the same time, produces equal level of accuracy to that of the FFT (evident from the negligible relative pricing errors).

Table 3.4 displays the cpu time comparison and the relative error information between the COS and FFT methods in calculating Asian option prices. In this example, we calculate for fixed strike Asian options across the different monitoring dates, ranging from quarterly ($n = 4$) to daily ($n = 252$), and grid size $N = 4096$. The COS method once again proves to be superior to the alternative FFT method in terms of efficiency for all monitoring dates (with $N_C = N$). However, the efficiency improvement is of a decreasing rate as we increase the monitoring frequency. Such patterns are not dissimilar to the results of Fang & Oosterlee (2008), whereby the COS method's rate of efficiency improvements was shown to decrease as the number of grid sizes, N, is increased.

Table 3.4: CPU time differences and relative error between COS method and FFT

	n	4	12	26	52	252
COS	sec	0.0396	0.0669	0.1063	0.1648	0.7042
FFT	sec	0.2745	0.3135	0.3675	0.4625	1.3861
relative error		9.7365e-06	1.3956e-05	1.5865e-05	1.6838e-05	1.7689e-05

Finally, when dealing with more involved stochastic price processes, such as (3.1) above, as well as (4.1) and (5.1) in the sequel, it should be noted that the resulting option prices across different Fourier-based techniques may converge even when the characteristic function may be inaccurate. Hence, we provide a price comparison for FFT- and COS-based Asian option prices when benchmarked to the result of various Monte Carlo (MC) simulations over a range of strike values. Tables 3.5 and 3.6 reports the % price differences (relative errors) for fixed strike arithmetic Asian options between the FFT / COS method and that of a MC simulation with 100,000 sample paths (which requires just over 28 seconds of CPU time). Notably, the maximum relative error across the range of strikes observed is less than 1%. Relative errors for floating strike arithmetic Asian option prices between the FFT / COS methods and a MC simulation of 100,000 paths is presented in Tables 3.7 and 3.8. Similar to the case fo fixed strikes, we observe that the maximum relative error is just over 1%.

Table 3.5: Relative prices differences between COS / FFT and MC fixed strike Asian call option prices (for $N_C = N = 4096$ and $n = 4$)

Strike	MC price	COS price	relative error	FFT price	relative error
0.9	0.1950	0.1937	0.0069	0.1937	0.0068
0.95	0.1684	0.1678	0.0036	0.1678	0.0036
1	0.1445	0.1446	0.0006	0.1446	0.0007
1.05	0.1241	0.1240	0.0012	0.1240	0.0011
1.1	0.1054	0.1057	0.0022	0.1057	0.0023

Table 3.6: Relative prices differences between COS / FFT and MC fixed strike Asian call option prices (for $N_C = N = 4096$ and $n = 252$)

Strike	MC price	COS price	relative error	FFT price	relative error
0.9	0.2029	0.2030	0.0004	0.2030	0.0005
0.95	0.1761	0.1773	0.0066	0.1773	0.0068
1	0.1530	0.1540	0.0065	0.1541	0.0066
1.05	0.1344	0.1332	0.0092	0.1332	0.0090
1.1	0.1139	0.1146	0.0058	0.1146	0.0060

Table 3.7: Relative prices differences between COS / FFT and MC floating strike Asian call option prices (for $N_C = N = 4096$ and $n = 4$)

Strike	MC price	COS price	relative error	FFT price	relative error
-0.10	0.1909	0.1916	0.0037	0.1917	0.0041
-0.05	0.1665	0.1666	0.0004	0.1667	0.0008
0	0.1448	0.1442	0.0038	0.1443	0.0035
0.05	0.1244	0.1244	0.0001	0.1244	0.0005
0.1	0.1080	0.1069	0.0103	0.1069	0.0100

Table 3.8: Relative prices differences between COS / FFT and MC floating strike Asian call option prices (for $N_C = N = 4096$ and $n = 252$)

Strike	MC price	COS price	relative error	FFT price	relative error
-0.10	0.2003	0.2006	0.0019	0.2010	0.0038
-0.05	0.1752	0.1757	0.0028	0.1335	0.0046
0	0.1546	0.1533	0.0079	0.1536	0.0064
0.05	0.1338	0.1333	0.0033	0.1335	0.0019
0.1	0.1153	0.1155	0.0015	0.1157	0.0027

3.4.3 Price Sensitivity to Changes in Model Parameters

Apart from our comparison on the pricing efficiency and accuracy, we evaluate the effect of parameter value changes on the Asian option price computed from the proposed COS method. This falls particularly in line with the analyses performed by Chung & Wong (2014). Inclusion of such analyses also provide further robust evidence on the stability of the COS method in comparison to the FFT if time-dependent parameters were advocated. The three parameters observed are (i) the jump intensity, (ii) the mean level, and (iii) the asset volatility of the proposed commodity price dynamic (3.1).

We plot the Asian call option prices against different values of the three parameters mentioned above. Parameter values in Table 3.2 are used as a base case and altered within a specified range to find different Asian option prices. The resulting prices are calculated using the COS method, with $K = S_0$ for fixed strike Asian options and $K = 0$ for the floating. Prices against each changing parameter are then plotted in Figures 3.2-3.4.

From Figure 3.2, it is clear that both fixed and floating Asian call option prices are increasing functions of jump intensity, λ. Such result may be deemed valid as an increase in jump intensity also introduces more variability into the underlying asset price dynamic, which in turn increases the value of the options.

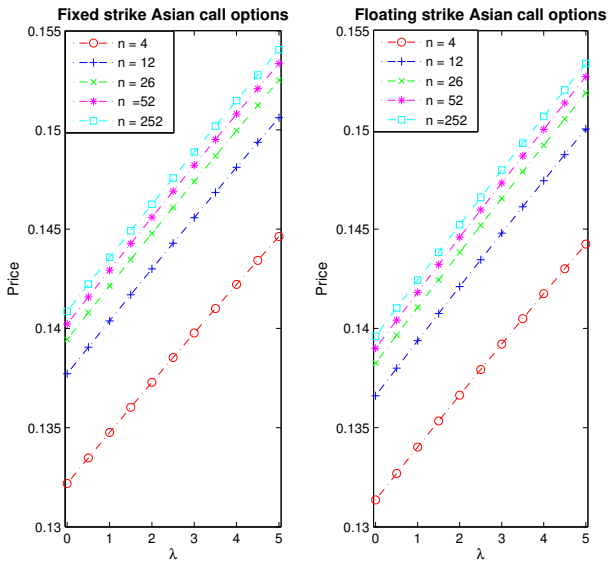

Figure 3.2: Asian option price against jump intensity under COS method

Long term mean levels should also have a positive relationship with call option prices, as greater long term mean levels implies asset prices will tend to remain at a higher level. Such notion is evident in Figure 3.3, which shows a higher Asian call option

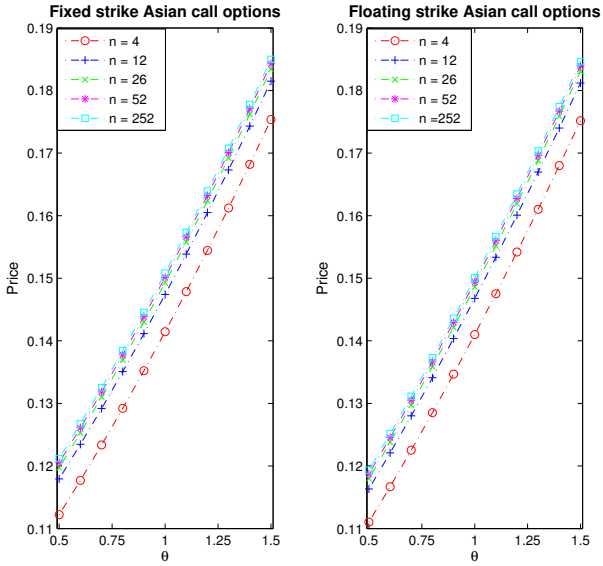

Figure 3.3: Asian option price against mean levels under COS method

prices for greater long term mean levels (the opposite will hold for puts). Finally, Figure 3.4 confirms the trivial notion of a positive relationship between volatility and option prices. The greater the volatility the more the variability there is in the asset price, and thus the greater the option value.

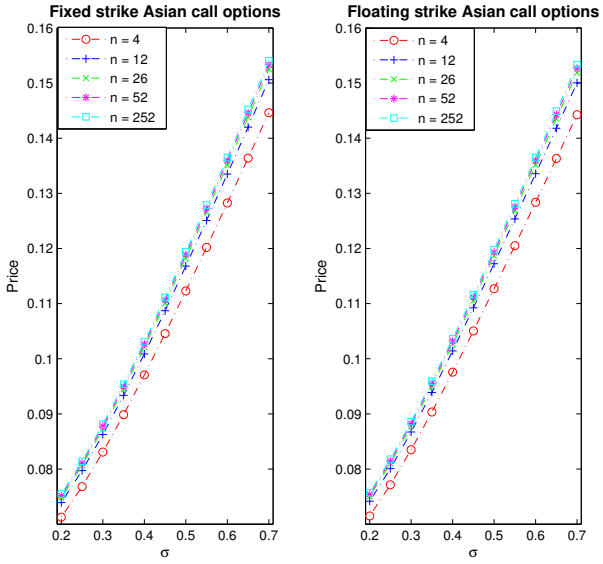

Figure 3.4: Asian option price against asset volatility under COS method

It is also worthwhile emphasising that the above results are consistent with the FFT case presented in Chung & Wong (2014), further reinforcing the stability in the pricing accuracy of the COS method to that of the alternative FFT. When the resulting COS

prices are compared to the FFT as a benchmark, the relative errors (or price differences) were also found to be negligible (not dissimilar to that of Figure 3.1).

Finally, in Figure 3.5 we present the rate of convergence of the COS Asian option prices when monitoring frequencies are increased. Both fixed and floating strike Asian option prices tend to converge or stabilise for weekly monitoring frequencies and above.

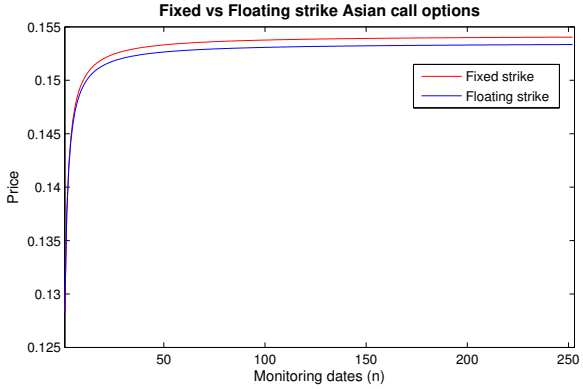

Figure 3.5: Asian option price under COS method for different monitoring dates

3.5 Summary

In this chapter, we proposed the pricing of arithmetic Asian options with the Fourier-cosine method. In particular, we assume a mean reverting jump diffusion process in modelling the underlying commodity price dynamics. Our main focus lies in the investigation of the efficiency and accuracy of the COS method in comparison to the widely accepted FFT. The COS method were shown through our numerical analyses to be more efficient than the benchmark FFT, while producing an equal level of accuracy. Such results are also of particular significance to high frequency traders in search of a better tradeoff between pricing accuracy and efficiency, and a superior method to that of the currently preferred FFT.

To further demonstrate the stability of the COS method, investigations on the price sensitivity to different underlying parameters were conducted. The results presented in this chapter further support the use of jumps in the price dynamic, and the inclusion of time-varying mean level and asset volatility. In addition, it demonstrated the stability of the COS method in comparison to the alternative FFT when underlying parameters vary. Further work may include the investigation of COS method pricing of early exercise Asian options of the arithmetic type, in particular, with mean reversion and jumps inherent in the underlying price dynamics.

Chapter 4

Efficient Option Pricing under the Double Jump Model with Stochastic Volatility and Stochastic Interest Rate Based on Fourier-Cosine Expansions

4.1 Introduction

In this chapter, we investigate the efficient valuation of options when the underlying asset follows a double exponential jump diffusion model with both stochastic volatility and stochastic interest rates. In particular, we focus on the novel pricing method under Fourier-cosine (COS) expansions proposed by Fang & Oosterlee (2008). We then compare our pricing results to that of its close competitor, the well known fast Fourier transform technique (FFT) of Carr & Madan (1999), for pricing efficiency. Both the above-mentioned numerical integration techniques have the advantage whereby the option is priced after transformation to the Fourier domain. This is particularly useful in our assumed model framework, as the density function of the asset price process is not readily available, but its characteristic function is. Hence, instead of utilising the density function of price process to value the conditional expectation of the option payoff in the original pricing domain, we may price options more efficiently by making use of its characteristic function under the Fourier domain. Our numerical experiment demonstrates that not only is the computational time significantly reduced under the COS method, the resulting option prices are also more accurate than the FFT when

benchmarked to the existing closed-form solution (CF) of Deng (2007), which has been shown to be more efficient than the widely used Monte Carlo simulation by Zhang & Wang (2013). Finally, we analyse the robustness of the COS method by evaluating the resulting error convergence, as well as the price changes under the impact of varying jump intensity and correlation coefficient for a range of strikes and maturities.

4.2 Model Specification and Characteristic Function Derivation

4.2.1 Model Specification

Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space where $\mathcal{F}_t$ is a filtration generated by three Brownian motion process W_t^1, W_t^2 and W_t^3, for $0 \leq t \leq T$. Suppose $\mathbb{Q}$ is the risk-neutral probability under which the asset price process S_t, volatility process v_t and interest rate process r_t are given by the following dynamics:

$$\begin{aligned} dS_t &= (r_t - \lambda\theta)S_t dt + \sigma\sqrt{v_t}S_t dW_t^1 + S_t d\Big(\sum_{i=1}^{N_t}(J_i - 1)\Big), \\ dv_t &= (\kappa_v - \alpha_v v_t)dt + \sigma_v\sqrt{v_t}dW_t^2, \\ dr_t &= (\kappa_r - \alpha_r r_t)dt + \sigma_r\sqrt{r_t}dW_t^3, \end{aligned} \tag{4.1}$$

where $\sigma, \kappa_v, \alpha_v, \sigma_v, \kappa_r, \alpha_r, \sigma_r$ are constants, and $S_0 = s$, $v_0 = v$, $r_0 = r$. Furthermore, W_t^1 and W_t^2 are a pair of correlated Brownian motions such that $dW_t^1 dW_t^2 = \rho dt$. The Brownian process W_t^3 is independent of both W_t^1 and W_t^2. N_t is a Poisson process with constant intensity $\lambda > 0$ and $\theta = E^{\mathbb{Q}}[J - 1]$, where $J = (J_i)_{i\geq 1}$ is a sequence of non-negative random variables that is independent and identically distributed. Hence, $Y = \ln J$ follows an asymmetric double exponential distribution with density

$$f_Z(z) = p[\eta_u e^{-\eta_u z}\mathbf{1}_{(z\geq 0)}] + q[\eta_d e^{\eta_d z}\mathbf{1}_{(z<0)}], \tag{4.2}$$

where p,$q \leq 1$, $p+q = 1$ are the the probability of up-move jump and down-move jump, respectively. $\eta_u > 1$, $\eta_d > 0$ are the mean of positive and negative jumps, respectively. Therefore, $\theta = \frac{p\eta_u}{\eta_u - 1} + \frac{q\eta_d}{\eta_d + 1} - 1$. Finally, we further assume that the process W_t^1, W_t^2, W_t^3 are independent of N_t and J.

