OPTION PRICING USING FOURIER SPACE TIME-STEPPING FRAMEWORK

by

Vladimir Surkov

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

Copyright © 2009 by Vladimir Surkov

Abstract

Option Pricing using Fourier Space Time-stepping Framework

Vladimir Surkov

Doctor of Philosophy

Graduate Department of Computer Science

University of Toronto

2009

This thesis develops a generic framework based on the Fourier transform for pricing and hedging of various options in equity, commodity, currency, and insurance markets. The pricing problem can be reduced to solving a partial integro-differential equation (PIDE). The Fourier Space Time-stepping (FST) framework developed in this thesis circumvents the problems associated with the existing finite difference methods by utilizing the Fourier transform to solve the PIDE. The FST framework-based methods are generic, highly efficient and rapidly convergent.

The Fourier transform can be applied to the pricing PIDE to obtain a linear system of ordinary differential equations that can be solved explicitly. Solving the PIDE in Fourier space allows for the integral term to be handled efficiently and avoids the asymmetrical treatment of diffusion and integral terms, common in the finite difference schemes found in the literature. For path-independent options, prices can be obtained for a range of stock prices in one iteration of the algorithm. For exotic, path-dependent options, a time-stepping methodology is developed to handle barriers, free boundaries, and exercise policies.

The thesis includes applications of the FST framework-based methods to a wide range of option pricing problems. Pricing of single- and multi-asset, European and path-dependent options under independent-increment exponential Lévy stock price models, common in equity and insurance markets, can be done efficiently via the cornerstone FST method. Mean-reverting Lévy spot price models, common in commodity markets, are handled by introducing a frequency transformation, which can be readily computed via scaling of the option value function. Generating stochastic volatility, to match the long-term equity options market data, and stochastic skew, observed in currency markets, is addressed by introducing a non-stationary extension of multi-dimensional Lévy processes using regime-switching. Finally, codependent jumps in

multi-asset models are introduced through copulas.

The FST methods are computationally efficient, running in $O(MN^d \log_2 N)$ time with M time steps and N space points in each dimension on a d -dimensional grid. The methods achieve second-order convergence in space; for American options, a penalty method is used to attain second-order convergence in time. Furthermore, graphics processing units are utilized to further reduce the computational time of FST methods.

Dedication

To my parents, for having the strength and conviction
to see me embark on this journey.

To my wife, for being supportive and loving
every step of the way.

Acknowledgements

To my supervisors Dr. Kenneth R. Jackson and Dr. Sebastian Jaimungal, I owe an immense debt of gratitude. Without their insight, guidance, and encouragement, this work would not be possible. From the initial ideas to the final draft, their doors were always open and I had the opportunity to choose my research direction while benefiting greatly from their sage advice.

My most sincere thanks to the University of Waterloo faculty Dr. Peter Forsyth, the University of Toronto faculty Dr. Christina Christara, and Boston University faculty Dr. Marcel Rindisbacher for their input into this work.

I am also thankful for the comments and suggestions I received from various conference participants and journal referees. Your input has been highly valuable.

Contents

1	Introduction	1
1.1	Market Models	3
1.1.1	Equity Markets	3
1.1.2	Currency Markets	7
1.1.3	Commodity Markets	8
1.1.4	Insurance Markets	10
1.2	Option Pricing Problem	11
1.3	Numerical Methods	12
1.3.1	Monte Carlo Simulation	13
1.3.2	Tree Methods	13
1.3.3	Finite Difference Methods	14
1.3.4	Fast Fourier Transform Methods	18
1.3.5	Other Methods	19
1.4	Motivation for Research	19
2	Fourier Space Time-stepping Method	22
2.1	Introduction	22
2.2	Spot Price Model	23
2.3	PIDE Solution	24
2.4	The Method	27
2.5	Numerical Properties	29
2.5.1	Stability	29
2.5.2	Convergence	30
2.5.3	Precision	32
2.6	Applications to Option Pricing	32
2.6.1	European Options	33
2.6.2	American Options	36

2.6.3	American Options with Penalty Method	40
2.6.4	Barrier Options	43
2.7	Applications to Hedging with Greeks	45
3	Mean-Reverting Fourier Space Time-stepping Method	51
3.1	Introduction	51
3.2	Spot Price Model	52
3.3	PIDE Solution	55
3.4	The Method	58
3.5	Applications to Option Pricing	60
4	Spot Price Model Extensions	69
4.1	Introduction	69
4.2	Regime-Switching	70
4.3	Codependent Jumps via Copulas	75
5	Exotic Options	80
5.1	Introduction	80
5.2	Shout Options	81
5.3	Swing Options	84
6	Graphics Processing Units	87
6.1	Introduction	87
6.2	Fast Fourier Transform Computation on Graphics Processing Units	88
6.3	Applications to Option Pricing	90
6.4	Applications to Parallel Option Pricing	94
7	Conclusions	97
7.1	Summary of Research	97
7.2	Future Work	98
A	Acronyms and Notation	100
A.1	Acronyms	100
A.2	Notation	101
B	Option and Model Parameters	103

C	Supplementary Results	105
C.1	Fourier Space Time-stepping Results	106
C.2	Fourier Space Time-stepping Time Convergence Results	109
C.3	Fourier Space Time-stepping Greeks Results	110
C.4	Mean-Reverting Fourier Space Time-stepping Results	111
C.5	Graphics Processing Units Pricing Results	112
	Bibliography	112

List of Tables

2.1	Lévy densities and characteristic exponents of various exponential Lévy models .	26
2.2	Pricing results for a European option under a Merton jump-diffusion model . . .	33
2.3	Pricing results for a European spread option under a 2D Merton jump-diffusion model	35
2.4	Pricing results for an American option under a Carr-Geman-Madan-Yor model .	38
2.5	Pricing results for an American spread option under a 2D Merton jump-diffusion model	38
2.6	Pricing results for an American option under a Carr-Geman-Madan-Yor model with penalty method	43
2.7	Pricing results for a barrier option under a Black-Scholes-Merton model	44
3.1	Pricing results for a European option under a mean-reverting Merton jump-diffusion model	62
3.2	Pricing results for an American option under a mean-reverting Merton jump-diffusion model	62
3.3	Pricing results for a discrete barrier option under a mean-reverting Merton jump-diffusion model	62
3.4	Pricing results for a European option under a mean-reverting Kou jump-diffusion with decoupled jumps model	64
3.5	Pricing results for a Bermudan option under a mean-reverting Kou jump-diffusion with decoupled jumps model	64
3.6	Pricing results for a European spread option under a 2D mean-reverting Kou jump-diffusion with Gaussian copula jumps model	64
3.7	Pricing results for a Bermudan spread option under a 2D mean-reverting Kou jump-diffusion with Gaussian copula jumps model	66
3.8	Pricing results for a European option under a geometric Brownian motion with mean-reverting reversion level model	66

3.9	Pricing results for a European option under a geometric Brownian motion with mean-reverting reversion level model	66
4.1	Cumulative and probability density functions for various copulas	77
6.1	Fast Fourier Transform execution performance on Central Processing Unit vs. Graphics Processing Unit	90
6.2	Pricing results for a European option under a Kou jump-diffusion model	92
6.3	Pricing results for a European catastrophe equity put option under a joint stock-loss model	92
6.4	Pricing results for an American option under a Variance Gamma model	94
6.5	Pricing results for an American double-trigger stop-loss option under a joint stock-loss model	94
6.6	Timing results for parallel pricing of European options	95
C.1	Pricing results for a European option under a Merton jump-diffusion model . . .	106
C.2	Pricing results for a European option under a Kou jump-diffusion model	106
C.3	Pricing results for a European option under a Variance Gamma model	106
C.4	Pricing results for a European option under a Carr-Geman-Madan-Yor model . .	107
C.5	Pricing results for an American option under a Merton jump-diffusion model . .	107
C.6	Pricing results for an American option under a Merton jump-diffusion model using a penalty method	107
C.7	Pricing results for an American option under a Variance Gamma model	108
C.8	Pricing results for an American option under a Variance Gamma model using a penalty method	108
C.9	Pricing results for a barrier option under a Merton jump-diffusion model	108
C.10	Pricing results for an American option under a Merton jump-diffusion model . .	109
C.11	Pricing results for an American option under a CGMY model	109
C.12	Pricing results for a European option under a mean-reverting Kou jump-diffusion model	111
C.13	Pricing results for an American option under a mean-reverting Kou jump-diffusion model	111
C.14	Pricing results for a discrete barrier option under a mean-reverting Kou jump-diffusion model	111
C.15	Pricing results for a European option under a Carr-Geman-Madan-Yor model . .	112
C.16	Pricing results for an American option under a Merton jump-diffusion model . .	112

C.17 Pricing results for a European spread option under a 2D Black-Scholes-Merton model	112
C.18 Pricing results for an American spread option under a 2-dimensional Black- Scholes-Merton model	113