The proposed model above is also a general setup of the usual market and is subject to a number of special cases. For instance, (i) the constant interest rate and volatility model of Black-Scholes (with $\lambda = 0$) and jump-diffusion model of Kou (2002); (ii) constant interest rate models with stochastic volatility (with $\lambda = 0$, see Schöbel & Zhu, 1999); (iii) stochastic interest rate with constant volatility model of Kim & Kunitomo (1999), among others.

4.2.2 Derivation of the Characteristic Function

Given the asset price process in (4.1), it is possible to obtain an analytical expression for the characteristic function $\psi(u)$ of $\ln S_\tau$ at terminal time $\tau := T - t$. We first define an explicit expression for the moment generating function (MGF) of the log-asset price at time τ under the risk-neutral measure $\mathbb{Q}$,

$$M_{\ln S_\tau}(\Phi) = E^{\mathbb{Q}}[e^{\Phi \ln S_\tau}]. \tag{4.3}$$

It then follows that the complex-valued characteristic function is given by $\psi(u) = M_{\ln S_\tau}(iu)$. It is worthwhile mentioning that $M_{\ln S_\tau}(\Phi)$ can also be interpreted as a contingent claim with payoff $e^{r\tau+\Phi \ln S_\tau}$ at maturity τ.

By solving the resulting PIDE after applying the Itô formula to (4.1) above, one can conclude that the moment generating function of the mean reverting process with stochastic volatility, jump and stochastic interest rate is (Deng, 2007):

$$M_{\ln S_\tau}(\Phi) = A(\Phi,\tau)e^{B(\Phi,\tau)+C(\Phi,\tau)+D(\Phi,\tau)+E(\Phi,\tau)+F(\Phi,\tau)v+\Phi \ln S_t}, \tag{4.4}$$

where

$$\begin{aligned}
A(\Phi,\tau) &= \left\{\frac{2\xi}{2\xi+(\alpha_v-\xi-\rho\sigma\sigma_v\Phi)(1-e^{-\xi\tau})}\right\}^{\frac{2\kappa_v}{\sigma_v^2}},\\
B(\Phi,\tau) &= \frac{\kappa_v}{\sigma_v^2}\left\{(\alpha_v-\xi-\rho\sigma\sigma_v\Phi)\tau\right\}+\lambda(\vartheta-\theta\Phi-1)\tau,\\
C(\Phi,\tau) &= r\left\{\frac{2(1-e^{-\xi_1\tau})}{2\xi_1+(\alpha_r-\xi_1)(1-e^{-\xi_1\tau})}-\frac{2(1-\Phi)(1-e^{-\xi_2\tau})}{2\xi_2+(\alpha_r-\xi_2)(1-e^{-\xi_2\tau})}\right\},\\
D(\Phi,\tau) &= \frac{\kappa_r}{\sigma_r^2}\left\{(\alpha_r-\xi_2)\tau+2\ln\frac{2\xi_2}{2\xi_2+(\alpha_r-\xi_2)(1-e^{-\xi_2\tau})}\right\},\\
E(\Phi,\tau) &= \frac{-\kappa_r}{\sigma_r^2}\left\{(\alpha_r-\xi_1)\tau+2\ln\frac{2\xi_1}{2\xi_1+(\alpha_r-\xi_1)(1-e^{-\xi_1\tau})}\right\},\\
F(\Phi,\tau) &= \frac{-\sigma^2(1-e^{-\xi\tau})\Phi(1-\Phi)}{2\xi+(\alpha_v-\xi-\rho\sigma\sigma_v\Phi)(1-e^{-\xi\tau})},\\
\theta &= \frac{p\eta_u}{\eta_u-1}+\frac{q\eta_d}{\eta_d+1}-1,\\
\vartheta &= \frac{p\eta_u}{\eta_u-\Phi}+\frac{q\eta_d}{\eta_d+\Phi},\\
\xi &= \sqrt{(\alpha_v-\rho\sigma\sigma_v\Phi)^2+\sigma^2\sigma_v^2\Phi(1-\Phi)},\\
\xi_1 &= \sqrt{\alpha_r^2+2\sigma_r^2},\\
\xi_2 &= \sqrt{\alpha_r^2+2\sigma_r^2(1-\Phi)}.
\end{aligned}$$

By definition, the characteristic function $\psi(u) = M(iu)$, hence, we have

$$\psi(u) = A(u,\tau)e^{B(u,\tau)+C(u,\tau)+D(u,\tau)+E(u,\tau)+F(u,\tau)v+iu \ln S_t}, \tag{4.5}$$

where

$$
\begin{aligned}
A(u,\tau) &= \left\{\frac{2\xi}{2\xi+(\alpha_v-\xi-iu\rho\sigma\sigma_v)(1-e^{-\xi\tau})}\right\}^{\frac{2\kappa_v}{\sigma_v^2}},\\
B(u,\tau) &= \frac{\kappa_v}{\sigma_v^2}\left\{(\alpha_v-\xi-iu\rho\sigma\sigma_v)\tau\right\}+\lambda(\vartheta-iu\theta-1)\tau,\\
C(u,\tau) &= r\left\{\frac{2(1-e^{-\xi_1\tau})}{2\xi_1+(\alpha_r-\xi_1)(1-e^{-\xi_1\tau})}-\frac{2(1-iu)(1-e^{-\xi_2\tau})}{2\xi_2+(\alpha_r-\xi_2)(1-e^{-\xi_2\tau})}\right\},\\
D(u,\tau) &= \frac{\kappa_r}{\sigma_r^2}\left\{(\alpha_r-\xi_2)\tau+2\ln\frac{2\xi_2}{2\xi_2+(\alpha_r-\xi_2)(1-e^{-\xi_2\tau})}\right\},\\
E(u,\tau) &= \frac{-\kappa_r}{\sigma_r^2}\left\{(\alpha_r-\xi_1)\tau+2\ln\frac{2\xi_1}{2\xi_1+(\alpha_r-\xi_1)(1-e^{-\xi_1\tau})}\right\},\\
F(u,\tau) &= \frac{-iu(1-iu)\sigma^2(1-e^{-\xi\tau})}{2\xi+(\alpha_v-\xi-iu\rho\sigma\sigma_v)(1-e^{-\xi\tau})},\\
\theta &= \frac{p\eta_u}{\eta_u-1}+\frac{q\eta_d}{\eta_d+1}-1,\\
\vartheta &= \frac{p\eta_u}{\eta_u-iu}+\frac{q\eta_d}{\eta_d+iu},\\
\xi &= \sqrt{(\alpha_v-iu\rho\sigma\sigma_v)^2+iu(1-iu)\sigma^2\sigma_v^2},\\
\xi_1 &= \sqrt{\alpha_r^2+2\sigma_r^2},\\
\xi_2 &= \sqrt{\alpha_r^2+2\sigma_r^2(1-iu)}.
\end{aligned}
$$

4.3 European Option Pricing Using the COS Method

The pricing of a European option under the COS method, as with all numerical integration techniques, follows from the discounted expected payoff approach under the risk-neutral measure $\mathbb{Q}$ (Grzelak et al., 2012):

$$
\begin{aligned}
v(x,t) &= \mathbb{E}\Big[e^{-\int_t^T r_s ds}v(y,T)|\mathcal{F}_t\Big]\\
&= P(\tau)\int_{\mathbb{R}} v(y,T)g(y|x)dy, \qquad (4.6)
\end{aligned}
$$

where $v(x,t)$ denotes the option value at time t, $g(y|x)$ the conditional density, and r is the interest rate. In addition, x and y are state variables at time t and expiry date T, respectively. Finally, the affine structure of the interest rate dynamics in (4.1) allows one to express the price of a zero coupon bond, maturing at time T, as the following exponential affine form of its state variables (Duffie et al., 2000):

$$P(\tau) = A(\tau)e^{-B(\tau)r}, \qquad (4.7)$$

where

$$\begin{aligned} A(\tau) &= \left[\frac{2\xi_1 e^{(\alpha_r+\xi_1)\frac{\tau}{2}}}{(\alpha_r+\xi_1)(e^{\xi_1\tau}-1)+2\xi_1}\right]^{\frac{2\kappa_r}{\sigma_r^2}}, \\ B(\tau) &= \frac{2(e^{\xi_1\tau}-1)}{(\alpha_r+\xi_1)(e^{\xi_1\tau}-1)+2\xi_1}, \\ \xi_1 &= \sqrt{\alpha_r^2+2\sigma_r^2}. \end{aligned}$$

Typically, the option's payoff function, $v(y,T)$, in (4.6) is usually known, however the transitional density function $g(\cdot)$ is not. Fang & Oosterlee (2008) proposed an approximation of the transition probability, under (4.6), by using a truncated Fourier-cosine series expansion with N_C terms and a truncated domain $[a,b]$, based on the conditional characteristic function, i.e.:

$$g(y|x) \approx \frac{2}{b-a}\sum_{h=0}^{N_C-1}{}' \operatorname{Re}\left[\psi\left(\frac{h\pi}{b-a};x\right)e^{-ih\pi\frac{a}{b-a}}\right]\cos\left(h\pi\frac{y-a}{b-a}\right), \tag{4.8}$$

where $\psi(\nu;x)$ is the conditional characteristic function of $g(y|x)$, and a, b denotes the integration boundaries in the original domain. The integration range $[a,b]$ may also be determined by making use of the cumulants (μ_j), such that the error of the approximation is within a user-defined tolerance level, Tol, i.e. $|\int_{\mathbb{R}} h(y|x)dy - \int_a^b h(y|x)dy| < Tol$ (see Fang & Oosterlee, 2008). Finally, replacing the conditional density function in (4.6) with its approximation (4.8), then interchanging the summation and integration, we have the following COS formula to price a European call option with payoff $v(y,T)$:

$$\hat{v}(x,t) = P(\tau)\sum_{h=0}^{N_C-1}{}' \operatorname{Re}\left[\psi\left(\frac{h\pi}{b-a};x\right)e^{-ih\pi\frac{a}{b-a}}\right]V_h \tag{4.9}$$

where $\hat{v}(x,t)$ is the approximation of the call value at time t, and

$$V_h := \frac{2}{b-a}\int_a^b v(y,T)\cos\left(h\pi\frac{y-a}{b-a}\right)dy \tag{4.10}$$

are the Fourier-cosine coefficients of payoff $v(y,T)$.

Let us now define the log-asset state prices by $x := \ln\frac{S_t}{K}$ and $y := \ln\frac{S_T}{K}$, where K denotes the option strike. It follows that the payoff of the European call option, in the log-asset price defined above, can be expressed as

$$v(y,T) \equiv K(e^y - 1). \tag{4.11}$$

Substituting (4.11) into (4.10) above, we obtain the following expression for V_k

$$V_h = \frac{2}{b-a}K[\chi_{1,h}(0,b) - \chi_{2,h}(0,b)], \tag{4.12}$$

where the cosine series coefficients $\chi_{1,h}$ and $\chi_{2,h}$ are

$$\chi_{1,h}(c,d) = \int_c^d e^y \cos\left(h\pi\frac{y-a}{b-a}\right)dy \tag{4.13}$$

and

$$\chi_{2,h}(c,d) = \int_c^d \cos\left(h\pi\frac{y-a}{b-a}\right)dy, \tag{4.14}$$

which can both be solved through straightforward calculus (see *Proposition* 1). Hence,

$$\begin{aligned}\chi_{1,h}(c,d) = \frac{1}{1+\left(\frac{h\pi}{b-a}\right)^2}\Bigg[&e^d\cos\left(h\pi\frac{d-a}{b-a}\right) - e^c\cos\left(h\pi\frac{c-a}{b-a}\right)\\ &+\frac{h\pi}{b-a}e^d\sin\left(h\pi\frac{d-a}{b-a}\right) - \frac{h\pi}{b-a}e^c\sin\left(h\pi\frac{c-a}{b-a}\right)\Bigg],\end{aligned}$$

and $\chi_{2,h}(c,d) = \Pi_{2,h}(c,d)$ as in *Proposition* 1.

Finally, given the characteristic function in (4.5), we can derive the characteristic function of the log-asset price $\ln(S_T/K)$, which can also be factorised to the form of:

$$\psi(\omega;\mathbf{x}) = \psi_0(\omega).e^{i\omega\mathbf{x}} \quad \text{with} \quad \psi_0(\omega) := \psi(\omega;0).$$

As a result, the pricing formula (4.9) can then be simplified to

$$\hat{c}(\mathbf{x},t) = P(\tau)\sum_{h=0}^{N_C-1}{}' \mathrm{Re}\left[\psi_0\left(\frac{h\pi}{b-a}\right)e^{ih\pi\frac{\mathbf{x}-a}{b-a}}\right]\mathbf{V_h}, \tag{4.15}$$

where $\mathbf{V_h} = U_h\mathbf{K}$ with $U_h = \frac{2}{b-a}[\chi_{1,h}(0,b) - \chi_{2,h}(0,b)]$. We can then express (4.15) as

$$\hat{c}(\mathbf{x},t) = P(\tau)\mathbf{K}\cdot\mathrm{Re}\left[\sum_{h=0}^{N_C-1}{}' \psi_0\left(\frac{h\pi}{b-a}\right)U_h\cdot e^{ih\pi\frac{\mathbf{x}-a}{b-a}}\right], \tag{4.16}$$

where $\{\cdot\}$ can be calculated as a matrix-vector multiplication for efficiency if $\mathbf{K}$ and $\mathbf{x}$ are vectors. It is worthwhile mentioning that, since the log-asset prices, $\mathbf{x}$, are independent variables, from (4.16) it is possible to calculate a range of option prices through a single numerical experiment by choosing a vector of strikes, $\mathbf{K}$, as inputs (similar to that of the Carr-Madan FFT approach). Under such conditions, the computation time for a range of options with different strikes can be greatly simplified. We will also demonstrate that even for very small N_C one can obtain high degrees of accuracy in the resulting option prices.

4.4 Numerical Results

In this section, we analyse the performance differences between the FFT and the COS method. In particular, we compare the accuracy in the resulting option prices when

benchmarked to the closed-form solution of Deng (2007). Moreover, we evaluate the efficiency of the two methods by investigating the time required to obtain the option prices. The numerical experiments are conducted in Matlab R2017b on a 3.5GHz quad-core Intel Core i7-4850HQ processor with 16GB RAM. Following Deng (2007), we set the parameters as: $S_0 = 100, \eta_u = \eta_d = 5, p = 0.4, \lambda = 1, \sigma = 0.2, r_0 = 0.05, \kappa_r = 0.035, \alpha_r = 0.4, \sigma_r = 0.095, \alpha_v = 0.3, \kappa_v = 0.6, \sigma_v = 0.1, \rho = -0.25, T = 0.5$. For the COS method we set $N_C = 64$ for the numerical example, and $N = 4096$ (with $\alpha = 1.18$) for the FFT. Table 4.1 compares the pricing accuracy and efficiency between the FFT (see, also, Zhang & Wang, 2013) and the COS methods when benchmarked to the CF solution. The resulting option prices across a range of strikes are presented, as well as the relative error in comparison to the CF price. Two main results may be

Table 4.1: cpu time differences and relative error between COS method and FFT

Option prices					
Strike	CF	FFT	relative error	COS	relative error
90	15.9476	15.8383	6.8537e-03	15.9519	2.7189e-04
95	12.6801	12.6436	2.8785e-03	12.6759	3.2936e-04
100	9.9301	9.8546	7.6031e-03	9.9222	7.9076e-04
105	7.7110	7.6502	7.8848e-03	7.7083	3.5052e-04
110	5.9872	5.9620	4.2090e-03	5.9927	9.2030e-04
115	4.6884	4.6589	6.2921e-03	4.6972	1.8724e-03

observed from the numerical experiment. Firstly, the COS method is more accurate than the alternative FFT approach when benchmarked to the CF solution. Second, the COS method is more efficient (faster with lower error) than the FFT. The FFT takes approximately 0.03 seconds to calculate 100 option prices (i.e. $K = 51, 52, ..., 150$), whereas the COS method only requires approximately 0.0032 seconds. Apart from a lower overall mean absolute error (MAE) at 0.14%, the COS method also produce prices with the maximum absolute error of only 0.55%, whereas the FFT produced prices with absolute errors less than 1%.

The above COS option prices are determined with the truncation range of

$$[a, b] = \left[\mu_1 - L\sqrt{\mu_2 + \sqrt{\mu_4}},\ \ \mu_1 + L\sqrt{\mu_2 + \sqrt{\mu_4}}\right], \tag{4.17}$$

with $L = 10$ as in (2.4). We denote μ_j the j^{th} cumulant of the log-asset price $\ln(S_T/K)$, and are determined by evaluating the j^{th} derivative of the MGF in (4.4) at 0. The 4th cumulants are included here to accommodate the sharp peaks and fat tails in the density function of the assumed double exponential jump price dynamic.

In Table 4.2, we demonstrate the convergence of the mean absolute error for incremental values of N_C and its corresponding computational times. The COS method is efficient and stable for even small values of N_C in comparison to the alternative FFT, which requires large values of N in order to reach adequate level of accuracy. Moreover, unlike the FFT, the selection of the number of terms in the Fourier-cosine series expansion, N_C, is not limited to powers of 2. It is also apparent that at $N_C = 96$, the mean absolute error of COS prices starts to converge, and provides much higher degree of accuracy than the FFT with negligible increments in computational time (< 1 msec.). The exponential rate of error convergence in the COS option prices is consistent with the results shown in Fang & Oosterlee (2008).

Table 4.2: Error Convergence and CPU time of COS method

N_C	40	64	96	120	140
Mean Abs. Error	8.3e-0.2	1.4e-0.3	2.6805e-04	2.6803e-04	2.6803e-04
CPU time (msec.)	1.3	2.0	2.9	3.5	4.3

For interested readers, we provide evidence of option price convergence to increasing truncation parameter L under the COS method in Figure 4.1. Resulting option prices are referenced to the case of $L = 20$. We observe that the prices exhibit exponential convergence to changes in parameter L, and already reaching satisfactory accuracy with $L = 10$. The findings further supports the choice of L in the numerical examples to follow.

In Table 4.3 and 4.4, we evaluate the effect of varying parameters on the COS option prices, such as changes in interest rates over different correlation coefficients or jump intensities, across a set of strikes and option maturities. By keeping our parameters equal to the above-mentioned, except for ρ and T, we compare the option prices for both 6-month ($T = 0.5$) and 3-year ($T = 3$) maturities.

We examine the effect of both positive and negative correlation coefficient on the resulting call option value. Results in Table 4.3 show that negative correlations produced larger price differences between the stochastic interest and fixed interest rate models than the corresponding zero and positive correlation scenarios, with the difference more pronounced for long maturity options (see Figure 4.2). It is also worthwhile emphasising that under a positive correlation the resulting out-of-the-money call prices are higher, and lower for in-the-money calls, in comparison to the negative correlation case. The findings are also presented graphically in Figure 4.3. Such observations are also consistent with Deng's closed-form solutions, further supporting the robustness of the COS method.

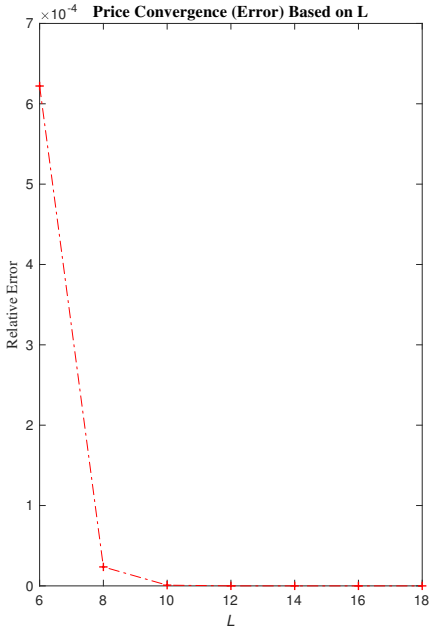

Figure 4.1: COS option price error convergence on increasing truncation range parameter L

Table 4.3: Effects of variability in correlation and interest rate volatility

		$\rho = -0.25$		$\rho = 0$		$\rho = 0.25$	
	Strike	Stoch. IR	Fixed IR	Stoch. IR	Fixed IR	Stoch. IR	Fixed IR
T=0.5	90	15.9519	15.8403	15.9430	15.8314	15.9340	15.8224
	95	12.6759	12.5742	12.6694	12.5678	12.6629	12.5615
	100	9.9222	9.8340	9.9213	9.8333	9.9205	9.8327
	105	7.7083	7.6351	7.7135	7.6404	7.7187	7.6458
	110	5.9927	5.9340	6.0021	5.9434	6.0113	5.9528
	115	4.6972	4.6511	4.7081	4.6621	4.7188	4.6728
T=3	90	36.5000	34.4981	36.4909	34.4916	36.4815	34.4848
	95	34.2075	32.2266	34.2017	32.2240	34.1956	32.2210
	100	32.0604	30.1092	32.0583	30.1108	32.0559	30.1121
	105	30.0528	28.1386	30.0547	28.1446	30.0564	28.1504
	110	28.1781	26.3070	28.1843	26.3175	28.1902	26.3279
	115	26.4298	24.6063	26.4402	24.6214	26.4504	24.6362

Finally, Table 4.4 reveals the effect of varying jump intensities on both short- and long-dated options. We observe that COS option prices increase with the jump intensity λ, and the results are consistent across the various strikes and maturities. This should be expected as an increase in jump intensity introduces more variability in the price of the underlying asset, which raises the value of the option. Similar to the case of correlations, we observe more significant price increases for options with longer time to

Table 4.4: Effects of variability in jump intensity λ

		$\lambda = 2$		$\lambda = 5$	
Maturity	Strike	Stoch. IR	Fixed IR	Stoch. IR	Fixed IR
T=0.1	90	11.6657	11.6586	13.1391	13.1328
	95	6.7907	6.7849	8.8048	8.7995
	100	3.7281	3.7251	5.9443	5.9413
	105	2.5857	2.5849	4.6194	4.6182
	110	2.3284	2.3280	4.0339	4.0331
	115	1.9126	1.9115	3.3933	3.3919
T=5	90	53.3148	50.1757	62.0938	60.2519
	95	51.7181	48.5538	61.1546	59.2656
	100	50.1919	47.0096	60.2416	58.3105
	105	48.7324	45.5387	59.3544	57.3855
	110	47.3361	44.1366	58.4924	56.4898
	115	45.9997	42.7993	57.6550	55.6221

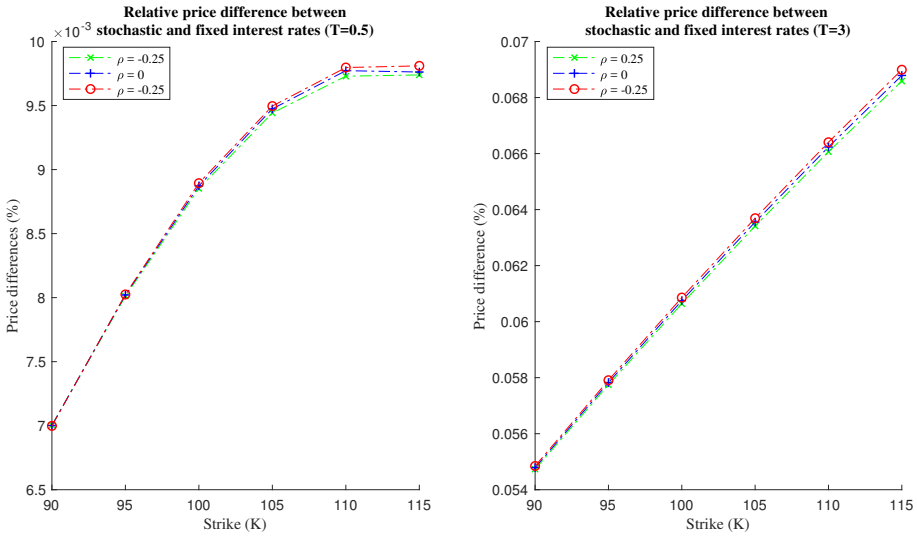

Figure 4.2: Relative Price differences between stochastic and fixed rates across different correlations

maturity. Moreover, our findings show that changes in intensity impose a non-linear effect on the resulting option values across the various strikes, with price differences more pronounced for Out-of-the-Money options.

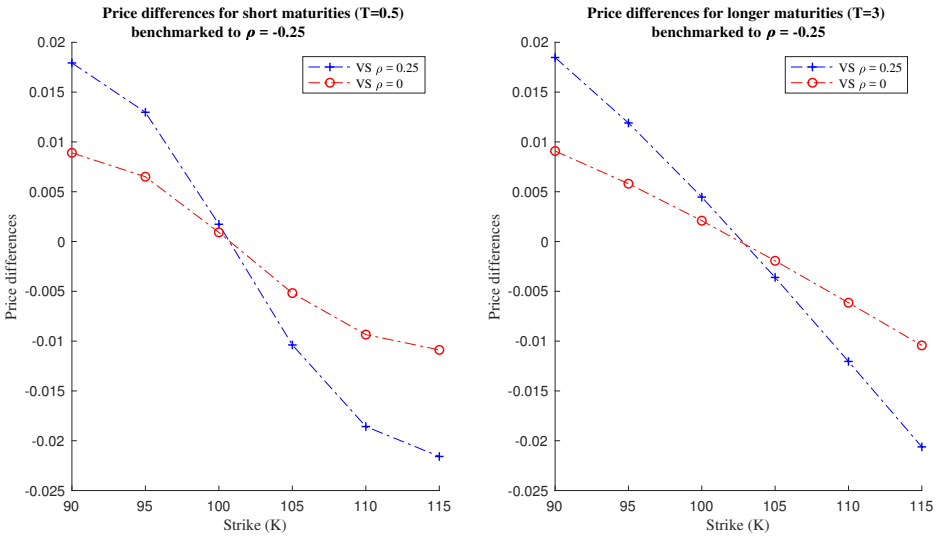

Figure 4.3: Price differences between levels of correlation coefficients across various option moneyness

4.5 Summary

In this chapter, we have demonstrated the efficiency of the COS method in pricing European options under the double exponential jump framework with stochastic volatility and stochastic interest rate. Our numerical results show that not only is the COS method faster than its close competitor the FFT, but it is also more accurate when benchmarked to the closed-form solution. The findings provides support for the use of the COS method even when a semi closed-form solution exists for the option price. Since the solution may require more computation time due to its sophisticated nature.

We demonstrate that the COS method is also robust with very fast error convergence rates in comparison to the alternative FFT. In addition, changes in the COS option price under variabilities in the underlying interest rate and the correlation coefficient has shown significant impact on the option prices. This is particularly true for long dated options, which is also consistent with the corresponding changes in the closed-form solution, further demonstrating model robustness. Moreover, the numerical results from variability in the underlying jump intensity, and its impact on the final option prices, provides evidence in support of stochastic intensity models for further studies. Finally, the efficiency of the COS method, together with its ease of implementation, as demonstrated in this chapter, provides practitioners an adequate method to value option under the double exponential jump framework with stochastic volatility and stochastic interest rate, especially in a high-frequency setting.

Chapter 5

Highly Efficient Power Option Valuation under the Double Jump Framework with Stochastic Volatility and Jump Intensity based on Shannon Wavelet Inverse Fourier Technique

In this chapter, we explore the highly efficient valuation of financial options under a double exponential jump framework with stochastic volatility. Moreover, we allow the jump intensity to be governed by a separate stochastic process. We analyse the efficiency of pricing options under the novel Shannon wavelet inverse Fourier technique (SWIFT) proposed by Ortiz-Gracia & Oosterlee (2016). The resulting accuracy and efficiency in the pricing are then measured against the well-known benchmark Fast-Fourier transform (FFT) method of Carr & Madan (1999), as well as the more efficient alternative proposed by Fang & Oosterlee (2008), the Fourier-cosine (COS) expansions.