List of Figures

1.1	IBM stock price vs. geometric Brownian motion sample price path	4
1.2	IBM stock price vs. Merton jump-diffusion sample price path	5
1.3	Russell 3000 Index vs. CBOE Volatility Index	6
1.4	US dollar vs. pound sterling and Japanese yen	7
1.5	West Texas Intermediate crude oil prices and NYMEX Henry Hub natural gas prices	8
1.6	Great Britain average system sell electricity prices	9
1.7	USA flood damage per capita and insurer sample stock price path	11
1.8	Spot price evolution in the binomial option pricing model	14
2.1	Errors for pricing a European option under a Merton jump-diffusion model and a barrier option under a Black-Scholes-Merton model	34
2.2	Payoff and value of a catastrophe equity put option under a joint stock-loss model	36
2.3	Schematic representation of the Fourier Space Time-stepping method	37
2.4	Exercise boundary of an American spread option	39
2.5	Payoff and value of an American double-trigger stop-loss option under a joint stock-loss model	40
2.6	Exercise boundary of an American double-trigger stop-loss option under a joint stock-loss model	41
2.7	Error in computing option price and Greeks for a European option under a Black-Scholes-Merton model	49
3.1	Sample price paths under mean-reverting Merton and Kou jump-diffusion models and a mean-reverting Kou jump-diffusion with decoupled jumps model	61
3.2	Sample price paths under a 2D mean-reverting Kou jump-diffusion with Gaus- sian copula jumps model and a geometric Brownian motion with mean-reverting reversion level model	65
3.3	Errors for pricing European options under mean-reverting geometric Brownian motion (with mean-reverting level) models	67

4.1	Price and exercise boundary of an American option under a regime-switching Black-Scholes-Merton model	73
4.2	Implied volatility smiles in regime-switching vs. stationary volatility Merton jump-diffusion models	74
4.3	Implied volatility smiles in regime-switching vs. stationary volatility and skew Variance Gamma models	75
4.4	Random samples from four commonly used copulas	78
4.5	Approximation error for the characteristic exponent of Gaussian copula jumps model	79
5.1	Prices and exercise boundaries of a multi-shout option under a Variance Gamma model	83
5.2	Optimal shout times of a single-shout option under a Variance Gamma model . .	84
5.3	Effect of mean-reversion level and speed on the value of a swing option under a mean-reverting Merton jump-diffusion model	86
6.1	Performance of various high-end Central and Graphics Processing Units, and Cell Broadband Engine	89
6.2	Timing results for batched Fast Fourier Transform computation	95
C.1	Error in computing option price and Greeks for a digital option under a Merton jump-diffusion model	110

Chapter 1

Introduction

The use of a growing array of derivatives and the related application of more-sophisticated approaches to measuring and managing risk are key factors underpinning the greater resilience of our largest financial institutions Derivatives have permitted the unbundling of financial risks.

(Federal Reserve Board Chairman Alan Greenspan, May 5, 2005)

Derivatives are financial contracts which have a value contingent on the evolution of the underlying assets. Derivatives are traded in various markets, including, but not limited to, equity, commodity, currency, credit, and interest rate markets. The primary intent of derivatives is to reduce the risk that the value of the underlying assets will change unexpectedly. However, derivatives can also be used to acquire risk by speculating on the value of the underlying assets. According to the Bank of International Settlement, the total outstanding notional amount of all derivatives contracts has rapidly grown from \$88 trillion in December 1998 to \$624 trillion in December 2007. Unfortunately, this unprecedented and primarily speculative growth in derivatives use, especially in credit and mortgage derivatives markets, is associated with the tremendous turmoil in global financial markets that dominated the period of writing of this thesis. Regulators, credit agencies, financial institutions and homeowner all share the blame for the crisis, whose effect will be felt for many years to come. And while there are growing calls to limit derivatives' use, the solution lies in better modeling, pricing and risk management techniques, in conjunction with more stringent regulation.

When a bridge collapses, no one demands the abolition of civil engineering If engineering is to blame, the solution is better—not less—engineering. Furthermore, it would be preposterous to replace the bridge with a slower, less efficient ferry rather than to rebuild the bridge and overcome the obstacle.

(Carnegie Mellon University professor Steven Shreve, October 8, 2008)

Better engineering is the driving principle of this work. While the field of research in derivatives is very broad, this thesis focuses on the computational aspect of derivatives pricing, by proposing a new and efficient algorithm for pricing various types of options (particular type of derivative contracts) in equity, commodity, currency and insurance markets. Moreover, the aim is to present a general framework for derivatives pricing that can be tailored and extended to a variety of applications. The method's versatility is presented through various applications and numerical examples.

The outline of this thesis is as follows. This chapter introduces the field of option pricing by presenting a survey of the financial markets, pricing theory and numerical methods available to perform the valuation. Also, the chapter motivates the research presented in this thesis and highlights the main contributions of this work to the area of numerical option pricing. Chapter 2 presents the *Fourier Space Time-stepping* (FST) method and applies it to pricing of various options in equity markets. The precision of the method and the order of convergence are established through numerical experiments. Also, the *Greeks Fourier Space Time-stepping* (greekFST) method is developed for computation of option value sensitivities to changes in market conditions and model parameters. Chapter 3 extends the FST method to handle the mean-reverting spot price processes, commonly used in commodity markets, via the *mean-reverting Fourier Space Time-stepping* (mrFST) method. Again, the order of convergence and precision under this extension is studied. Chapter 4 discusses two spot price model extensions — regime-switching and copula driven jumps. Introducing regime-switching into a stationary model allows it to generate stochastic skew/volatility behavior, commonly observed in currency markets. Introducing jumps driven by a copula is essential for joint models on several assets that respond in a codependent fashion to changes in market conditions or arrival of information, typically seen in commodity markets. The chapter develops the *regime-switching Fourier Space Time-stepping* (rsFST) method for pricing under regime-switching models and discusses an efficient technique for working with copula jumps model. Chapter 5 applies the FST algorithm to pricing of two exotic options — shout options, which provide enhanced protection by allowing the strike price of the option to be reset, and swing options, which provide constrained flexibility with respect to the amount and timing of commodity delivered. Chapter 6 shows how graphics processing units (GPUs) can be utilized to increase the computational efficiency of the FST method. Finally, Chapter 7 summarizes the main contributions of this work and gives possible avenues for further research. Commonly used acronyms are presented in Appendix A, option and model parameters are defined in Appendix B, and further numerical results obtained with the FST framework-based methods are provided in Appendix C.

The outline of the remainder of this chapter is as follows. Section 1.1 examines the stylistic features of various processes, such as equity and commodity spot prices, foreign exchange rates

and catastrophe events, on which financial derivatives are traded. Also, models used in the literature to describe the behavior of such processes are presented. Section 1.2 gives an overview of mechanics of option markets and introduces, through formal mathematical language, the problem of pricing standard options. Section 1.3 surveys various numerical methods currently available for tackling the option pricing problem. Lastly, Section 1.4 motivates this research by describing the major challenges in the area of numerical option pricing.

1.1 Market Models

A stochastic process is a variable whose value changes over time in a non-deterministic way. This section presents continuous-variable, continuous-time stochastic processes for prices in various markets. Although in practice these variables can only be observed when the exchanges are open and their values are restricted to discrete values (e.g., multiples of a cent), such processes provide accurate approximations for the real-world processes. Stochastic processes used in modeling of market factors in credit and interest rate markets are not discussed as this work does not cover pricing of derivatives in these markets.

1.1.1 Equity Markets

The weak form of market efficiency states that the current price of a stock impounds all publicly available information past and present, i.e., there are no patterns to stock prices. If it were not so, above average returns could be made by investors through technical analysis of the past history of stock prices; there is little evidence that anyone can do this consistently. Thus, it is usually assumed that stock prices are Markov processes, i.e., the distribution of future returns is independent of past history and only depends on the value of the stock price at this instant.

Bachelier (1900) was the first to address the problem of modeling stock prices and assumed driftless Brownian motion dynamics for the prices. Under such a model, however, negative realizations of stock prices are possible. Osborne (1959) modified the Bachelier model to assume that the returns, not stock prices, follow a Brownian motion:

$$\frac{dS(t)}{S(t)} = \gamma dt + \sigma dW(t), \quad (1.1)$$

where γ and σ are the drift and volatility of returns and $W(t)$ is a Brownian motion. Consequently, stock prices follow log-normal distribution, instead of the normal distribution proposed by Bachelier. This model is known as the geometric Brownian motion (GBM) model or the Black-Scholes-Merton (BSM) model (due to work of Black and Scholes (1973) and Merton (1973)). In Figure 1.1 a sample stock price path under the BSM model is plotted against the

Figure 1.1: IBM stock price vs. geometric Brownian motion sample price path.

actual price of IBM stock over the same time period. Both curves have similar small-scale behavior and without prior knowledge it is practically impossible to tell which curve is the price of IBM stock and which curve is the simulated GBM process.

Today, the BSM model is widely used to model asset prices in a wide array of markets, owing a large part of its popularity to mathematical tractability of pricing formulas that are based on this model. However, the log-normal distribution of stock returns, as implied by the BSM model, is not supported by the empirical evidence, which points to a distribution with a higher probability for outliers. Also, empirical evidence suggests that the stock price changes can be classified as either marginal changes (due to supply-demand imbalance or changes in market economic conditions and outlook) or large changes (due to arrival of information that is usually company or industry specific) in price. Merton (1976) introduced a discontinuous stock price model, where the stock price changes are comprised of marginal changes, modeled by geometric Brownian motion, and large variations in price, modeled by a jump process. In differential form the stock price process can be described via

$$\frac{dS(t)}{S(t-)} = \gamma dt + \sigma dW(t) + dJ(t), \quad (1.2)$$

where $S(t-)$ is the stock price at time t before the jump. The jump process is given by $J(t) \triangleq \sum_{n=1}^{N(t)} j_n$, where $N(t)$ is a Poisson process with activity rate λ which governs the arrival of i.i.d. jumps j_n arriving at Poisson times t_n . Merton (1976) uses a log-normal distribution for the jumps j_n , while Kou (2002) accounts for skewness of returns by using a double-exponential distribution. Looking at the stock price history of IBM during the year 1987 in Figure 1.2,

Figure 1.2: IBM stock price vs. Merton jump-diffusion sample price path.

jump models can be effectively used to generate large movements in the stock price, as seen on October 19, 1987¹.

A more general class of models based on Lévy processes² is becoming increasingly popular in modeling stock prices. Lévy processes allow models to be built that accurately reflect the fat tails and skewness of asset returns observed in the real world. Unlike the jump models, where the number of jumps on any finite interval is finite, Lévy processes may have an infinite number of jumps on a finite interval with most jumps being infinitesimally small. Moreover, their dynamics are rich enough to generate Brownian-like behavior on a small time scale with the benefit of analytical tractability. Huang and Wu (2004) report results of numerous statistical tests which demonstrate that models with infinitesimal jumps outperform jump-diffusion models for equity options.

The exponential Lévy model defines the stock price process to be

$$S(t) = S(0) e^{X(t)}, \quad (1.3)$$

where $X(t)$ is a Lévy process. While various classes of Lévy processes exist, this work focuses on the Variance Gamma (VG) model of Madan and Seneta (1990) and the Normal Inverse Gaussian model of Barndorff-Nielsen (1997), belonging to the class of generalized hyperbolic models introduced by Eberlein and Prause (2002), and the CGMY model of Carr, Geman, Madan, and Yor (2002), belonging to the class of tempered stable models. These three models

¹See Bates (1991) for analysis of performance of jump-diffusion models in predicting the crash of 1987

²See Sato (1999) for further mathematical background on Lévy processes

Figure 1.3: Russell 3000 Index vs. CBOE Volatility Index.

are currently the models of choice both in academia and industry for modeling equity prices³.

Another line of research addresses the biases inherent in the BSM model by elevating volatility to a continuous stochastic variable. A number of approaches have been suggested by Hull and White (1987), Scott (1987), Stein and Stein (1991) and Heston (1993). The latter approach, being the most suitable for equity markets, models the spot price by

$$\frac{dS(t)}{S(t)} = \gamma dt + \sqrt{v(t)} dW_s(t), \quad (1.4a)$$

$$dv(t) = \kappa (\theta - v(t)) dt + \sigma \sqrt{v(t)} dW_v(t), \quad (1.4b)$$

where $dW_s(t)dW_v(t) = \rho dt$. The model allows arbitrary correlations between volatility and asset returns. In equity markets this correlation is typically negative. Periods of low volatility are associated with steady, upward movement in stock prices, which incidentally causes investors to become complacent to risk. In contrast, significant macroeconomic events lead to repricing of market risk and redistribution of capital resulting in higher volatility and downward movement in stock prices. This phenomena is well exemplified by the stock market behavior over the past eight years. Figure 1.3 depicts the Russell 3000 Index, representing the U.S. broad market, as compared to the CBOE VIX Index, representing the implied market volatility, during the years 2001 to 2009. The steady rise of stock prices during the bull market of 2003-2007 was accomplished with extremely small volatility (below 20%), while the tumultuous market declines of 2001-2002 and 2008 are associated with extremely high volatility (exceeding 60% in late 2008).

³See Cont and Tankov (2004) and Papapantoleon (2005) for an exhaustive survey of exponential Lévy models and their applications to asset price modeling

Figure 1.4: US dollar vs. pound sterling and Japanese yen

Another technique to incorporate stochastic volatility into a stationary volatility model is via regime-switching, an approach first proposed by Naik (1993). The essential idea is to assume that the world switches between different states, representing, for instance, low, medium and high volatilities. Regime states can either be visible or hidden from market participants and the transition between the states is governed by a continuous-time Markov chain. One limitation of regime-switching models is their inability to incorporate correlation between asset returns and market volatility. BSM and jump-diffusion models with regime-switching volatility are used for stochastic volatility modeling in this thesis.

1.1.2 Currency Markets

The global currency, or foreign exchange, market is much larger than the equity market, yet despite its size, the currency market is not as widely studied, primarily due to the fact that the trading in the market is done directly between the counterparties (over-the-counter trading), as opposed to equity markets, which are primarily exchange traded. While a plot of various foreign exchange rates may suggest that the BSM model is an adequate model in this environment, a closer look suggests the need for more complicated models. Market prices of various currency derivatives imply that the volatility varies randomly over time and the volatility of volatility (kurtosis) is significant. Models that strive to capture these effects are the aforementioned Merton (1976) jump-diffusion and Heston (1993) stochastic volatility models; Bates (1996) model generalizes the previous two models by incorporating both jump-diffusion and stochastic

volatility.

Moreover, unlike in the equity markets, there is a substantial variation over time in the skewness of returns distribution, which suggest a need for a stochastic skewness component in the model (in the aforementioned three models the skew is non-stochastic). Carr and Wu (2007) propose a stochastic skew model that captures both the stochastic volatility and stochastic skew of currency dynamics and has the advantage of being highly tractable.

Another approach to incorporate a stochastic variable into a stationary model is via regime-switching, as discussed in the previous section. In this thesis, BSM and jump-diffusion models with regime-switching volatility and exponential Lévy models with regime-switching volatility and skewness are used for foreign exchange rate modeling.

1.1.3 Commodity Markets

Energy commodities, such oil, gas and electricity, lack the liquidity of equity and currency markets, have large costs associated with storage and exhibit high volatilities and sudden price jumps. Finally, and most importantly, commodities tend to revert to a long run equilibrium price. These stylized empirical facts are well documented in Clewlow and Strickland (2000), Carmona and Durrleman (2003), Eydeland and Wolyniec (2003), Hull (2005), and Geman (2005), for example.

During periods of strong economic growth, demand for raw materials and energy outstrips

Figure 1.5: West Texas Intermediate crude oil prices and NYMEX Henry Hub natural gas prices.

Figure 1.6: Great Britain average system sell electricity prices

supply, driving the prices upwards. Eventually, high commodity (and consequently final product) prices lead to economic slowdown and diminished demand from consumers, which coupled with increased supply, cause prices to mean revert to their long-run equilibrium. Such behavior can be observed in various commodity markets. Figure 1.5 shows West Texas Intermediate crude oil and NYMEX Henry Hub natural gas prices during the years 2004 to 2008. While both crude oil and natural gas prices spiked during the commodity run-up of 2007-2008, only natural gas spiked following Hurricane Katrina in 2005. Both oil and gas fell significantly by the end of 2008, following the global economic downturn. Electricity, being difficult and expensive to store, is prone to spikes due to asymmetry in supply and demand. Bringing additional power into the grid in response to a spike in demand (e.g., increased use of air conditioners on a hot day) or disruption in supply (e.g., generator malfunction) can be expensive. Still, these large spikes in electricity prices (orders of magnitude higher than typical jumps) quickly revert to the mean once the balance is restored. Figure 1.6 depicts the frequent spikes in Great Britain electricity prices, with the magnitude of jumps oftentimes exceeding 200 – 300% of the ‘normal’ price. Such behavior is common to many other electricity markets around the world.

Some of the pioneering work in modeling commodity spot prices has been carried out by Gibson and Schwartz (1990), Cortazar and Schwartz (1994), Pilipović (1997) and Schwartz (1997), where the spot prices are modeled using geometric Brownian motion with stochastic convenience yield and/or interest rate. The Clewlow and Strickland (2000) model is widely

used since it incorporates mean-reversion into the spot price dynamics and allows for jumps:

$$d \ln S(t) = \kappa (\theta - \ln S(t)) dt + \sigma dW(t) + dJ(t). \quad (1.5)$$

For electricity prices, the Hikspoors and Jaimungal (2007) model incorporates different mean-reversion speeds for diffusion and jump components, since spikes are typically pulled back more quickly than the diffusion components:

$$S(t) = e^{X_1(t) + X_2(t)}, \quad (1.6a)$$

$$dX_1(t) = \kappa_1 (\theta - X_1(t)) dt + \sigma dW(t), \quad (1.6b)$$

$$dX_2(t) = -\kappa_2 X_2(t) dt + dJ(t). \quad (1.6c)$$

Section 3.2 develops a general, multi-factor model with jumps that generalizes the models discussed above and allows for rich price dynamics.

1.1.4 Insurance Markets

The effect of catastrophe events, such as a flood, an earthquake or a fire, on the share price of insurance companies is a contentious issue. Different studies arrived at conflicting conclusions regarding the correlation between the market value of an insurance company and catastrophe losses incurred on its policies. A possible explanation could be that, on the one hand, portfolio losses constrain future growth of the insurer, while, on the other hand, such losses could be more than offset by premium increases⁴. To add to the complexity, it seems that the market differentiates between the different types of natural disasters when factoring them into the insurer's stock price. Nonetheless, it is essential that a good model for the evolution of share prices of insurance companies should take losses from catastrophic events into account.