Wavelet-based pricing has gained much attraction over the more recent past. Such method may provide both efficient and robust results with exponential error convergence (see Kirkby, 2015; Ortiz-Gracia & Oosterlee, 2016). While Kirkby's method successfully prices options more efficiently than that of the FFT and COS methods, it is restricted by certain drawbacks. One such drawback is the truncation of the infinite integration range to a finite domain when obtaining the density coefficients through

the method of Parseval's identity instead of the Cauchy's integral theorem as in Ortiz-Gracia & Oosterlee (2013). While the exact coefficient expression may be obtained under such a method, the convergence in the number of wavelet terms is only algebraic (see, Ortiz-Gracia & Oosterlee, 2016). The SWIFT method on the other hand does not require prior decisions on integration range truncation. In addition, the decision on domain size will not cause deterioration when approximating the underlying density function, and the method automatically determines the number of terms required in the expansion.

We demonstrate that not only is the SWIFT method more efficient, it is also accurate with exponential error convergence for both call and put valuations. Our investigation extends to the pricing of exotic type options, more precisely, the valuation of power options under the proposed double exponential jump framework. Finally, we conclude by presenting further evidence of model robustness and stability through a price sensitivity analysis, where the significant impact of changing model parameters to the resulting option values are investigated.

5.1 Model Specification and Characteristic Function Derivation

5.1.1 Model Specification

Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a complete probability space on which three Brownian motion processes W_t^s, W_t^v and W_t^λ, for $0 \leq t \leq T$, are defined. Let $\mathcal{F}_t$ be the filtration generated by the Brownian motions and the jump process, and suppose $\mathbb{Q}$ is a risk-neutral probability. The underlying price process S_t, volatility process v_t and the jump intensity λ_t are given by the following dynamics:

$$\begin{aligned} dS_t &= (r - d - \lambda_t \delta) S_t dt + \sqrt{v_t} S_t dW_t^s + S_t (e^Y - 1) dN_t, \\ dv_t &= (\theta_v - \alpha_v v_t) dt + \sigma_v \sqrt{v_t} dW_t^v, \\ d\lambda_t &= (\theta_\lambda - \alpha_\lambda \lambda_t) dt + \sigma_\lambda \sqrt{\lambda_t} dW_t^\lambda, \end{aligned} \tag{5.1}$$

where $\theta_v, \theta_\lambda, \alpha_v, \alpha_\lambda, \sigma_v, \sigma_\lambda$ are constants, and r and d are the risk-free interest and dividend rates, respectively. N_t is a Poisson process with stochastic jump intensity λ_t, and Y is a random variable denoting the jump size, with the mean jump amplitude given by the expectation $\delta = \mathbb{E}[e^Y - 1]$ under the risk-neutral measure $\mathbb{Q}$. Moreover, dW_t^s and dW_t^v are a pair of correlated Brownian motions with $dW_t^s dW_t^v = \rho dt$, and W_t^λ is a Brownian motion independent of both W_t^s and W_t^v. Finally, we assume that

the jump size Y follows an asymmetric double exponential distribution with density function

$$f(Y) = p[\eta_u e^{-\eta_u Y}\mathbf{1}_{(Y\geq 0)}] + q[\eta_d e^{\eta_d Y}\mathbf{1}_{(Y<0)}], \tag{5.2}$$

where p, $q \geq 0$ are the the probability of up-move and down-move jumps, respectively, with $p+q=1$. Furthermore, $\eta_u > 1$, $\eta_d > 0$ are the mean positive and negative jumps, respectively. Therefore, $\delta = \frac{p\eta_u}{\eta_u - 1} + \frac{q\eta_d}{\eta_d + 1} - 1$. We further suppose that the processes W_t^s, W_t^v, W_t^λ are independent of N_t and Y.

5.1.2 Derivation of the Characteristic Function

To obtain the characteristic function for the log-asset price $\ln S_\tau$ with $\tau := T - t$, we first define the moment generating function (MGF) under the risk-neutral measure $\mathbb{Q}$,

$$M(\Phi) = \mathbb{E}[e^{\Phi \ln S_\tau}|\mathcal{F}_t] = e^{-r\tau}\mathbb{E}[e^{r\tau}e^{\Phi \ln S_\tau}|\ln S_t = \ln s, v_t = v, \lambda_t = \lambda]. \tag{5.3}$$

It is clear that the MGF above may also be expressed as a contingent claim with final payoff $e^{r\tau + \Phi X_\tau}$ at maturity τ. Hence, we can solve for the MGF via the Feynman-Kac formula, and thereafter obtain the characteristic function given by $\psi(u) = M(iu)$.

Following the implementation of the generalised Feynman-Kac theorem to the MGF above, and using (5.1) (see Duffie et al., 2000; Huang et al., 2014), we can show that $M(\Phi)$ solves the following partial integro-differential equation (PIDE):

$$\begin{aligned} - \quad & \frac{\partial M}{\partial \tau} + \Big(r - d - \lambda\delta - \frac{v}{2}\Big)\frac{\partial M}{\partial X} + \frac{1}{2}v\frac{\partial^2 M}{\partial X^2} + (\theta_v - \alpha_v v)\frac{\partial M}{\partial v} + \frac{1}{2}\sigma_v^2 v\frac{\partial^2 M}{\partial v^2} \\ & + \rho\sigma_v v\frac{\partial^2 M}{\partial X \partial v} + (\theta_\lambda - \alpha_\lambda \lambda)\frac{\partial M}{\partial \lambda} + \frac{1}{2}\sigma_\lambda^2 \lambda \frac{\partial^2 M}{\partial \lambda^2} \\ & + \lambda \int_{-\infty}^{\infty} [M(X+Y) - M(X)] f(Y) dY = 0. \end{aligned} \tag{5.4}$$

We can postulate a solution of the form:

$$M(\Phi) = e^{\Phi(r-d)\tau + A(\Phi,\tau) + B(\Phi,\tau) + C(\Phi,\tau)v + D(\Phi,\tau)\lambda + \Phi \ln s}, \tag{5.5}$$

with boundary conditions $A(\Phi, 0) = 0$, $B(\Phi, 0) = 0$, $C(\Phi, 0) = 0$, $D(\Phi, 0) = 0$ to solve the above PIDE, and obtain the characteristic function.

We first consider the integral term in (5.4):

$$\begin{aligned} \lambda \int_{-\infty}^{\infty} [M(X+Y) - M(X)] f(Y) dY &= \lambda \int_{-\infty}^{\infty} \mathbb{E}[e^{\Phi X}]\mathbb{E}[e^{\Phi Y} - 1] f(Y) dY \\ &= \lambda M(\Phi)\delta_1 \end{aligned} \tag{5.6}$$

where $\delta_1 = \frac{p\eta_u}{\eta_u - \Phi} + \frac{q\eta_d}{\eta_d + \Phi} - 1$. Substituting equations (5.5) and (5.6) into PIDE (5.4) we obtain an equation that holds for all τ, $\ln S_\tau$, v and λ. Hence, we can reduce the

problem to solving the following system of five, much simpler, ordinary differential equations:

$$\frac{\partial A(\Phi,\tau)}{\partial \tau} + \frac{\partial B(\Phi,\tau)}{\partial \tau} = \theta_v C(\Phi,\tau) + \theta_\lambda D(\Phi,\tau), \tag{5.7}$$

$$\frac{\partial C(\Phi,\tau)}{\partial \tau} = \frac{1}{2}(\Phi^2 - \Phi) - (\alpha_v - \rho\sigma_v\Phi)C(\Phi,\tau) + \frac{1}{2}\sigma_v^2 C^2(\Phi,\tau), \tag{5.8}$$

$$\frac{\partial D(\Phi,\tau)}{\partial \tau} = (\delta_1 - \delta\Phi) - \alpha_v D(\Phi,\tau) + \frac{1}{2}\sigma_\lambda^2 D^2(\Phi,\tau), \tag{5.9}$$

where

$$\frac{\partial A(\Phi,\tau)}{\partial \tau} = \theta_v C(\Phi,\tau), \tag{5.10}$$

$$\frac{\partial B(\Phi,\tau)}{\partial \tau} = \theta_\lambda D(\Phi,\tau). \tag{5.11}$$

From the above system of ODE's, we can solve for (5.9) followed by (5.11). Thereafter, by analogy, we can solve for $C(\Phi,\tau)$ and $A(\Phi,\tau)$, respectively. To solve for (5.9), we will require a particular solution from which to derive the general solution for the Riccati equation. Hence, we make the following substitution:

$$D(\Phi,\tau) = \frac{-2w'(\tau)}{\sigma_\lambda^2 w}. \tag{5.12}$$

Substituting (5.12) into (5.9) above and simplifying we obtain the following ODE:

$$w''(\tau) + \alpha_\lambda w'(\tau) - \frac{1}{2}\sigma_\lambda^2(\delta\Phi - \delta_1)w(\tau) = 0, \tag{5.13}$$

which has a general solution of the form:

$$w(\tau) = U_1 e^{\frac{1}{2}\zeta_- \tau} + U_2 e^{\frac{1}{2}\zeta_+ \tau}, \tag{5.14}$$

where $\zeta_\pm = \varphi \mp \alpha_\lambda$ and $\varphi = \sqrt{\alpha_\lambda^2 + 2\sigma_\lambda^2(\delta\Phi - \delta_1)}$. Both U_1 and U_2 are constants to be determined from the initial conditions $w(0) = U_1 + U_2$ and $w'(0) = 0$, since $D(\Phi,0) = 0$. Solving for U_1 and U_2, we obtain the following solution:

$$U_1 = \frac{w(0)\zeta_+}{2\varphi} \quad \text{and} \quad U_2 = \frac{w(0)\zeta_-}{2\varphi}.$$

Substituting U_1 and U_2 into (5.12) above, we find the exact solution to (5.9):

$$D(\Phi,\tau) = 2(\delta\Phi - \delta_1)\left\{\frac{1 - e^{\varphi\tau}}{\zeta_+ + \zeta_- e^{\varphi\tau}}\right\}. \tag{5.15}$$

Therefore, from (5.11), together with (5.12) and the solution in (5.15), we obtain:

$$\begin{aligned} B(\Phi,\tau) &= \theta_\lambda \int_0^\tau D(\Phi,s)ds \\ &= -\frac{2\theta_\lambda}{\sigma_\lambda^2}\int_0^\tau \frac{w'(s)}{w(s)}ds \end{aligned}$$

$$
\begin{aligned}
&= -\frac{2\theta_\lambda}{\sigma_\lambda^2}\ln\left[\frac{w(\tau)}{w(0)}\right] \\
&= -\frac{\theta_\lambda}{\sigma_\lambda^2}\left\{\zeta_-\tau + 2\ln\left[\frac{\zeta_+ + \zeta_- e^{\varphi\tau}}{2\varphi}\right]\right\}. \qquad (5.16)
\end{aligned}
$$

Finally, solving for ODE's $A(\Phi,\tau)$ and $C(\Phi,\tau)$ by analogy, we obtain the following solutions:

$$
\begin{aligned}
A(\Phi,\tau) &= -\frac{\theta_v}{\sigma_v^2}\left\{\gamma_-\tau + 2\ln\left[\frac{\gamma_+ + \gamma_- e^{\vartheta\tau}}{2\vartheta}\right]\right\}, \\
C(\Phi,\tau) &= (\Phi - \Phi^2)\left\{\frac{1 - e^{\vartheta\tau}}{\gamma_+ + \gamma_- e^{\vartheta\tau}}\right\}.
\end{aligned}
$$

where

$$
\begin{aligned}
\gamma_\pm &= \pm(\rho\sigma_v\Phi - \alpha_v) + \vartheta, \\
\vartheta &= \sqrt{(\rho\sigma_v\Phi - \alpha_v)^2 + \sigma_v^2(\Phi - \Phi^2)}.
\end{aligned}
$$

We conclude that the expression for characteristic function, $\psi(u) = M(iu)$, of the proposed price process (5.1) is:

$$
\psi(u) = e^{iu(r-d)\tau + A(u,\tau) + B(u,\tau) + C(u,\tau)v + D(u,\tau)\lambda + iu\ln s}, \qquad (5.17)
$$

where

$$
\begin{aligned}
A(u,\tau) &= -\frac{\theta_v}{\sigma_v^2}\left\{\gamma_-\tau + 2\ln\left[\frac{\gamma_+ + \gamma_- e^{\vartheta\tau}}{2\vartheta}\right]\right\}, \\
B(u,\tau) &= -\frac{\theta_\lambda}{\sigma_\lambda^2}\left\{\zeta_-\tau + 2\ln\left[\frac{\zeta_+ + \zeta_- e^{\varphi\tau}}{2\varphi}\right]\right\}. \\
C(u,\tau) &= (iu - u^2)\left\{\frac{1 - e^{\vartheta\tau}}{\gamma_+ + \gamma_- e^{\vartheta\tau}}\right\}. \\
D(u,\tau) &= 2(iu\delta - \delta_1)\left\{\frac{1 - e^{\varphi\tau}}{\zeta_+ + \zeta_- e^{\varphi\tau}}\right\}, \\
\gamma_\pm &= \pm(iu\rho\sigma_v - \alpha_v) + \vartheta, \\
\vartheta &= \sqrt{(iu\rho\sigma_v - \alpha_v)^2 + \sigma_v^2(iu + u^2)}, \\
\zeta_\pm &= \varphi \mp \alpha_\lambda, \\
\varphi &= \sqrt{\alpha_\lambda^2 + 2\sigma_\lambda^2(iu\delta - \delta_1)}, \\
\delta &= \frac{p\eta_u}{\eta_u - 1} + \frac{q\eta_d}{\eta_d + 1} - 1, \\
\delta_1 &= \frac{p\eta_u}{\eta_u - iu} + \frac{q\eta_d}{\eta_d + iu} - 1.
\end{aligned}
$$

5.2 European Option Pricing with SWIFT

Let us now define the scaled log-asset state prices by $x := \ln \frac{S_t}{K}$ and $y := \ln \frac{S_T}{K}$, where the scaling factor K denotes the option strike. It then follows that the final payoff of an European option can be expressed as:

$$v(y,T) = \begin{cases} [K(e^y - 1)]^+ & \text{for a call option,} \\ [K(1 - e^y)]^+ & \text{for a put option.} \end{cases} \tag{5.18}$$

The above scaling of the log-asset price will allow for a series of option values to be calculated in the same instance with a vector of target strikes. This contributes to a more efficient process when determining multiple option prices across a range of different strikes, which we shall demonstrate in the sequel.