Cox, Fairchild, and Pedersen (2004) introduce a simple model for the joint evolution of catastrophe losses and stock price where the size of the loss is fixed. Jaimungal and Wang (2006) extend the model by introducing random losses:

$$S(t) = S(0) e^{\gamma t - \chi L(t) + \sigma W(t)}, \quad (1.7a)$$

$$L(t) = \sum_{n=1}^{N(t)} l_n, \quad (1.7b)$$

where χ represents the percentage drop in the share price per unit of loss, $N(t)$ is a Poisson process with activity rate λ and l_i are i.i.d. random variables (with probability density $f_L(l)$ having support on $\mathbb{R}_+$) representing loss size. The Gamma distribution, with mean m_l and variance v_l as parameters, is commonly used for modeling such losses. The presence of losses,

⁴See Yang, Wang, and Chen (2008) for a survey of studies on the effect of catastrophes on the stock markets

Figure 1.7: USA flood damage per capita and insurer sample stock price path.

and not independent jump processes, drives the jumps in the stock price. While it is possible to add a separate jump component to the stock price which is not due to a catastrophe, the most important contribution to the price is the catastrophe effect. Figure 1.7 depicts the size of losses due to floods in the United States (scaled by population size) and a sample stock price that responds negatively to such losses generated by the above model. Note that the stock price curve has logarithmic scale.

1.2 Option Pricing Problem

As previously mentioned, options are financial contracts whose value depends on the value of other underlying variables. These variables are typically prices of tradable securities, such as the price of an IBM share, or commodities, such as the price of natural gas delivered to the New York Henry hub. However, there exist options that depend on variables that are not traded on any exchange, such as the amount of monthly rain or temperature in a particular location or losses incurred by an insurance company due to natural disasters.

The issuer, or the seller, of the option gives its holder, or the buyer, a right, but not an obligation, to do something — *exercise* the option by a certain date (*expiration* or *maturity* date), denoted by T . For tradable underlying variables, this embedded right is typically a right to purchase (*call* option) or sell (*put* option) the underlying security or commodity at a prescribed price, referred to as the *strike price*, denoted by K . For a call option, if the market value of the underlying asset at the time of exercise is larger than the strike price, the holder of

the option can make money by exercising the option (hereby buying the underlying asset at the strike price) and then selling the asset in the market and pocketing the difference. Conversely, for a put option, if the value of the underlying asset at the time of exercise is smaller than the strike price, the holder of the option can make money by buying the underlying asset in the market and exercising the option (hereby selling the underlying asset for the strike price) and pocketing the difference. For non-tradable securities the option gives the holder a right to a certain payoff that depends on the non-tradeable underlying (e.g., heating or cooling degree days for temperature derivatives).

Thus, at expiry or at some exercise time, an option can potentially have non-zero value to the holder (this value is deterministic given the value of the underlying either at maturity or throughout the life of the contract). The question at the heart of this research is computing the value of the option to the holder at any time prior to expiry or exercise. Ever since the seminal works of Black and Scholes (1973) and Merton (1973), who revolutionized our understanding of financial contracts with embedded options, a tremendous amount of research has been done in this area.

Let $V(t, \mathbf{S}(t))$ denote the price at time t of an option, written on a vector of d underlying price indices $\mathbf{S}(t)$, whose components are $S_k(t)$, with a T -maturity payoff of $\varphi(\mathbf{S}(T))$. It is well known that, in an arbitrage-free and frictionless market, the value of a European option is the discounted expectation of its payoff under a, not necessarily unique, risk-neutral measure $\mathbb{Q}$ ⁵. That is,

$$V(t, \mathbf{S}(t)) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \varphi(\mathbf{S}(T)) \right], \quad (1.8)$$

where the expectation is taken with respect to the information, or filtration, $\mathcal{F}_t$, available at time t . Here and in the remainder of this thesis, the risk-free interest rate r is assumed to be constant. Other types of derivatives can be priced by deriving appropriate extensions of the above equation. When the underlying index follows a diffusion process, the risk-neutral measure is indeed unique; however, in the more interesting case of exponential Lévy models, many equivalent risk-neutral measures exist. Nonetheless, the point of view taken here is that a trader is using such a model to price derivative instruments and therefore is modeling directly under a particular risk-neutral measure — possibly induced through a calibration procedure.

1.3 Numerical Methods

Black and Scholes (1973) and Merton (1973) demonstrated, based on the assumption that stock prices follow a geometric Brownian motion, that a replicating strategy reduces the option pricing

⁵See Harrison and Pliska (1981) for more details on valuation of contingent claims

problem to solving a partial differential equation (PDE), which is independent of the return of the asset. Under the other models for the underlying discussed in the previous section, the option price satisfies a more complex partial integro-differential equation (PIDE). The differential formulation is dual and equivalent to the expectation form (1.8). The P(I)DE can be solved analytically to obtain a closed-form solution to the option price only in a few simple cases. For pricing of complex derivatives and/or under complex underlying asset models a number of alternative numerical methods have been developed. Monte Carlo methods are effective if the option payoff is dependent on several underlying variables or if there is complex dependence of option payoff on the history of the underlying variable(s). Tree methods and finite difference methods are efficient in pricing options if the option can be exercised prior to maturity. Fast Fourier transform (FFT) methods are especially effective if the underlying variables are modeled using jump-diffusion or exponential Lévy models. Finally, several other methods address various computational challenges and improve the performance of the aforementioned approaches. A brief overview of the main methods used in option pricing is presented below.

1.3.1 Monte Carlo Simulation

Boyle (1977) was the first to apply the Monte Carlo method to price options. The underlying idea of the method is to generate N trajectories $\mathbf{S}^{\{n\}}$ of the underlying asset under the pricing measure $\mathbb{Q}$. The expectation in equation (1.8) can then be found by averaging the discounted option payoff:

$$\mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \varphi(\mathbf{S}(T)) \right] \approx \frac{1}{N} \sum_{n=1}^N e^{-r(T-t)} \varphi(\mathbf{S}^{\{n\}}(T)). \quad (1.9)$$

Different types of derivatives can be priced this way as the only requirement is the ability to generate samples from the distribution of the underlying asset. Hence, multi-dimensional problems can be handled naturally by the Monte Carlo method and it does not suffer from the curse of dimensionality⁶. The error of the method is $O(1/\sqrt{N})$; using variance reduction techniques, such as antithetic variables and control variates, can significantly reduce the pricing error. The use of deterministic low-discrepancy sequences can improve the convergence order to $O((\log N)^d/N)$, however, for large d the theoretical benefits come into effect only at large sample sizes.

1.3.2 Tree Methods

The underlying idea behind tree methods is to approximate the continuous-time model for the evolution of spot prices by a discrete-time framework. The first such framework was the

⁶See Boyle, Broadie, and Glasserman (1997) and Glasserman (2003) for an overview of Monte Carlo methods in the context of option pricing

binomial option pricing model introduced by Cox, Ross, and Rubinstein (1979), followed by the trinomial option pricing model of Boyle (1986). Madan, Milne, and Shefrin (1989) generalize these models further to the multinomial case. Figure 1.8 depicts the evolution tree of a stock price under the binomial pricing model.

Figure 1.8: Spot price evolution in the binomial option pricing model

In a discrete-time framework the evolution of a spot price is traced on a lattice, where each node in the lattice represents a possible value for the price at a given point in time. Once a tree is generated, option valuation is done iteratively, computing option values at earlier nodes from the option values at later nodes. Such methodology allows for path-dependent and early-exercise options to be readily handled, in addition to having faster convergence rate than Monte Carlo methods. Another advantage of the discrete-time framework is the ability to handle discrete dividends paid by the underlying asset.

1.3.3 Finite Difference Methods

Finite difference methods have been used extensively to obtain a numerical solution to the pricing PIDE. Such methods are efficient, precise and can be applied to price various types of options. The underlying idea behind finite difference methods is to discretize the space-time domain of the PIDE into a finite mesh of points and approximate the differential operator by finite differences. The solution to the equation at any time prior to maturity is obtained on the mesh points by taking steps backwards in time, with the terminal condition at maturity is given by the option payoff.

The operator acting on the option value function V in the pricing PIDE can be written succinctly as a sum of diffusion and integral components $\mathcal{D}$ and $\mathcal{J}$, respectively:

$$(\partial_t + \mathcal{D} + \mathcal{J}) V(t, \mathbf{x}) = 0. \quad (1.10)$$

A finite difference scheme for solving the pricing equation can be expressed as

$$[\mathbf{I} - (1 - \alpha_d)\Delta t \mathcal{D} - (1 - \alpha_j)\Delta t \mathcal{J}] \mathbf{V}_{m-1} = [\mathbf{I} + \alpha_d\Delta t \mathcal{D} + \alpha_j\Delta t \mathcal{J}] \mathbf{V}_m, \quad (1.11)$$

where $\mathbf{V}_m$ are the option values on a discrete grid at time t_m , $\mathbf{D}$ and $\mathbf{J}$ are the matrices associated with the discretization of $\mathcal{D}$ and $\mathcal{J}$, respectively, and $\mathbf{I}$ is the identity matrix. The constants α_d and α_j specify whether the diffusion and integral terms are treated using explicit Euler ($\alpha = 1$), implicit Euler ($\alpha = 0$) or Crank-Nicolson ($\alpha = 1/2$) schemes⁷. While the explicit Euler scheme is the most straightforward to implement, stability considerations impose a time step restriction, which can become severe on a large grid. The implicit Euler scheme overcomes this stability restriction at a cost of solving a system of linear equations at each time step. The Crank-Nicolson scheme, like the implicit Euler scheme, overcomes the stability restriction of the explicit Euler scheme and has second-order convergence in time (as opposed to first-order convergence in time for the explicit and implicit Euler schemes).