From (5.18), our pricing equation (2.7) then follows with

$$\begin{aligned} v(x,t) \approx \tilde{v}_1(x,t) &= e^{-r(T-t)} \int_{\mathcal{D}_m} v(y,T) g_m(y|x) dy \\ &= e^{-r(T-t)} \sum_{k=k_1}^{k_2} c_{m,k}(x) \cdot V_{m,k}, \end{aligned} \tag{5.19}$$

where $g_m(y|x) \approx \sum_{k=k_1}^{k_2} c_{m,k}(x)\phi_{m,k}(y)$ is the approximation of density function $g(y|x)$, and the infinite integration truncated to a finite range $\mathcal{D}_m = [k_1/2^m, k_2/2^m]$ with specific k_1 and k_2 values. Finally, $V_{m,k}$ are the payoff coefficients defined as:

$$V_{m,k} := \int_{\mathcal{D}_m} v(y,T)\phi_{m,k}(y) dy.$$

Given the payoff functions in (5.18) and our choice of $\phi_{m,k}(y)$, the payoff coefficients may be expressed as

$$V_{m,k} = \begin{cases} \int_{\mathcal{D}_m \cap [0,+\infty)} K(e^y - 1) \cdot 2^{m/2} \operatorname{sinc}(2^m y - k) dy & \text{for a call option,} \\ \int_{(-\infty,0] \cap \mathcal{D}_m} K(1 - e^y) \cdot 2^{m/2} \operatorname{sinc}(2^m y - k) dy & \text{for a put option.} \end{cases} \tag{5.20}$$

Through the approximation of the cardinal sine function by Gearhart & Schultz (1990), together with the cosine product-to-sum identity of Quine & Abrarov (2013), and by defining $\bar{k}_1 := \max(k_1, 0)$ and $\bar{k}_2 := \min(k_2, 0)$, we can deduce that

$$V_{m,k} \approx V^*_{m,k} := \begin{cases} \frac{K 2^{m/2}}{2^{J-1}} \sum_{j=1}^{2^{J-1}} \left[\Pi_{1,k}\left(\frac{\bar{k}_1}{2^m}, \frac{k_2}{2^m}\right) - \Pi_{2,k}\left(\frac{\bar{k}_1}{2^m}, \frac{k_2}{2^m}\right) \right] & \text{if } k_2 > 0, \\ 0 & \text{if } k_2 \leq 0, \end{cases}$$

for a call option, and

$$V_{m,k} \approx V^*_{m,k} := \begin{cases} \frac{-K 2^{m/2}}{2^{J-1}} \sum_{j=1}^{2^{J-1}} \left[\Pi_{1,k}\left(\frac{k_1}{2^m}, \frac{\bar{k}_2}{2^m}\right) - \Pi_{2,k}\left(\frac{k_1}{2^m}, \frac{\bar{k}_2}{2^m}\right) \right] & \text{if } k_1 < 0, \\ 0 & \text{if } k_1 \geq 0, \end{cases}$$

for a put. The functions $\Pi_{1,k}$ and $\Pi_{2,k}$ are from the mathematical results below.

Proposition 2. The payoff coefficients, $\Pi_{1,k}$, of a function $\mathcal{H}(y) = e^y$ on $[c, d]$ is given by

$$\Pi_{1,k}(c, d) := \int_c^d e^y \cos(\chi_j(2^m y - k))dy, \tag{5.21}$$

and the payoff coefficient, $\Pi_{2,k}$, of another function $\mathcal{H}(y) = 1$ on $[c, d]$ is given by

$$\Pi_{2,k}(c, d) := \int_c^d \cos(\chi_j(2^m y - k))dy, \tag{5.22}$$

where $\chi_j = \frac{2j-1}{2^J}\pi$, are both known analytically.

Proof. Firstly, as with the derivation of the density coefficients, the classical Vieta formula allows us to express the cardinal sine function as the following infinite product of cosine terms (see Gearhart & Schultz, 1990). By truncating the infinite product to a domain with J factors only, the cosine product-to-sum identity allows us to approximate the cardinal sine function as in (2.10) (see Quine & Abrarov, 2013). Hence, by replacing the sinc functions in (5.20) with the approximation above, and defining $\chi_j = \frac{2j-1}{2^J}\pi$, we can obtain the following expressions:

$$\Pi_{1,k}(c, d) := \int_c^d e^y \cos(\chi_j(2^m y - k))dy,$$

and

$$\Pi_{2,k}(c, d) := \int_c^d \cos(\chi_j(2^m y - k))dy.$$

Thereafter, a straight forward calculation of the above integrals show that

$$\begin{aligned}\Pi_{1,k}(c, d) = \frac{\chi_j 2^m}{1 + (\chi_j 2^m)^2}\Big[&e^d \sin(\chi_j(2^m d - k)) - e^c \sin(\chi_j(2^m c - k)) \\ &+ \frac{1}{\chi_j 2^m}\Big(e^d \cos(\chi_j(2^m d - k)) - e^c \cos(\chi_j(2^m c - k))\Big)\Big],\end{aligned}$$

and

$$\Pi_{2,k}(c, d) = \frac{1}{\chi_j 2^m}\Big(\sin(\chi_j(2^m d - k)) - \sin(\chi_j(2^m c - k))\Big).$$

□

It is worthwhile mentioning that the computational complexity in obtaining the payoff coefficients, $V_{m,k}$, may be greatly reduced by avoiding a straight forward calculation, and choosing an appropriate constant value of J over all k instead, i.e., $\tilde{J} := \lceil \log_2(\pi N) \rceil$, where $N := \max(|k_1|, |k_2|)$. We may then apply the FFT to speed up the computation process. For detailed derivations, readers are referred to Appendix B of Ortiz-Gracia & Oosterlee (2016).

Finally, with our characteristic function (5.17), we can obtain the characteristic function of $z = \ln(\frac{S_T}{K})$, which can also be factorised to the form of:

$$\psi_t(\omega; z) = \psi_t(\omega; 0) \cdot e^{i\omega z}.$$

Hence, we can further simplify the SWIFT pricing formula to allow for the pricing of a vector of strikes, $\mathbf{K}$, in a single numerical experiment with

$$\begin{aligned}\tilde{v}(\mathbf{x},t) &= e^{-r(T-t)}\frac{2^m\mathbf{K}}{2^{\bar{J}+\tilde{J}-2}} \\ &\quad\times\sum_{k=k_1}^{k_2}\mathrm{Re}\left[e^{\frac{ik\pi}{2^{\bar{J}}}}\sum_{j=0}^{2^{\bar{J}}-1}\psi_t^*\left(\frac{(2j+1)\pi 2^m}{2^{\bar{J}}};0\right)e^{-\frac{i(2j+1)\pi 2^m}{2^J}\cdot\mathbf{x}}e^{\frac{2\pi ikj}{2^J}}\right]\tilde{V}^*_{m,k},\end{aligned} \tag{5.23}$$

where

$$\tilde{V}^*_{m,k} := \begin{cases}\sum_{j=1}^{2^{J-1}}\left[\Pi_{1,k}\left(\frac{\bar{k}_1}{2^m},\frac{k_2}{2^m}\right)-\Pi_{2,k}\left(\frac{\bar{k}_1}{2^m},\frac{k_2}{2^m}\right)\right] & \text{if } k_2>0,\\ 0 & \text{if } k_2\le 0,\end{cases}$$

for a call option, and

$$\tilde{V}^*_{m,k} := \begin{cases}\sum_{j=1}^{2^{J-1}}\left[\Pi_{1,k}\left(\frac{k_1}{2^m},\frac{\bar{k}_2}{2^m}\right)-\Pi_{2,k}\left(\frac{k_1}{2^m},\frac{\bar{k}_2}{2^m}\right)\right] & \text{if } k_1<0,\\ 0 & \text{if } k_1\ge 0,\end{cases}$$

for a put. Functions $\Pi_{1,k}$ and $\Pi_{2,k}$ are as defined in (5.21) and (5.22), respectively. The above simplification relaxes each $V^*_{m,k}$ from its dependence on the option strikes, K, thus requiring only one computation for $k = k_1, ..., k_2$ via an FFT application. Moreover, utilising only a further FFT algorithm n_K times to compute the payoff coefficients completes the overall calculation. Under the above conditions, the final CPU time required for a range of options with different strikes is greatly simplified from that of a direct calculation for each individual option.

5.3 Power Option Pricing with SWIFT

Similar to the European option scenario, we can define the log-asset state prices by $x := \ln S_t^\beta/K$ and $y := \ln S_T^\beta/K$, where K denotes the strike price and β the constant power term, then the payoff of power options can also be represented by (5.18). Finally, the characteristic function of the scaled log-asset price $y = \ln S_T^\beta/K$, required in (2.11) may be obtained via (5.3) and Proposition 3 below.

Proposition 3. Given the moment generating function, $M_X(\nu)$, of $X = \ln S_T$, then for constants $\beta, K > 0$, the moment generating function of $Y = \ln\left(\frac{S_T^\beta}{K}\right)$ is given by $M_Y(\nu) = e^{-\nu\ln K}M_X(\omega)$, where $\omega = \nu\beta$, and the resulting characteristic function is given by $\psi_Y(\nu) = e^{-i\nu\ln K}\psi_X(\omega)$.

Proof. The MGF of Y is given by $M_Y(\nu) = \mathbb{E}\left(e^{\nu\ln(S_T^\beta/K)}\right) = e^{-\nu\ln K}M_X(\omega)$, since $e^{\nu\beta\ln S_T}$ is a monotonic function of $X = \ln S_T$. Hence, we have

$$M_Y(\nu) = e^{-\nu\ln K}M_X(\omega), \tag{5.24}$$

where $\omega = \nu\beta$. Finally, substituting $i\nu$ for ν in (5.24) above completes the proof. □

Finally, to price a basket of power options in a single numerical experiment, we can apply (5.23) as with the vanilla case, together with the results of Proposition 3.

5.4 Numerical Results

In this section, we analyse the performance of the SWIFT method relative to the alternative FFT and COS methods when evaluating options. In particular, we benchmark the resulting SWIFT option prices to that of the FFT and COS prices, and investigate the resulting accuracy and the computational time required by each method. We demonstrate the superior efficiency of the SWIFT method, as well as the robustness of the model in a separate section through a sensitivity analysis on option prices to changes in model parameters. Finally, we conclude the section by extending our pricing methodology to the valuation of power option.

5.4.1 Comparison of SWIFT to FFT and COS for Plain Vanilla Option Pricing

Following Huang et al. (2014), we consider the parameters at time $t = 0$: $S_0 = 1, v_0 = 0.15, \lambda_0 = 3, r = 0.05, d = 0.05, \theta_v = 0.18, \alpha_v = 0.3, \sigma_v = 0.1, \theta_\lambda = 3, \alpha_\lambda = 5, \sigma_\lambda = 0.3, \rho = -0.25, \eta_u = 33.33, \eta_d = 7.69, T = 0.5$. For both the SWIFT and COS methods, the truncation of the integration range in (3.4) to obtain the pay-off coefficients requires the cumulants of the log-asset price $\ln(\frac{S_T}{K})$. The j^{th} cumulant, μ_j, is defined as the j^{th} derivative of the underlying MGF, as defined in Section 2, evaluated at 0. Following Fang & Oosterlee (2008), we have the following truncated domain to approximate the option prices:

$$[a, b] = \left[\mu_1 - L\sqrt{\mu_2 + \sqrt{\mu_4}},\ \ \mu_1 + L\sqrt{\mu_2 + \sqrt{\mu_4}}\right], \tag{5.25}$$

with $L = 10$.[1] Table 5.1 presents the number of terms required in the SWIFT method for different choice of L values, obtained via (5.25), and the equivalent number of COS terms required for the comparison can also determined as a result. We include the 4^{th} cumulant in (5.25) above to accommodate for heavy tails and excess kurtosis exhibit in the density function of our proposed double exponential jump pricing model. Interested readers may refer to Fang & Oosterlee (2008, 2009) for further discussions on the choice of parameter L and the number of the cumulants to include when determining the truncation range.

[1] The choice of parameter L is further justified in our numerical experiments in the sequel

Table 5.1: Number of terms used in SWIFT calculation based on size of interval determined by L with scale of approximation $m = 5$

L	10	12	14	18	26
k_1	-131	-156	-182	-233	-339
k_2	128	153	179	230	336

Following the guidelines in Ortiz-Gracia & Oosterlee (2016), together with our numerical experiment to demonstrate the rate of error convergence in the next section, we fix the scale of approximation $m = 5$, and set $\bar{J} = 9$ to capture the heavy tails of the distribution, and accommodate for the low rate of decay in the characteristic function, under our proposed jump model. Hence, we also set $N = 4096$ for the FFT, and use $N_C = 260$ for the COS method in order to match the number of terms to the SWIFT calculation (see Table 5.1). Unlike the FFT, both the COS and SWIFT methods do not require a large number of terms in order to reach adequate level of pricing accuracy. Table 5.2 below compares the resulting option prices from the SWIFT method to that of the FFT and COS.

Table 5.2: Call price differences and relative error between SWIFT vs FFT / COS methods

Call option prices					
Strike	SWIFT	FFT	abs. error	COS	abs. error
85	20.3618	20.3618	2.1953e-07	20.3618	2.2522e-07
90	17.3847	17.3847	1.8851e-07	17.3847	1.8852e-07
95	14.7344	14.7344	2.3659e-07	14.7344	2.3660e-07
100	12.4025	12.4025	4.7290e-07	12.4025	4.7291e-07
105	10.3731	10.3730	4.7991e-07	10.3730	4.7992e-07
110	8.62449	8.62449	6.8148e-07	8.62449	6.8149e-07
115	7.13180	7.13180	9.5387e-07	7.13180	9.5391e-07

We observe that the SWIFT method has a high degree of accuracy when benchmarked to the alternative FFT and COS methods. Interestingly, the resulting absolute error on the SWIFT call prices tends to be higher for Out-of-the-Money options in comparison to At-the-Money and In-the-Money options.

We further compare the option prices in Table 5.2 to the numerical results of Monte Carlo simulations for a measure of pricing accuracy. The Monte Carlo method can be implemented once system (5.1) has been discretised (see appendix A for details). Using 100,000 sample paths, and simulation steps of $\Delta t = 0.001$, we obtain 12.4112 (8.6198e-

03 abs. error) for the price of an At-the-Money option (Strike=100). Similar range of errors are obtained for In-the-Money and Out-of-the-Money options (with MAE of 3.6994e-03 across all strikes).

All numerical experiments were conducted on a 3.9GHz quad-core Intel Core i7 machine with 16GB RAM. We confirm that the COS method, which spends just under 2 milliseconds to obtain the option prices, once again proving to be more efficient than the widely accepted FFT method, which requires just over 0.03 seconds. However, both models are outperformed by the highly efficient SWIFT method, which takes a mere 0.6 milliseconds to obtain the option prices.

5.4.2 Error Convergence of Plain Vanilla Option Pricing

We justify our choice of the scale of approximation (m) above through three important numerical experiments to demonstrate the convergence of the approximation errors. Firstly, as per Maree et al. (2017), and our discussions in Chapter 2, we can define a tolerance level, TOL, a priori, and perform an iterative procedure using (2.16) to identify the wavelet scale of approximation, m, such that the mass of the tails of the characteristic function is less than TOL. Secondly, we demonstrate that the rate of convergence of the SWIFT option prices to increasing wavelet scale of approximation m is exponential for both calls and puts, and we illustrate this in Figure 5.2. Finally, recall that a major advantage of utilising Shannon scaling function is the ability to determine the area under the approximated density function with minimal efforts once the coefficients $c_{m,k}$ are calculated. The approximated area may then be determined by:

$$\tilde{H} = \frac{1}{2^{m/2}}\left(\frac{c_{m,k_1}}{2} + \sum_{k_1<k<k_2} c_{m,k} + \frac{c_{m,k_2}}{2}\right).$$

Hence, to complement the numerical experiments mentioned in the first and second instance above, we further investigate the amount of area under the density lost through our approximation.

From Figure 5.1 we observe that, should we set tolerance level as $TOL = 10^{-99}$, say, then the iterative procedure with (2.16) through choices of $m = 0, 1, 2, ...$ would identify $m = 5$ as an adequate wavelet scale of approximation. Notably, a lower choice of tolerance level by the user may deem $m = 3$, say, as an appropriate choice for the scale of approximation.

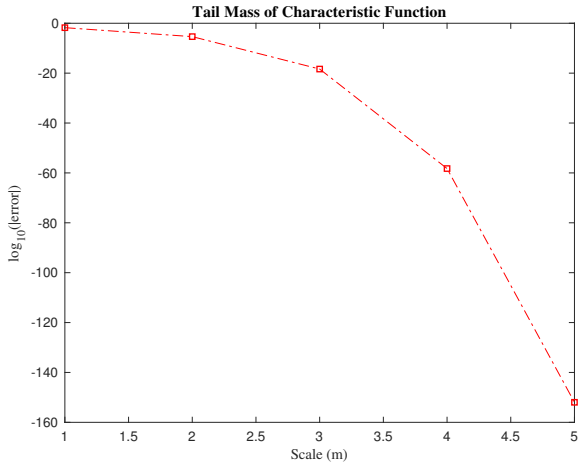

Figure 5.1: Tail Mass of Characteristic Function Not Recovered

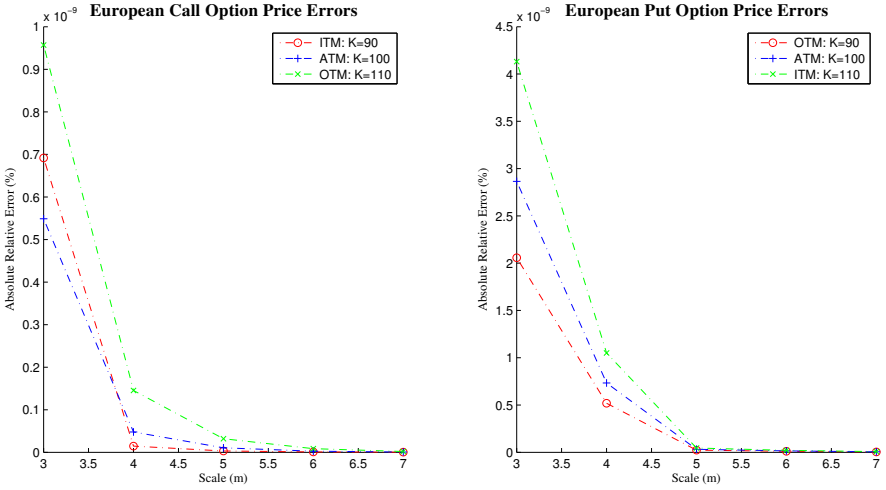

Figure 5.2: SWIFT option price error convergence on increasing scale of approximation (m)

In addition to the above, Figure 5.3 presents the error convergence in the density approximation with the SWIFT method. Evidently, the error convergence in the density approximation is also exponential, with negligible differences for $m > 5$. Moreover, the insignificance of the area not recovered further supports our selection of $L = 10$ in identifying the truncation range. It is worthwhile emphasising that, depending on the error tolerance of the user, lower scales of approximation may already reach satisfactory accuracy, with the benefit of requiring less computational time. For instance, $m = 3$ comes with the tradeoff of requiring just under half the calculation time required for our choice of $m = 5$.

For interested readers, further justifications for our choice of the L parameter, when identifying the truncation range, can be provided by conducting a numerical experiment

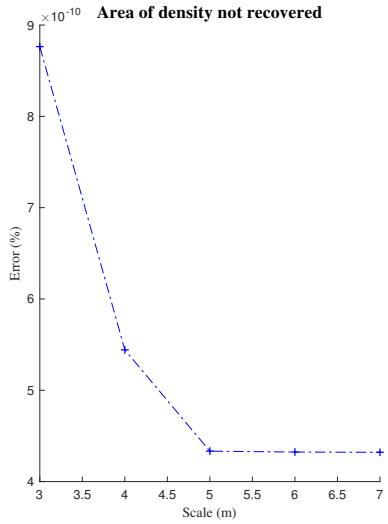

Figure 5.3: Error convergence of density approximation to increasing scale of approximation (m)

to evaluate the price convergence in the parameter L. We demonstrate this in Figure 5.4, which shows robustness for the choice of $L = 10 \sim 12$ across a range of scale of approximation m.

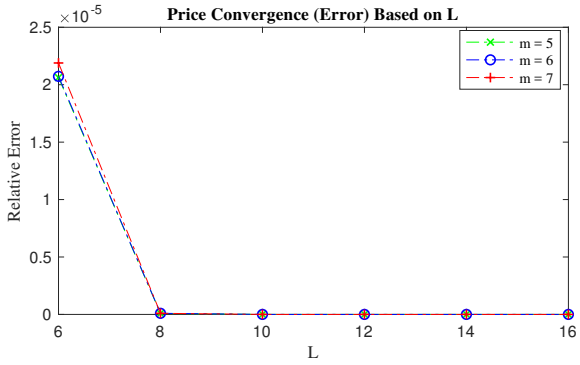

Figure 5.4: SWIFT option price error convergence on increasing truncation range parameter L

5.4.3 Price sensitivity to changes in model parameters for plain vanilla options

We conduct various sensitivity analyses on the resulting SWIFT call and put option values to changes in the underlying model parameters. In particular, resulting SWIFT prices are benchmarked to prices of the alternative COS method. Such analyses provide evidence of robustness and stability of the SWIFT method under different market conditions, as represented by their corresponding model parameters. We shall analyse

the sensitivity of SWIFT option prices to changes in (i) mean-reversion rate of volatility, (ii) mean-reversion rate of jump intensity, (iii) volatility of volatility, (iv) volatility of jump intensity, and finally (v) time to maturity of the options.

Figures 5.5 and 5.6 demonstrate changes in the differences between the SWIFT call (put) prices and the corresponding call (put) values of the COS method. We observe that the price differences are less for lower levels of mean reversion of volatility, and for higher levels of mean reversion for jump intensity. Such observation is consistent across all levels of initial volatility and jump intensity, with errors more pronounced for larger initial parameter values.

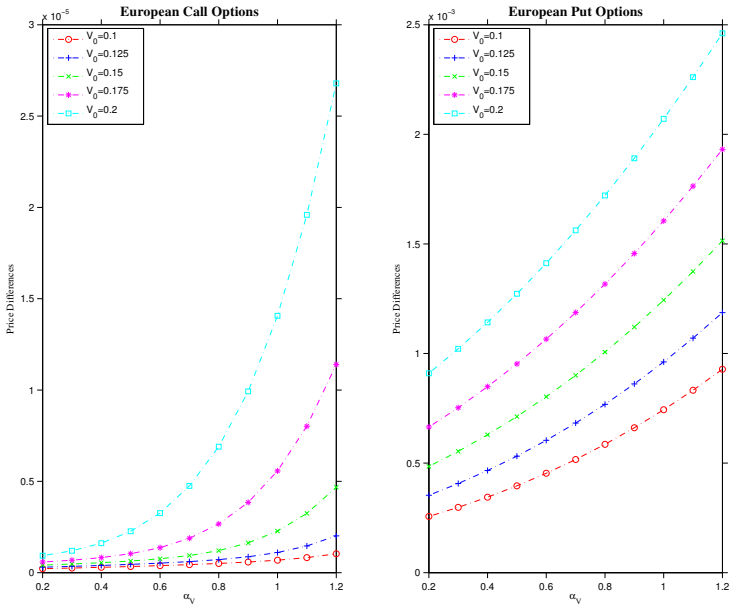

Figure 5.5: Option price sensitivity analysis for changes in mean-reversion of volatility

Similar results are observed for price changes across different volatility of volatility, as well as the volatility of jump intensity (see figures 5.7 and 5.8). Higher levels of volatility, of both volatility and jump intensity, tends to generate more significant price discrepancies between the COS and SWIFT methods. Interestingly, the price differences between the two methods remains stable across different levels of volatility of jump intensity, and over all values of initial jump intensity. Lastly, we analyse the percentage price differences between the SWIFT and COS methods in pricing both European calls and puts for both long and short-dated options, and present our results in Figure 5.9. While the relative error between the two methods are minimal, we observe more significant price differences for long-dated options in comparison to the shorter maturity options.

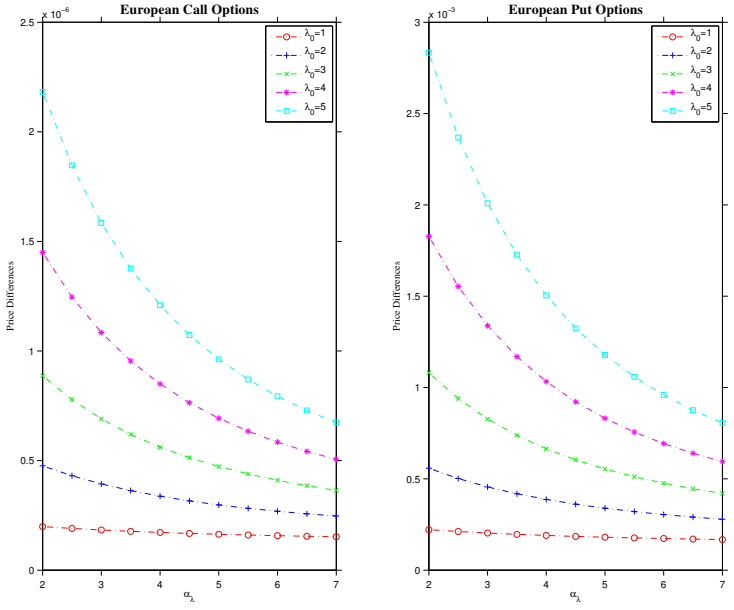

Figure 5.6: Option price sensitivity analysis for changes in mean-reversion of jump intensity

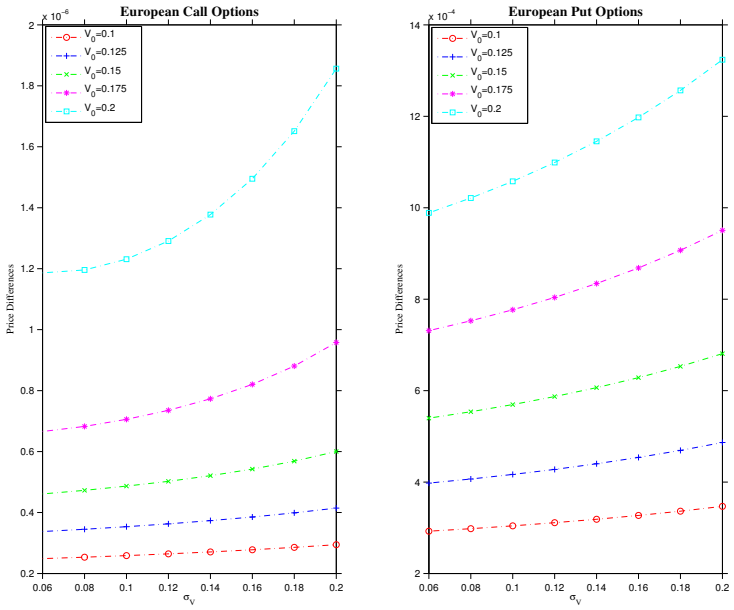

Figure 5.7: Option price sensitivity analysis for changes in volatility of volatility

Overall, we conclude that the SWIFT method is robust, with negligible price differences in comparison to the alternative COS method for both European calls and puts. The methods are consistent even when tested under both low and high model parameter values, with less relative errors under low initial values for volatility and jump intensity.