In the absence of jumps, implicit treatment of the diffusion term can be done efficiently since the matrix $\mathbf{D}$ is banded — the pricing PDE can be discretized using standard divided differences to approximate the first- and second-order derivatives. In one spatial dimension, whether the approximation scheme is carried out explicitly, implicitly or through a weighted scheme, the resulting system is banded (often tri-diagonal) and leads to very efficient numerical approximations. When jumps are present, the matrix $\mathbf{J}$ is dense (and quite complex in the multi-dimensional setting) and implicit treatment of the integral term is computationally expensive. Thus, a number of approaches have been proposed to effectively treat the integral term while retaining high-order convergence. Although the methods are quite diverse, they all treat the integral and diffusion terms of the PIDE separately. Invariably, the integral term is evaluated explicitly to avoid solving a dense system of linear equations. In addition, the FFT algorithm may be employed to speed up the computation of the integral term (which can be regarded as a convolution) and/or its inverse. Unfortunately, these methods require several approximations such as:

- in infinite activity processes, where the probability density function blows up in the neighborhood of zero, small jumps are approximated by a diffusion and incorporated into the diffusion term;
- the integral term must be localized to the bounded domain of the diffusion term, i.e., large jumps are truncated;
- the option price behavior outside the solution domain must be assumed; and
- the separate treatment of diffusion and integral components requires that function values be interpolated between the diffusion and integral grids to compute the convolution term.

⁷See Wilmott, Howison, and Dewynne (1995) for a discussion of finite difference schemes in the context of option pricing without jumps

These factors together make finite difference methods for option pricing under jump models quite complex, and potentially prone to accuracy and stability problems, especially for path-dependent claims. As a consequence, many methods are tuned to a specific class of Lévy model. Moreover, for infinite activity Lévy processes, finite difference methods typically suffer from slow convergence. A summary of such finite difference methods is presented below.

Explicit-Implicit Method

Explicit-implicit methods have traditionally been used to solve diffusion-reaction or diffusion-convection PDEs⁸. Typically, the diffusion term is treated implicitly while the reaction/convection term is treated explicitly to avoid the expensive solution of the associated dense linear system. Cont and Tankov (2004) and Cont and Voltchkova (2005) propose an explicit-implicit scheme that treats the diffusion term and the integral term asymmetrically. This algorithm corresponds to $\alpha_d = 0$ and $\alpha_j = 1$ in equation (1.11) and results in the following time-stepping scheme:

$$[\mathbf{I} - \Delta t \mathbf{D}] \mathbf{V}_{m-1} = [\mathbf{I} + \Delta t \mathbf{J}] \mathbf{V}_m. \quad (1.12)$$

The advantages of this method is its stability and efficiency since a costly inversion of a dense matrix is not required. However, the differential and integral parts are treated asymmetrically which leads to loss of accuracy, particularly for long dated options. Cont and Voltchkova (2005) show that the finite difference scheme is consistent, stable and converges to a viscosity solution of the PIDE (i.e., the option price). The explicit-implicit method is independent of the chosen jump process and can handle a wide variety of path-dependent options. Unfortunately, the method does not use the FFT algorithm to speed up the computation of the integral term, reducing its performance.

FFT - Alternating Directions Implicit Method

The FFT - Alternating Direction Implicit (FFT-ADI) algorithm, developed by Andersen and Andreasen (2000), aims to improve upon the explicit-implicit method by treating the differential and integral components of the differential operator symmetrically. This is done by splitting each time step into two half-steps and alternating the direction of the method for each component on the two half-steps.

The first half-step is advanced using a scheme for which the diffusion term is treated implicitly while the integral term is treated explicitly, corresponding to $\alpha_d = 0, \alpha_j = 1$:

$$\left[\mathbf{I} - \frac{\Delta t}{2} \mathbf{D} \right] \mathbf{V}_{m-1/2} = \left[\mathbf{I} + \frac{\Delta t}{2} \mathbf{J} \right] \mathbf{V}_m. \quad (1.13)$$

⁸See Ascher, Ruuth, and Wetton (1997) for an overview of various explicit-implicit schemes

An important improvement upon the explicit-implicit method is to use the FFT algorithm to compute the integral term $\mathbf{J}\mathbf{v}$ by regarding it as a convolution of the option value and jump density functions. On the second half-step, the differential term is treated explicitly while the integral term is treated implicitly, corresponding to $\alpha_d = 1, \alpha_j = 0$:

$$\left[\mathbf{I} - \frac{\Delta t}{2} \mathbf{J} \right] \mathbf{V}_{m-1} = \left[\mathbf{I} + \frac{\Delta t}{2} \mathbf{D} \right] \mathbf{V}_{m-1/2}. \quad (1.14)$$

Even though the inversion of the operator $\mathbf{J}$ is required, it can be performed efficiently using the FFT algorithm (in the first half-step, the FFT is used to compute the product $\mathbf{J}\mathbf{V}$). By regarding $\mathbf{J}\mathbf{V}$ as a convolution of the option value and the distribution density functions and using the convolution property of Fourier transforms (Fourier transform of a convolution of two functions is the product of their respective Fourier transforms), the inversion can be performed very efficiently using forward and backward FFT evaluations. Thus, the FFT-ADI method is very efficient, requiring only a few evaluations of the FFT method per time step.

Fixed-Point Iteration Method

The fixed-point iteration method improves upon the explicit-implicit method by computing the solution to the PIDE at the next time step using multiple iterations of a finite difference scheme. In general terms, the fixed-point iteration method can be expressed as

$$\begin{aligned} [\mathbf{I} - (1 - \alpha_d)\Delta t \mathbf{D}] \mathbf{V}_{m-1}^{(k)} &= [(1 - \alpha_j)\Delta t \mathbf{J}] \mathbf{V}_{m-1}^{(k-1)} + \\ &[\mathbf{I} + \alpha_d\Delta t \mathbf{D} + \alpha_j\Delta t \mathbf{J}] \mathbf{V}_m, \end{aligned} \quad (1.15)$$

where $\mathbf{V}_{m-1}^{(k)}$ is the solution at iteration k . The algorithm starts with $\mathbf{V}_{m-1}^{(0)} = \mathbf{V}_m$ and iterates until the error $e^{(k)} = |\mathbf{V}_{m-1}^{(k)} - \mathbf{V}_{m-1}^{(k-1)}|$ satisfies a given tolerance requirement.

d'Halluin, Forsyth, and Vetzal (2005) use Crank-Nicolson time-stepping ($\alpha_d = \alpha_j = 1/2$). For the evaluation of the integral term $\mathbf{J}\mathbf{v}$, an efficient procedure using Lagrange interpolation is developed, along the lines of the convolution approach of the FFT-ADI method discussed prior. An additional advantage of the fixed-point iteration method is the ability to handle early-exercise in American options by incorporating a penalty method to solve the associated linear complementarity problem.

Almendral and Oosterlee (2005) propose a different iterative finite difference method using second-order backward-differentiation for the time discretization. Again, an FFT-based approach is used to evaluate the convolution term.

Implicit-Explicit Runge-Kutta Method

Implicit-explicit Runge-Kutta schemes have been extensively used for solution of problems with a stiff or non-local term⁹. An implicit-explicit Runge-Kutta method can be defined for PIDE (2.2) by

$$\begin{cases} \mathbf{V}^{\{k\}} &= \mathbf{V}_m + \Delta t \sum_{l=1}^{k-1} \tilde{a}_{kl} \mathbf{J} \mathbf{V}^{\{l\}} + \Delta t \sum_{l=1}^k a_{kl} \mathbf{D} \mathbf{V}^{\{l\}}, \quad k = 1, \dots, c \\ \mathbf{V}_{m-1} &= \mathbf{V}_m + \Delta t \sum_{k=1}^c \tilde{b}_k \mathbf{J} \mathbf{V}^{\{k\}} + \Delta t \sum_{k=1}^c b_k \mathbf{D} \mathbf{V}^{\{k\}}, \end{cases} \quad (1.16)$$

where c is the number of stages and the different choices of $a_{kl}, \tilde{a}_{kl}, b_k, \tilde{b}_k$ yield distinct numerical methods.

Briani, Natalini, and Russo (2007) apply various second- and third-order numerical schemes based on the above equation in the context of option pricing. They consider the stability and convergence properties of such numerical schemes, showing that implicit-explicit Runge-Kutta schemes are high-order schemes under weak stability time step restrictions.

1.3.4 Fast Fourier Transform Methods

Carr and Madan (1999) first applied the FFT algorithm to pricing options under exponential Lévy models. The underlying idea of their approach is to develop an analytic expression for the Fourier transform of the option value function. The derivation presented below is slightly different than the one presented in their work. The expectation in equation (1.8) can be expressed as a convolution of the option payoff φ and stock price density $f_{\mathbf{X}}$:

$$\mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \varphi(\mathbf{S}(T)) \right] = e^{-r(T-t)} \int_{-\infty}^{\infty} \varphi(\mathbf{S}(t) e^{\mathbf{y}}) \cdot f_{\mathbf{X}}(\mathbf{y}) d\mathbf{y}. \quad (1.17)$$

The desired expression for the Fourier transform of the option value, V , is obtained by applying the Fourier transform to the integral above and using the convolution property of Fourier transforms:

$$\mathcal{F}[V](t, \boldsymbol{\omega}) = e^{-r(T-t)} \mathcal{F}[\varphi](\boldsymbol{\omega}) \mathcal{F}[f_{\mathbf{X}}](-\boldsymbol{\omega}). \quad (1.18)$$

The Fourier transform of the density function $\mathcal{F}[f_{\mathbf{X}}](\boldsymbol{\omega})$ (also known as the characteristic function) is available in closed form for all exponential Lévy models. The Fourier transform of simple payoff functions $\mathcal{F}[\varphi](\boldsymbol{\omega})$ can be computed analytically (a damping factor may be required to avoid singularities along the real axis in the Fourier transforms of some payoff functions). The Fourier transform of the option value function $\mathcal{F}[V](\boldsymbol{\omega})$ can then be computed

⁹See Pareschi and Russo (2000), Pareschi and Russo (2005), and Liu and Zou (2006)

analytically and option values in real space can be then obtained by applying the inverse Fourier transform. Such computation can be done efficiently using the FFT algorithm.

A number of important extensions to the above method have been developed by Dempster and Hong (2000), Raible (2000), Lewis (2001), Reiner (2001), Andricopoulos, Widdicks, Duck, and Newton (2003), and O’Sullivan (2005) to price multi-asset and path-dependent options with general payoffs. The method of Lord, Fang, Bervoets, and Oosterlee (2008) is discussed in greater detail in Section 2.1.

1.3.5 Other Methods

Other methods that do not fall into either of the above categories are briefly summarized below. Chiarella, El-Hassan, and Kucera (1999) and Chiarella and Ziogas (2005) develop a path-integral framework using the Fourier-Hermite series expansion for the continuous representation of the underlying asset price. Albanese, Jaimungal, and Rubisov (2001) introduce a pricing model, based on the method of lines, which postulates only the discretization of calendar time, with each key date in the model corresponding to a continuous line for stock prices. Boyarchenko and Levendorskii (2002) utilize the Wiener-Hopf factorization to obtain option prices in terms of the resolvents of the supremum and infimum processes and derive explicit formulas for these factors. Matache, von Petersdorff, and Schwab (2004) and Matache, Nitsche, and Schwab (2005) discretize the pricing PIDE in space using a wavelet Galerkin method with compression of the moment matrix of the jump component. The Hilbert transform is utilized by Feng and Linetsky (2008) to derive a computational algorithm which relies on FFT-based Toeplitz matrix-vector multiplication. Finally, Fang and Oosterlee (2008) develop an algorithm based on the Fourier-cosine series expansion of the density function.

1.4 Motivation for Research

As outlined in the previous section, a wide array of numerical methods for pricing of financial derivatives under different spot price models have been developed. A competitive method should rank high in all (or at least most) of the following criteria:

- **Precision, Speed and Convergence**

Precise computation of option prices and Greeks (sensitivities to model parameters and spot price movements) is the *raison d’être* of any numerical method. In the context of institutional and block trading, where large blocks of derivatives change hands, differences of a fraction of a cent in the computed price can lead to significant portfolio gains/losses and thus can not be ignored. While precise computation is paramount, the speed of

the calculation is also important. In electronic/algorithmic trading setting, market participants must be able to react to changes in market conditions within hundredths of second. The tradeoff between speed and precision characterizes the overall performance of a method. The order of convergence of a numerical method, that is, the rate at which the error reduces as the number of data increases, plays an important role in the performance of the method. In most cases, the higher order of convergence is preferred, as it usually leads to methods that give rise to smaller errors in the same amount of time (or to the same error in less time) as lower order methods. The majority of the methods for financial problems have either first- or second-order convergence, therefore, in this thesis, the second-order of convergence is highly desirable.

- **Efficient handling of path-independent and discretely-monitored derivatives**

Current state-of-the-art numerical methods for path-independent derivatives (depending only on the terminal value of the asset price) are quite efficient. Any competitive method must be able to compute derivative prices quickly using a single time step (for finite difference type methods), several FFT evaluations (for transform-based methods) or a small number of samples from a random number generator (for Monte Carlo-based methods). For discretely-monitored derivatives, the computational time should scale linearly with the number of monitoring dates. Furthermore, the method should be applicable to pricing of options with non-standard payoff functions.

- **Ability to handle path-dependent and multi-asset derivatives**

While a number of efficient methods already exist to price path-independent options, a more salient problem is the quick and precise valuation of highly path-dependent and/or multi-asset derivatives, possibly with embedded features, such as exercise prior to maturity, termination on barrier breach and others.

- **Generic handling of various spot price models**

With financial derivatives being traded in distinct markets, it is important for a method to handle the multitude of spot price models available in the literature. Also, the method should be general enough so that different models can be used without significant modification to the algorithm.

- **Utilization of multi-core architectures**

Moore's law states that the number of transistors that can be placed inexpensively on a computer processor doubles approximately every 18 months. The trend, first observed in 1965, has consistently persisted, although recently, as transistors have become smaller, heat dissipation and power consumption have become major problems. Multi-core archi-

tectures, such as Graphics Processing Units (GPUs), Cell Broadband Engines (Cell BEs) and multi-core Central Processing Units (CPUs), circumvent these problems by increasing the number of chips in a processor, rather than increasing the number of transistors in a chip. Numerical methods must be able to utilize the multiple computing cores effectively to reap the full benefit of such architectures.

Unfortunately, the vast majority of the methods summarized in Section 1.3 fall short in more than one of the above categories. The motivation for this research is to develop a fast and precise numerical method to price complex, path-dependent derivative contracts. The method should be general enough to handle various spot price models and parallelizable on multi-core architectures.

In line with these goals, this thesis develops a comprehensive Fourier transform-based framework to compute the evolution of derivative prices in time. The developed methods are precise, quick, and rapidly convergent. European options can be priced efficiently using one time step of the algorithm, Bermudan options do not require time-stepping between the monitoring dates, and other highly path-dependent options, such as American and barrier options, can be handled efficiently. The methodology allows pricing under various spot price models by simply supplying the appropriate characteristic exponent, without necessitating changes to the pricing algorithm. Furthermore, the developed methods are easily parallelizable due to the inherently parallel nature of the numerical computation of Fourier transforms.

While the Fourier transform methodology can be readily applied to independent-increment exponential Lévy processes, state-dependent models (where drift or volatility terms are functions of the current asset price) require special consideration. For mean-reverting models discussed in Chapter 3 a particular frequency transformation is performed to handle the state-dependent drift term. However, the Fourier transform-based framework developed in this thesis currently can not be applied to other classes of models, such as the stochastic volatility model of Heston (1993) and the local volatility model of Derman and Kani (1994).

Chapter 2

Fourier Space Time-stepping Method

2.1 Introduction

In this chapter, the FST method for pricing options is developed. The method avoids the problems associated with finite difference methods and utilizes the advantages of Fourier transform methods by transforming the PIDE into Fourier space. One of the advantages of working directly in Fourier space is that the characteristic exponent of an independent-increment stochastic process can be factored out of the Fourier transform of the PIDE. Consequently, the Fourier transform can be applied to the PIDE to obtain a linear system of easily solvable ordinary differential equations (ODE). Furthermore, the characteristic exponent is available in closed form for all independent-increment processes through the Lévy-Khintchine formula. This makes the FST method quite flexible and generic — contingent claims on any exponential Lévy stock price processes can be priced with no additional modifications to the algorithm. The FST method naturally leads to a symmetric treatment of the diffusion and jump terms. Moreover, it can be efficiently applied to pricing of multi-dimensional options with path-dependency.

For path-independent options, prices for a range of spots can be obtained in a single time step. The closed-form expression for the Fourier transform of the option payoff is not required, making the FST method easily applicable to options with non-standard payoffs. For exotic, path-dependent options, the FST method is demonstrated to handle Bermudan, American and barrier styled clauses. Since the FST method provides exact pricing results between monitoring times, it is significantly more efficient and accurate when compared with finite-difference methods for valuing Bermudan options. Furthermore, the method allows prices from one monitoring time to be projected back to a second monitoring time in one step of the algorithm. Contrastingly, finite-difference schemes require time-stepping between monitoring dates resulting

in further pricing biases and speed reduction.

During the refereeing stage of Jackson, Jaimungal, and Surkov (2008), after the first draft was made publicly available online, the authors learned of the contemporaneous work of Lord, Fang, Bervoets, and Oosterlee (2008), who independently developed a similar method, called CONV. The derivation of CONV, however, is quite different, as the authors utilize the convolution representation to derive their pricing method. While the two approaches lead to the same pricing formula, there are several important advantages of the PIDE transform approach. In particular, it leads naturally to a penalty method for American options, outlined in Section 2.6.2, and enables the FST method to be extended to mean-reverting and regime-switching frameworks, detailed in Chapter 3 and Section 4.2, respectively.

Through numerical experiments, this chapter establishes that the order of convergence of the method for pricing single-asset options is 2 in space and 1 in time, for path-dependent options, such as American and barrier options. For American options, the penalty method of Forsyth and Vetzal (2002) is extended to the FST framework and attains quadratic convergence in time. Moreover, the FST method is computationally efficient: since only two FFTs are required per time step, its computational complexity is $O(M N^d \log_2 N)$, where N is the number of spatial grid points in each dimension, d is the number of spatial dimensions and M is the number of time steps.