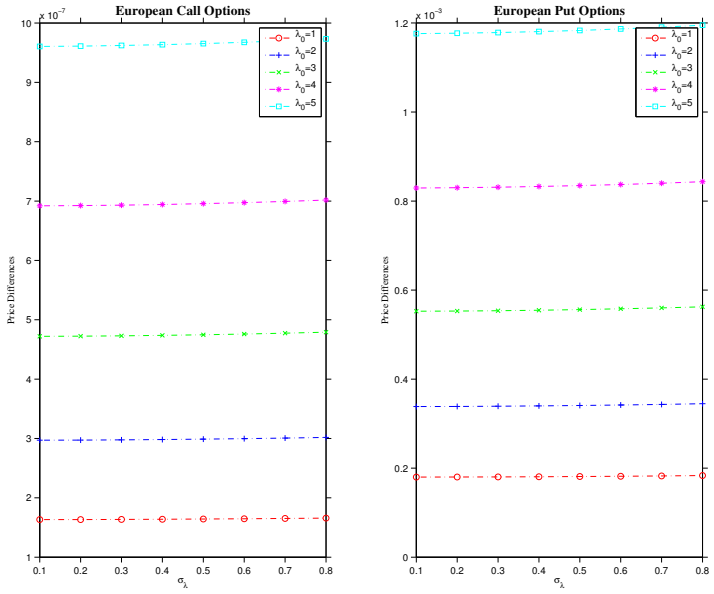

Figure 5.8: Option price sensitivity analysis for changes in volatility of jump intensity

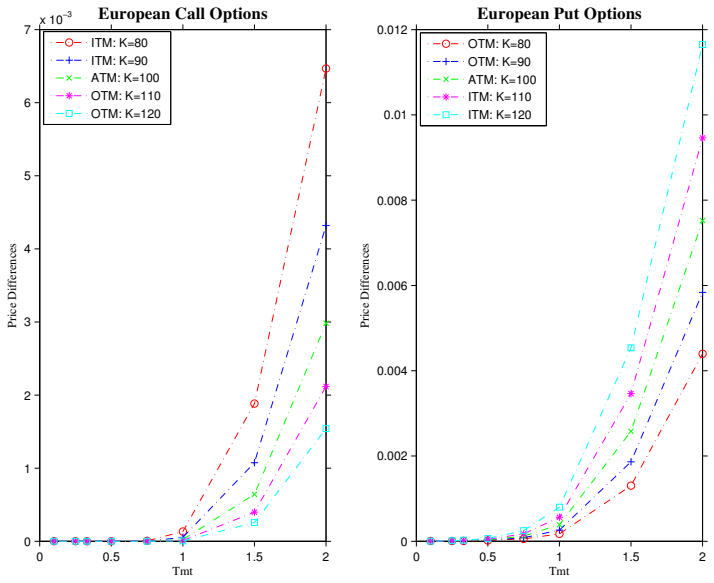

Figure 5.9: Option price sensitivity analysis for changes in time to maturity

5.4.4 Power Option Valuation Using SWIFT

Extending our analysis for the plain vanilla case, we price power call options with the SWIFT method and benchmark our results to the corresponding COS option prices. Adapting to the parameter values of the plain vanilla case, we demonstrate the exponential convergence of the resulting SWIFT prices to increasing scale m in Figure 5.10. In addition, we study the density mass loss from the approximation depending on the

rate of decay of the characteristic function, and presents our results in Figure 5.11. The above findings suggest that $m = 5$, with the corresponding $\bar{J} = 9$, may be the optimal tradeoff between accuracy and the speed of calculation, as the computational time required almost doubles, without significant recovery of the density mass lost, when increasing the scale of approximation to $m = 6$.

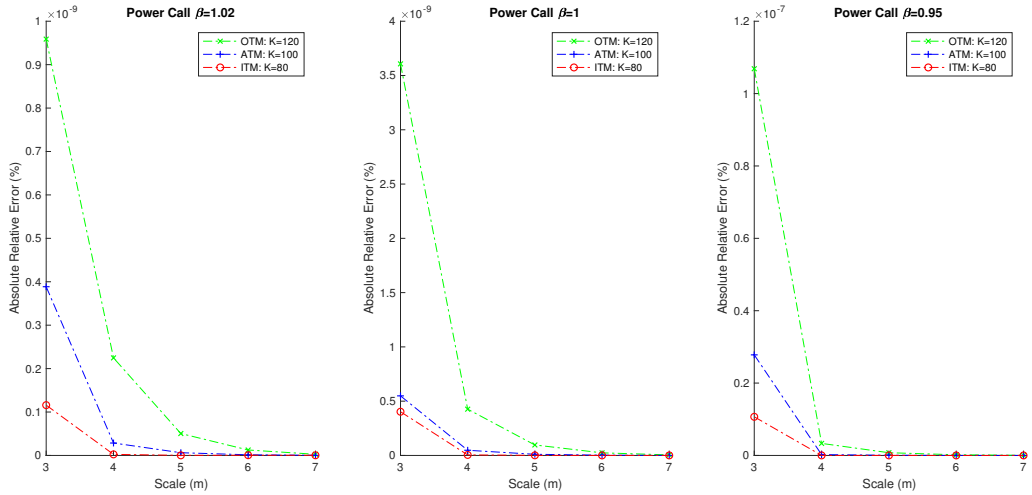

Figure 5.10: SWIFT option price error convergence on increasing scale of approximation (m)

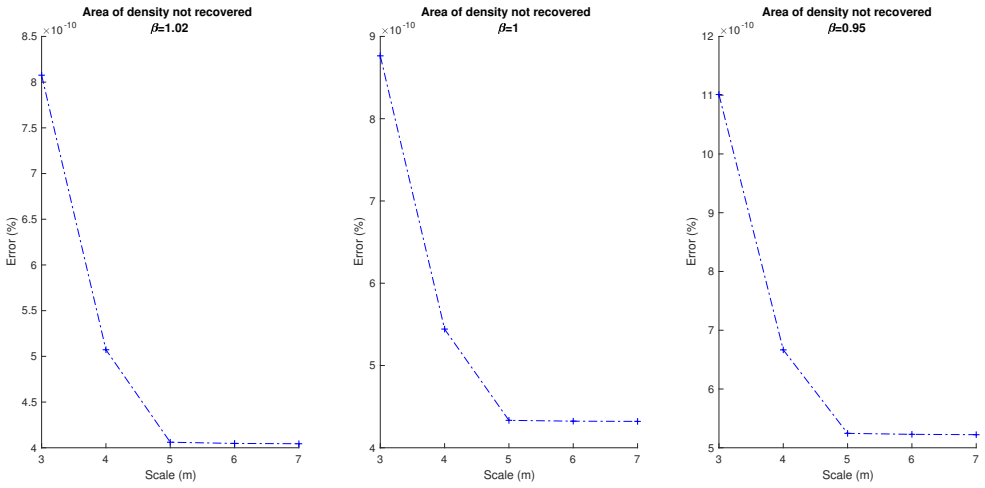

Figure 5.11: Convergence of density mass lost from approximation under the SWIFT method

Figure 5.12 presents the resulting prices of a basket of options across a range of strikes. Moreover, we reveal the price differences between the SWIFT and the benchmark COS methods. We observe that the SWIFT method is accurate and robust with minimal

pricing errors across different moneyness of the option, as well as the various powers, when benchmarked to the resulting COS prices. In comparison to prices from the Monte Carlo method, with 100,000 sample paths, we obtain a MAE of 9.3994e-03 across the various strikes. In terms of efficiency, the COS method consumes 3.1 milliseconds on average, while the SWIFT method takes a mere 0.63 milliseconds to obtain the option prices.[2]

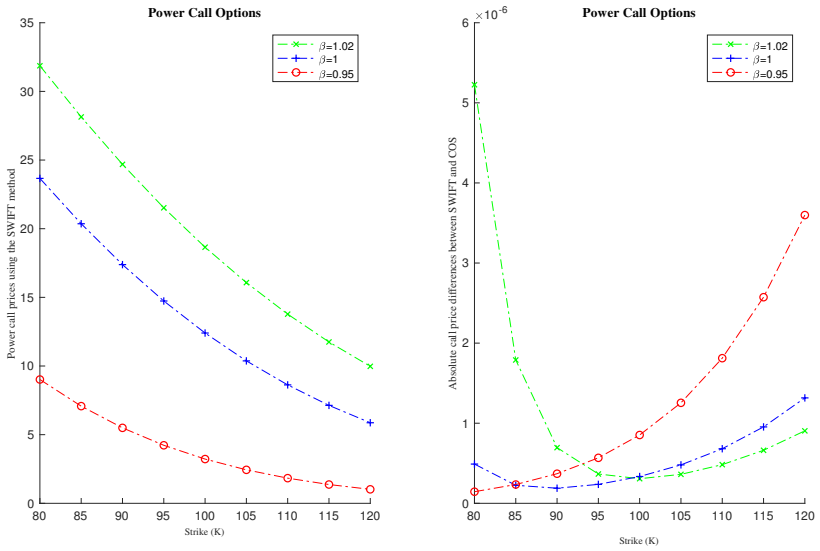

Figure 5.12: Power call option prices under the SWIFT method, and the resulting price differences between SWIFT and COS methods

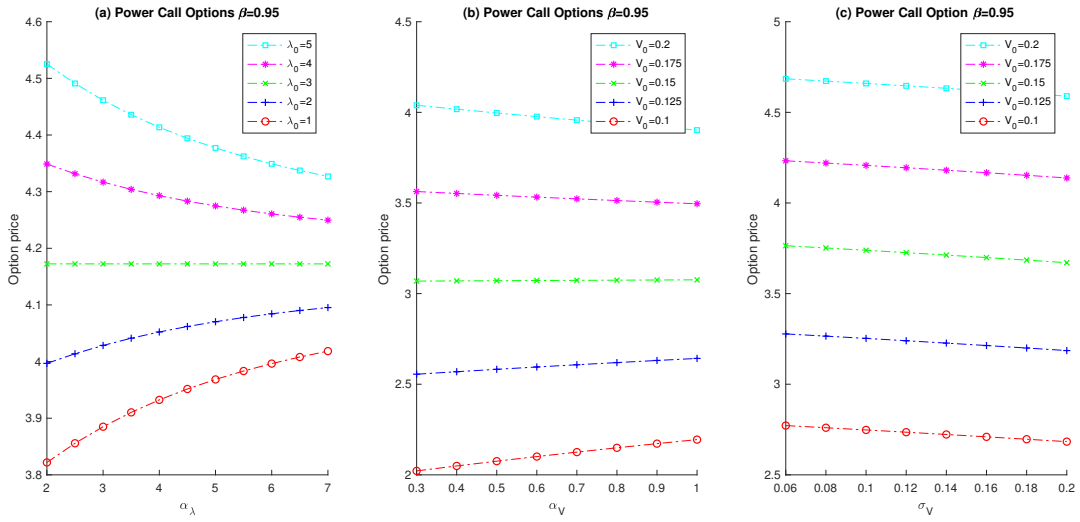

Figure 5.13: Power call option price sensitivity to changes in underlying model parameters with β=0.95

[2]Note: computing times were averaged over 10^4 iterations.

Finally, we evaluate the SWIFT power option prices to changes in the underlying model parameters. Figure 5.13.(a) demonstrates the changes in option prices to increases in the mean-reversion rate of jump intensity. By setting the long term intensity to 3, option prices with initial intensity less than the long term intensity increases with increasing rate of reversion, and the opposite is true for higher initial intensity. This is reasonable as increases (decreases) in the number of jumps increases (decreases) the variability in the underlying asset price, and increases (decreases) the resulting call value. Similar evidence is found in Figure 5.13.(b) when evaluating a long term volatility of 0.15. Faster rates of reversion increases the value of calls with a lower initial volatility, and the opposite is true for higher initial volatility values. Lastly, Figure 5.13.(c) demonstrate the changes in power call option prices to changes in the volatility of volatility.

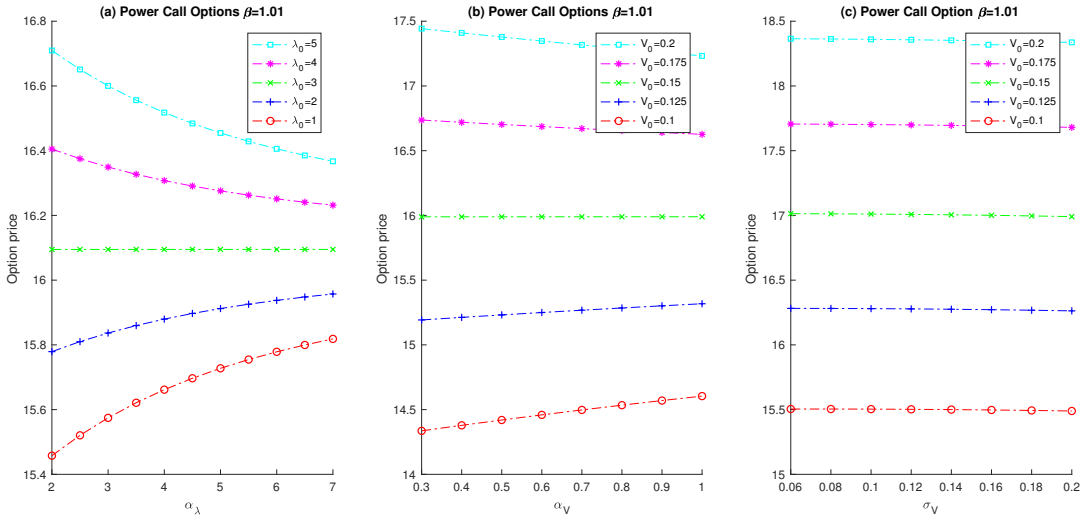

Figure 5.14: Power call option price sensitivity to changes in underlying model parameters with β=1.01

In addition to the above, we analyse the consistency in the model sensitivities for changes in power β, and present our findings in Figure 5.14. Our results demonstrate similar sensitivity patterns to those observed in the case where $\beta = 0.95$, providing further evidence in support of model robustness with the SWIFT method across a wide range of market conditions under our proposed framework.

5.5 Summary

In this chapter, we investigated the accuracy and efficiency of the COS and SWIFT methods in pricing both European calls and puts. Moreover, we extend our pricing method to the valuation of power call options. Under our proposed framework, the underlying asset price process is governed by a double exponential jump model with stochastic volatility and stochastic jump intensity. We show that, when benchmarked to the well-received FFT and COS methods, the SWIFT method is both accurate and significantly more efficient than its aforementioned predecessors. In addition, similar to the COS, our numerical experiments shows that the SWIFT method is robust and exhibits exponential error convergence. Furthermore, results from our Monte Carlo simulation suggest that the pricing methods above are indeed accurate in pricing options under our proposed framework.

The follow-up sensitivity analysis provides evidence of negligible pricing differences between the SWIFT and COS methods, across different levels of model parameters and their initial values. Such findings further demonstrates the stability and consistency of the SWIFT method under our proposed double exponential jump framework. Finally, we show that power option prices, as with the vanilla case, responds well to changes in model parameters, further supporting the use of our proposed stochastic model across a wide range of market conditions.

Chapter 6

Conclusions

In this thesis we evaluate highly efficient pricing methods for vanilla and exotic-type options. In particular, our option valuations takes place under a jump-diffusion framework with mean-reversion, stochastic volatility, and/or stochastic jump intensity. Since the density functions governing the proposed asset price dynamics are not readily available, but their associated characteristic functions are, we perform our pricing through numerical integration techniques in the Fourier domain instead. Two novel method of pricing are utilised, namely the Fourier-cosine expansion (COS) method and the Shannon inverse Fourier technique (SWIFT). While the former relies on the approximation of density functions with Fourier-cosine expansions, the latter is a wavelet-based pricing method, which has become prominent in the more recent past. We show through extensive numerical experiments and error analyses that not only are the two methods highly efficient when pricing exotic options, but are also accurate and robust with exponential error convergence. While the widely-acclaimed FFT method may require a large number of terms in the quadrature, which is also restricted to powers of 2, to obtain satisfactory accuracy, the alternative COS and SWIFT methods demonstrated high level of accuracy with significantly less coefficients in its summation. In addition, unlike the FFT, both the COS and SWIFT method does not rely on an arbitrary selection of damping factor for convergence. Moreover, the COS and SWIFT methods allows for the pricing of a basket of options, with a range of user-defined strikes, through a single computation.