The outline of the remainder of this chapter is as follows. Section 2.2 presents the class of exponential Lévy models for equity prices. Section 2.3 introduces the option pricing PIDE and solves it in Fourier space. Section 2.4 formulates the FST option pricing method based on the PIDE solution. Section 2.5 discusses the stability of the FST algorithm and develops methods for estimating its convergence order. Section 2.6 presents pricing results for various options and stock price models. Finally, Section 2.7 develops the greekFST method for computation of option Greeks.

2.2 Spot Price Model

If the underlying index follows an exponential Lévy process, then the price process can be written as $\mathbf{S}(t) = \mathbf{S}(0)e^{\mathbf{X}(t)}$, where $\mathbf{X}(t)$ is a Lévy process with characteristic triplet $(\boldsymbol{\gamma}, \boldsymbol{\Sigma}, \boldsymbol{\nu})$, $\boldsymbol{\gamma}$ represents the vector of unadjusted-drifts, $\boldsymbol{\Sigma}$ represents the variance-covariance matrix of the diffusions, and $\boldsymbol{\nu}$ is the multi-dimensional Lévy density. In this case, the process $\mathbf{X}(t)$ admits

the following canonical Lévy-Itô decomposition into its diffusion and jump components:

$$\begin{aligned}\mathbf{X}(t) &= \gamma t + \mathbf{W}(t) + \mathbf{J}^l(t) + \lim_{\epsilon \searrow 0} \mathbf{J}^\epsilon(t), \\ \mathbf{J}^l(t) &= \int_0^t \int_{|y| \geq 1} \mathbf{y} \tilde{\nu}(d\mathbf{y} \times ds), \\ \mathbf{J}^\epsilon(t) &= \int_0^t \int_{\epsilon \leq |y| < 1} \mathbf{y} [\tilde{\nu}(d\mathbf{y} \times ds) - \nu(d\mathbf{y} \times ds)] .\end{aligned}$$

Here $\mathbf{W}(t)$ is a standard Brownian motion, $\tilde{\nu}(d\mathbf{y} \times ds)$ is a Poisson random measure counting the number of jumps of size $\mathbf{y}$ occurring at time s , and $\nu(d\mathbf{y} \times ds) = \nu(d\mathbf{y}) ds$ is its compensator. Note that $\mathbf{J}^l(t)$ and $\mathbf{J}^\epsilon(t)$ carry the interpretation of large and small jumps, respectively. If the model has finite activity ($\int_{\mathbb{R}/\{0\}} (|\mathbf{y}| \wedge 1) \nu(d\mathbf{y}) < +\infty$) then there is no need to truncate small jumps and they can be lumped together with large jumps. If the model has infinitely many small jumps, ν may have a singularity at 0. Thus, the small jumps integral must be centered, i.e., replaced by the compensated version, to obtain convergence.

By enforcing the risk-neutrality condition, the drift is uniquely determined once the volatility and Lévy density are specified. In particular, γ is chosen such that

$$\mathbb{E}_0 \left[e^{\mathbf{X}_j(1)} \right] = e^r \quad \Rightarrow \quad \Psi(-i\mathbf{1}_j) = r ,$$

for each $j = 1, \dots, d$, where $\mathbf{1}_j$ is the vector with zeros everywhere except a single entry of 1 at dimension j and $\Psi(\boldsymbol{\omega})$ denotes the characteristic exponent of the d -dimensional Lévy process, provided explicitly by the Lévy-Khintchine formula:

$$\Psi(\boldsymbol{\omega}) = i\boldsymbol{\gamma}'\boldsymbol{\omega} - \frac{1}{2}\boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega} + \int_{\mathbb{R}^n} \left(e^{i\boldsymbol{\omega}'\mathbf{y}} - 1 - i\mathbb{1}_{\{|\mathbf{y}| < 1\}} \boldsymbol{\omega}'\mathbf{y} \right) \nu(d\mathbf{y}) , \quad (2.1)$$

where $\boldsymbol{\omega}'$ represents the transpose of the vector $\boldsymbol{\omega}$.

Within this framework, the BSM model is recovered by setting the Lévy density to zero. Furthermore, (one-dimensional) jump-diffusion models, in which the log-stock price contains a diffusive component together with jumps occurring at Poisson times, are recovered by setting $\nu(dy) = \lambda f_Y(y) dy$ where λ is the activity rate of the Poisson process and $f_Y(y)$ is the probability density of the jumps.

2.3 PIDE Solution

Using the fundamental theorem of asset pricing, it is well-known that the discount-adjusted and log-transformed price process $v(t, \mathbf{X}(t)) \triangleq e^{r(T-t)} V(t, \mathbf{S}(0) e^{\mathbf{X}(t)})$ is a martingale under the measure $\mathbb{Q}$. Consequently, the associated drift term of its defining SDE is identically zero.

Applying the zero-drift condition on $v(t, \mathbf{x})$, together with its boundary condition at maturity, leads to the pricing PIDE:

$$\begin{cases} (\partial_t + \mathcal{L}) v(t, \mathbf{x}) &= 0, \\ v(T, \mathbf{x}) &= \varphi(\mathbf{S}(0) e^{\mathbf{x}}), \end{cases} \quad (2.2)$$

where $\mathcal{L}$ is the infinitesimal generator of the multi-dimensional Lévy process and acts on twice-differentiable functions $g(\mathbf{x})$ as follows:

$$\mathcal{L}g(\mathbf{x}) = (\gamma' \partial_{\mathbf{x}} + \frac{1}{2} \partial_{\mathbf{x}}' \Sigma \partial_{\mathbf{x}}) g(\mathbf{x}) + \int_{\mathbb{R}^n / \{\mathbf{0}\}} (g(\mathbf{x} + \mathbf{y}) - g(\mathbf{x}) - \mathbb{1}_{\{|\mathbf{y}| < 1\}} \mathbf{y}' \partial_{\mathbf{x}} g(\mathbf{x})) \nu(d\mathbf{y}). \quad (2.3)$$

Fourier and Laplace transforms have been used extensively to solve PDEs, either by transforming the equation into an ODE or expressing the solution as an infinite series¹. The aim of this section is to develop a Fourier transform-based methodology for solving PIDEs of the form (2.2). The main advantage of such approach is that the PIDE can be handled efficiently, without the additional complexities associated with the integral term. Additionally, the algorithm is applicable to any independent-increment stock price model which admits a closed-form characteristic function.

A function in the space domain $g(\mathbf{x})$ can be transformed to a function in the frequency domain $\hat{g}(\boldsymbol{\omega})$, where $\boldsymbol{\omega}$ is given in radians per second, and vice-versa using the continuous Fourier transform (CFT):

$$\mathcal{F}[g](\boldsymbol{\omega}) \triangleq \int_{-\infty}^{\infty} g(\mathbf{x}) e^{-i\boldsymbol{\omega}'\mathbf{x}} d\mathbf{x} \quad \text{and} \quad \mathcal{F}^{-1}[\hat{g}](\mathbf{x}) \triangleq \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{g}(\boldsymbol{\omega}) e^{i\boldsymbol{\omega}'\mathbf{x}} d\boldsymbol{\omega}. \quad (2.4)$$

CFT is a linear operator that maps spatial derivatives ∂_x into multiplications in the frequency domain:

$$\mathcal{F}[\partial_{\mathbf{x}}^n g](\boldsymbol{\omega}) = i\boldsymbol{\omega} \mathcal{F}[\partial_{\mathbf{x}}^{n-1} g](\boldsymbol{\omega}) = \dots = (i\boldsymbol{\omega})^n \mathcal{F}[g](\boldsymbol{\omega}). \quad (2.5)$$

Consequently, applying the CFT to the infinitesimal generator $\mathcal{L}$ of $\mathbf{X}(t)$, defined by equation (2.3), allows the characteristic exponent of $\mathbf{X}(t)$ to be factored out:

$$\begin{aligned} \mathcal{F}[\mathcal{L}v](t, \boldsymbol{\omega}) &= \left\{ i\boldsymbol{\gamma}'\boldsymbol{\omega} - \frac{1}{2} \boldsymbol{\omega}' \Sigma \boldsymbol{\omega} + \int_{\mathbb{R}^n} \left(e^{i\boldsymbol{\omega}'\mathbf{y}} - 1 - i\mathbb{1}_{\{|\mathbf{y}| < 1\}} \boldsymbol{\omega}'\mathbf{y} \right) \nu(d\mathbf{y}) \right\} \mathcal{F}[v](t, \boldsymbol{\omega}) \\ &= \Psi(\boldsymbol{\omega}) \mathcal{F}[v](t, \boldsymbol{\omega}). \end{aligned} \quad (2.6)$$

The Lévy densities and characteristic exponents of various stock price models are provided in Table 2.1. Furthermore, taking the Fourier transform of both sides of the PIDE (2.2) leads to

$$\begin{cases} \partial_t \mathcal{F}[v](t, \boldsymbol{\omega}) + \Psi(\boldsymbol{\omega}) \mathcal{F}[v](t, \boldsymbol{\omega}) &= \mathbf{0}, \\ \mathcal{F}[v](T, \boldsymbol{\omega}) &= \mathcal{F}[\varphi](\boldsymbol{\omega}). \end{cases} \quad (2.7)$$