We begin our analyses with the efficient pricing of discretely monitored arithmetic Asian options in Chapter 3. In particular, we allow for mean-reversion and jumps in the underlying price dynamics. We show that not only is the COS method more efficient than the widely-acclaimed fast Fourier transform (FFT) method, but it is also accurate with exponential rate of error convergence. In addition, the same level of accuracy

may be achieved with significantly less terms in the COS expansion in comparison to the FFT. Finally, we demonstrated the model's robustness in pricing arithmetic Asian options through an extensive sensitivity analysis. The results presented further support for the use of jumps in the underlying price dynamics, and the inclusion of time-varying mean level and asset volatility.

In Chapter 4, we proposed the use of the COS method in pricing options in a double exponential jump framework with stochastic volatility and interest rate. Our particular interest is the COS method's efficiency, accuracy and robustness when benchmarked to the an existing closed-form solution, instead of the commonly utilised Monte Carlo simulation, under such a framework. While a closed-form solution for the option value under the proposed framework exists, the embedded complexity in the solution demands considerable computational time to obtain the resulting price. Indeed, prior research has confirmed the efficiency of the FFT over the closed-form solution in such a regard. However, we advocated the use of the COS method instead, and show that the model is not only more efficient than the alternative FFT, but it is also more accurate when benchmarked to the closed-form solution. While the FFT requires thousands of grid points and terms to obtain an acceptable level of accuracy, the COS method does not suffer the same setback. Instead, the COS method shows exponential convergence with significantly less number of terms in its expansion. We conclude with a set of sensitivity analyses to provide further evidence of robustness for the COS method.

A natural continuation for a more practical framework, as motivated at the end of Chapter 4, is the inclusion of stochastic jump intensity in the asset price process. Hence, in Chapter 5 we proposed the efficient pricing of options under a double exponential jump framework with stochastic volatility and stochastic jump intensity. In particular, we explore the highly efficient valuation of options through wavelet-based pricing via the novel SWIFT method. Over and above using cumulants to define a truncation range from the real line a priori, as with the COS method, the SWIFT method has the advantage of calculating the area under the approximated density curve on the go. This allows us to adjust the approximation if the calculated area does not satisfy some pre-specified level of tolerance. In addition, the number of coefficient terms required in the approximation is determined automatically by the length of the interval.

Once the associated characteristic function is developed, we demonstrate the efficiency and accuracy of the SWIFT method in pricing options under the aforementioned double exponential jump framework. More precisely, we begin with the vanilla case, and expand our valuation to power options under the same framework. The extensive numerical experiments and detailed error analyses that followed demonstrated the SWIFT method's ability to outperform its aforementioned predecessors. Not only is the SWIFT

more efficient than both the FFT and COS methods under our proposed framework, but it is also accurate with exponential error convergence. In addition, the model is robust in pricing both the vanilla calls and puts, as well as exotics such as power options. Due to the absence of a closed-form solution, we compare the resulting prices to that of a Monte Carlo simulation to check for accuracy. Our numerical results indicate that the novel methods are indeed accurate for option pricing under our proposed double exponential jump framework. Moreover, evidence of pricing consistency and robustness under a wide range of market conditions are provided in our sensitivity analyses prior to the conclusion of the chapter.

Further work may include extending the pricing to early-exercise and other path-dependent options, and calculating the associated option greeks, under the double exponential jump framework with the SWIFT method. In addition, one can explore the application of the SWIFT method to higher dimensions. Another interesting alternative is to explore different approximation techniques for the sinc integral, and achieve more appropriate balances between the pricing accuracy and computational time required. Finally, outside the scope of option pricing, one may also explore the SWIFT computation of risk measures within the context of risk management.

List of Tables

List of Figures

Bibliography

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies*, 9(1), 69–107.

Beliaeva, N. & Nawalkha, S. (2012). Pricing American interest rate options under the jump-extended constant-elasticity-of-variance short rate models. *Journal of Banking & Finance*, 36(1), 151–163.

Bessembinder, H., Coughenour, J. F., Seguin, P. J., & Smoller, M. M. (1995). Mean reversion in equilibrium asset prices: evidence from the futures term structure. *Journal of Finance*, (pp. 361–375).

Boyle, P. & Boyle, F. (2001). *Derivatives: The tools that changed Finance*. Risk Books London.

Broadie, M. & Yamamoto, Y. (2003). Application of the fast Gauss transform to option pricing. *Management Science*, 49(8), 1071–1088.

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

Casassus, J. & Collin-Dufresne, P. (2005). Stochastic convenience yield implied from commodity futures and interest rates. *The Journal of Finance*, 60(5), 2283–2331.

Cattani, C. (2008). Shannon wavelets theory. *Mathematical Problems in Engineering*, 2008.

Chang, C., Fuh, C.-D., & Lin, S.-K. (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 & Finance*, 37(8), 3204–3217.

Chaudhuri, K. & Wu, Y. (2003). Mean reversion in stock prices: evidence from emerging markets. *Managerial Finance*, 29(10), 22–37.

Chourdakis, K. (2005). Option pricing using the fractional FFT. *Journal of Computational Finance*, 8(2), 1–18.

Chung, S. F. & Wong, H. Y. (2014). Analytical pricing of discrete arithmetic Asian options with mean reversion and jumps. *Journal of Banking & Finance*, 44, 130–140.

Cont, R. & Tankov, P. (1975). *Financial Modelling with jump processes, 2004*. Chapman & Hall/CRC: London.

Dempster, M. A. H. & Hong, S. G. (2002). Spread option valuation and the fast Fourier transform. In *Mathematical Finance—Bachelier Congress 2000* (pp. 203–220).: Springer.

Deng, G. (2007). Pricing European option in a double exponential jump-diffusion model with two market structure risks and its comparisons. *Applied Mathematics-A Journal of Chinese Universities*, 22(2), 127–137.

Deng, S. (2000). *Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes*. University of California Energy Institute Berkeley.

Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6), 1343–1376.

Espinosa, F. & Vives, J. (2006). A volatility-varying and jump-diffusion Merton type model of interest rate risk. *Insurance: Mathematics and Economics*, 38(1), 157–166.

Eydeland, A. & Wolyniec, K. (2003). *Energy and power risk management: New developments in modeling, pricing, and hedging*, volume 206. Hoboken, New Jersey: John Wiley & Sons Inc.

Fang, F. & Oosterlee, C. (2010). Pricing options under stochastic volatility with Fourier-cosine series expansions. *Progress in Industrial Mathematics at ECMI 2008*, (pp. 833–838).

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

Fang, F. & Oosterlee, C. W. (2009). Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik*, 114(1), 27.

Fusai, G., Marena, M., & Roncoroni, A. (2008). Analytical pricing of discretely monitored Asian-style options: Theory and application to commodity markets. *Journal of Banking & Finance*, 32(10), 2033–2045.

Gearhart, W. B. & Schultz, H. (1990). The function sin(x)/x. *The College Mathematics Journal*, 2(2), 90–99.

Geman, H. & Roncoroni, A. (2006). Understanding the fine structure of electricity prices. *The Journal of Business*, 79(3), 1225–1261.

Grzelak, L. A., Oosterlee, C. W., & Van Weeren, S. (2012). Extension of stochastic volatility equity models with the Hull–White interest rate process. *Quantitative Finance*, 12(1), 89–105.

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.

Hilliard, J. E. & Reis, J. A. (1999). Jump processes in commodity futures prices and options pricing. *American Journal of Agricultural Economics*, 81(2), 273–286.

Hoepfner, R. (2009). A time inhomogeneous Cox-Ingersoll-Ross diffusion with jumps. *arXiv:0906.1856*.

Huang, C.-S., O'Hara, J. G., & Mataramvura, S. (2017). Efficient pricing of discrete arithmetic Asian options under mean reversion and jumps based on Fourier-cosine expansions. *Journal of Computational and Applied Mathematics*, 311, 230–238.

Huang, J., Zhu, W., & Ruan, X. (2013). Fast Fourier transform based power option pricing with stochastic interest rate, volatility, and jump intensity. *Journal of Applied Mathematics*, 2013.

Huang, J., Zhu, W., & Ruan, X. (2014). Option pricing using the fast Fourier transform under the double exponential jump model with stochastic volatility and stochastic intensity. *Journal of Computational and Applied Mathematics*, 263, 152–159.

Hull, J. & White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2), 281–300.

Ibrahim, S. N., O'Hara, J. G., & Constantinou, N. (2013). Pricing power options under the Heston dynamics using the FFT. *New Trends in Mathematical Sciences*, 1(1), 1–9.

Ibrahim, S. N. I., O'Hara, J. G., & Constantinou, N. (2014). Pricing extendible options using the fast Fourier transform. *Mathematical Problems in Engineering*, 2014.

Jiang, G. J. (2002). Testing option pricing models with stochastic volatility, random jumps and stochastic interest rates. *International Review of Finance*, 3(3-4), 233–272.

Jorion, P. (1988). On jump processes in the foreign exchange and stock markets. *Review of Financial Studies*, 1(4), 427–445.

Jorion, P. & Sweeney, R. J. (1996). Mean reversion in real exchange rates: evidence and implications for forecasting. *Journal of International Money and Finance*, 15(4), 535–550.

Kim, Y.-J. & Kunitomo, N. (1999). Pricing options under stochastic interest rates: a new approach. *Asia-Pacific Financial Markets*, 6(1), 49–70.

Kirkby, J. L. (2015). Efficient option pricing by frame duality with the fast Fourier transform. *SIAM Journal on Financial Mathematics*, 6(1), 713–747.

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

Kou, S. G. & Wang, H. (2004). Option pricing under a double exponential jump diffusion model. *Management Science*, 50(9), 1178–1192.

Lord, R., Fang, F., Bervoets, F., & Oosterlee, C. W. (2008). A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes. *SIAM Journal on Scientific Computing*, 30(4), 1678–1705.

Maree, S. C., Ortiz-Gracia, L., & Oosterlee, C. W. (2017). Pricing early-exercise and discrete barrier options by Shannon wavelet expansions. *Numerische Mathematik*, 136(4), 1035–1070.

Marena, M., Fusai, G., & Longo, G. (2014). *Handbook of multi-commodity markets and products: structuring, trading, and risk management, chapter Asian options in commodity markets: structuring, pricing, and hedging.* John Wiley & Sons.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2), 125–144.

Mori, M. & Sugihara, M. (2001). The double-exponential transformation in numerical analysis. *Journal of Computational and Applied Mathematics*, 127(1), 287–296.

Ortiz-Gracia, L. & Oosterlee, C. W. (2013). Robust pricing of European options with wavelets and the characteristic function. *SIAM Journal on Scientific Computing*, 35(5), B1055–B1084.

Ortiz-Gracia, L. & Oosterlee, C. W. (2016). A highly efficient Shannon wavelet inverse Fourier technique for pricing European options. *SIAM Journal on Scientific Computing*, 38(1), B118–B143.

Pillay, E. & O'Hara, J. G. (2011). FFT based option pricing under a mean reverting process with stochastic volatility and jumps. *Journal of Computational and Applied Mathematics*, 235(12), 3378–3384.

Quine, B. & Abrarov, S. (2013). Application of the spectrally integrated Voigt function to line-by-line radiative transfer modelling. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 127, 37–48.

Santa-Clara, P. & Yan, S. (2010). Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *The Review of Economics and Statistics*, 92(2), 435–451.

Schmitz, A., Wang, Z., & Kimn, J.-H. (2014). A jump diffusion model for agricultural commodities with Bayesian analysis. *Journal of Futures Markets*, 34(3), 235–260.

Schöbel, R. & Zhu, J. (1999). Stochastic volatility with an Ornstein–Uhlenbeck process: an extension. *European Finance Review*, 3(1), 23–46.

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3), 923–973.

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.

Seifert, J. & Uhrig-Homburg, M. (2007). Modelling jumps in electricity prices: theory and empirical evidence. *Review of Derivatives Research*, 10(1), 59–85.

Stein, E. M. & Stein, J. C. (1991). Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies*, 4(4), 727–752.

Yamamoto, Y. (2005). Double-exponential fast Gauss transform algorithms for pricing discrete lookback options. *Publications of the Research Institute for Mathematical Sciences*, 41(4), 989–1006.

Zhang, B. & Oosterlee, C. (2014). Pricing of early-exercise Asian options under Lévy processes based on Fourier cosine expansions. *Applied Numerical Mathematics*, 78, 14–30.

Zhang, B. & Oosterlee, C. W. (2013). Efficient pricing of European-style Asian options under exponential Lévy processes based on Fourier cosine expansions. *SIAM Journal on Financial Mathematics*, 4(1), 399–426.

Zhang, S. & Geng, J. (2016). Fourier-cosine method for pricing forward starting options with stochastic volatility and jumps. *Communications in Statistics-Theory and Methods*, 46(20), 9995–10004.

Zhang, S. & Wang, L. (2013). Fast Fourier transform option pricing with stochastic interest rate, stochastic volatility and double jumps. *Applied Mathematics and Computation*, 219(23), 10928–10933.

Appendix A. Discretisation of asset price dynamics (5.1)

The asset price dynamics (5.1) can be discretised as follows:

$$\begin{aligned}
\ln S_{t+\Delta t} &= \ln S_t + (r - d - \lambda_t \delta)\Delta t + \sqrt{v_t}\epsilon_1\sqrt{\Delta t} + (e^Y - 1)\ln S_t(N_{t+\Delta t} - N_t), \\
v_{t+\Delta t} &= v_t + (\theta_v - \alpha_v v_t)\Delta t + \sigma_v\sqrt{v_t}(\rho\epsilon_1 + \sqrt{1-\rho^2}\epsilon_2)\sqrt{\Delta t}, \\
\lambda_{t+\Delta t} &= \lambda_t + (\theta_\lambda - \alpha_\lambda\lambda_t)\Delta t + \sigma_\lambda\sqrt{\lambda_t}\epsilon_3\sqrt{\Delta t},
\end{aligned}$$

where t is the current time step, and Δt represents the length of time per simulated step. ϵ_1, ϵ_2, and ϵ_3 are independent random variables sampled from a standard normal distribution. Option values can be obtained once the various sample paths for the final asset price at maturity T have been simulated with the discretisation above. The valuation then follows with the usual discounted expectation of the final option payoff.