¹See Strauss (1992) and Taylor (1997) for further mathematical background

Model	Lévy Density $\nu(dy)$	Characteristic Exponent $\Psi(\omega)$
Black-Scholes-Merton	0	$i\gamma\omega - \frac{\sigma^2\omega^2}{2}$
Merton jump-diffusion	$\frac{\lambda}{\sqrt{2\pi\tilde{\sigma}^2}} e^{-\frac{1}{2}((y-\tilde{\mu})/\tilde{\sigma})^2}$	$i\gamma\omega - \frac{\sigma^2\omega^2}{2} + \lambda(e^{i\tilde{\mu}\omega - \tilde{\sigma}^2\omega^2/2} - 1)$
Kou jump-diffusion	$\lambda \left(\frac{\eta_p}{\eta_+} e^{-y/\eta_+} \mathbb{1}_{\{y>0\}} + \frac{1-\eta_p}{\eta_-} e^{- y /\eta_-} \mathbb{1}_{\{y<0\}} \right)$	$i\gamma\omega - \frac{\sigma^2\omega^2}{2} + \lambda \left(\frac{\eta_p}{1-i\omega\eta_+} + \frac{1-\eta_p}{1+i\omega\eta_-} - 1 \right)$
Variance Gamma	$\frac{1}{\mu y } e^{C_1 y - C_2 y }$	$-\frac{1}{\mu} \ln(1 - i\gamma\mu\omega + \frac{\sigma^2\mu\omega^2}{2})$
Normal Inverse Gaussian	$\frac{C_3}{ y } e^{C_1 y} K_1(C_4 y)$	$-\frac{1}{\mu} (\sqrt{1 - 2i\gamma\mu\omega + \sigma^2\mu\omega^2} - 1)$
Carr-Geman-Madan-Yor	$\frac{C}{ y ^{1+Y}} (e^{-G y } \mathbb{1}_{\{y<0\}} + e^{-M y} \mathbb{1}_{\{y>0\}})$	$C\Gamma(-Y) [(M - i\omega)^Y - M^Y + (G + i\omega)^Y - G^Y]$

Table 2.1: Lévy densities and characteristic exponents of various exponential Lévy models. Here γ and σ are the drift and volatility of the driving Brownian motion (if applicable), $C_1 = \frac{\gamma}{\sigma^2}$, $C_2 = \frac{\sqrt{\gamma^2 + 2\sigma^2/\mu}}{\sigma^2}$, $C_3 = \frac{\sqrt{\gamma^2 + \sigma^2/\mu}}{\pi\sigma\sqrt{\mu}}$, $C_4 = \frac{\sqrt{\gamma^2 + \sigma^2/\mu}}{\sigma^2}$ and $K_p(x)$ is the modified Bessel function of the second kind.

The PIDE is therefore transformed into a d -parameter family of ODEs (2.7) parameterized by $\boldsymbol{\omega}$. Given the value of $\mathcal{F}[v](t, \boldsymbol{\omega})$ at time $t_2 \leq T$, the system is easily solved to find the value at time $t_1 < t_2$:

$$\mathcal{F}[v](t_1, \boldsymbol{\omega}) = \mathcal{F}[v](t_2, \boldsymbol{\omega}) \cdot e^{\Psi(\boldsymbol{\omega})(t_2 - t_1)}. \quad (2.8)$$

Taking the inverse transform leads to the final result

$$v(t_1, \mathbf{x}) = \mathcal{F}^{-1} \left[\mathcal{F}[v](t_2, \boldsymbol{\omega}) \cdot e^{\Psi(\boldsymbol{\omega})(t_2 - t_1)} \right] (\mathbf{x}). \quad (2.9)$$

Alternatively, it is possible to derive (2.9) directly from the expectation representation of the options prices rather than going through the PIDE. Recall that v is a $\mathbb{Q}$ martingale; consequently,

$$\begin{aligned} v(t_1, \mathbf{X}(t_1)) &= \mathbb{E}_{t_1}^{\mathbb{Q}} [v(t_2, \mathbf{X}(t_2))] \\ &= \int_{-\infty}^{\infty} v(t_2, \mathbf{X}(t_1) + \mathbf{y}) \cdot f_{\mathbf{X}(t_2) - \mathbf{X}(t_1)}(\mathbf{y}) d\mathbf{y} \\ &= \int_{-\infty}^{\infty} v(t_2, \mathbf{X}(t_1) + \mathbf{y}) \cdot f_{\mathbf{X}(t_2 - t_1)}(\mathbf{y}) d\mathbf{y}. \end{aligned}$$

Here, $f_{\mathbf{X}(t)}$ denotes the probability density function (pdf) of the process $\mathbf{X}(t)$ and the third line follows from the independent-increment property of the process $\mathbf{X}(t)$. Furthermore, $\mathcal{F}[f_{\mathbf{X}(t)}](\boldsymbol{\omega}) = e^{t\Psi(-\boldsymbol{\omega})}$ and, since a convolution in real space corresponds to multiplication in Fourier space, equation (2.8) is obtained.

2.4 The Method

Armed with the Fourier transform-based solution (2.9), the numerical algorithm is straightforward. For path-independent options the price is obtained in one step by directly applying equation (2.9) similar in spirit to Carr and Madan (1999). For path-dependent options a time-stepping algorithm is used to apply boundary conditions or impose exercise constraints.

Consider a partition of the truncated stock price domain $\Omega = [\mathbf{x}^{\min}, \mathbf{x}^{\max}]$ into a finite mesh of points $\{\mathbf{x}_{\mathbf{n}} | \mathbf{n} \in [0, \dots, N-1]^d\}$, where $\mathbf{x}_{\mathbf{n}} = \mathbf{x}^{\min} + \mathbf{n}\Delta\mathbf{x}$ and $\Delta\mathbf{x} = (\mathbf{x}^{\max} - \mathbf{x}^{\min})/(N-1)$. Recall that $\mathbf{x} = \ln(\mathbf{S}/\mathbf{S}(0))$; alternatively, if pricing around the strike price is required, the scaling $\mathbf{x} = \ln(\mathbf{S}/\mathbf{K})$ can be chosen. Analogously, consider a partitioning of the time and frequency domain $\widehat{\Omega} = [0, \boldsymbol{\omega}^{\max}]$ into a finite mesh of points $\{\boldsymbol{\omega}_{\mathbf{n}} | \mathbf{n} \in [0, \dots, N/2]^d\}$, where $\boldsymbol{\omega}_{\mathbf{n}} = \mathbf{n}\Delta\boldsymbol{\omega}$ and $\Delta\boldsymbol{\omega} = 2\boldsymbol{\omega}^{\max}/N$. $\boldsymbol{\omega}^{\max}$ is chosen to be the Nyquist critical frequency, such that $\boldsymbol{\omega}^{\max} = \frac{1}{2\Delta\mathbf{x}}$. Note that $v(t, \mathbf{x})$ is a real-valued function and thus $\mathcal{F}[v](t, -\boldsymbol{\omega}) = \overline{\mathcal{F}[v](t, \boldsymbol{\omega})}$. The Fourier transform for negative frequencies is not required nor computed and therefore the frequency grid has half as many points as the spatial grid. Also, Let $t = t_0, t_1, \dots, t_M = T$, $\Delta t_m = t_m - t_{m-1}$ be a discretization of the time domain into M intervals.

Optimal grid selection is a non-trivial task and only heuristic guidelines are given below. Completion of a thorough study of grid selection and error analysis for FST framework-based methods is left for future work. As previously noted, the relationship between the space and frequency grids, as implied by the Nyquist critical frequency, is given by $\omega^{\max} \cdot (\mathbf{x}^{\max} - \mathbf{x}^{\min}) = (N - 1)/2$. Also, an appropriate transformation into log variables can be chosen so that pricing is done in the neighborhood of $\mathbf{x} = 0$. Then it is natural to choose $\mathbf{x}^{\min} = -\mathbf{x}^{\max}$, so that the area of interest is in the center of the grid, and thus $\omega^{\max} \cdot \mathbf{x}^{\max} = (N - 1)/4$. There are several competing factors at play: the real space boundary $\{-\mathbf{x}^{\max}, \mathbf{x}^{\max}\}$ should be chosen large enough to capture the overall behavior of the option value function, yet small enough to maintain the accuracy of the computed option price in the range of interest. Similarly for the frequency space, $\omega^{\max}$ should be chosen large enough to capture the high frequencies of the characteristic function, yet not too large, since having large $\Delta\omega$ would cause inaccuracies in the general shape of the stock price process and thus lead to inaccuracies in option values. Numerical experiments in this chapter suggest that $\mathbf{x}^{\max} \in [4, 7]$ works well for diffusion models with low volatility and short maturity, while $\mathbf{x}^{\max} \in [5, 10]$ is preferable for models with a large volatility term or a dominant jump component.

The one-dimensional case is considered in greater detail below. Let $\langle \mathbf{v}_m \rangle_n \triangleq v(t_m, x_n)$ represent $v(t, x)$ at the node points of the partition of Ω at time t_m , and let $\langle \hat{\mathbf{v}}_m \rangle_n \triangleq \hat{v}(t_m, \omega_n)$ represent $\mathcal{F}[v](t, \omega)$ at the node points of the partition $\hat{\Omega}$ at time t_m . The frequency domain prices are obtained from the spatial domain prices by approximating the CFT:

$$\begin{aligned} \langle \hat{\mathbf{v}}_m \rangle_n &\approx \sum_{k=0}^{N-1} v(t_m, x_k) e^{-i\omega_n x_k} \Delta x \\ &= \alpha_n \sum_{k=0}^{N-1} \langle \mathbf{v}_m \rangle_k e^{-ink/N} \\ &= \alpha_n \langle \text{FFT}[\mathbf{v}_m] \rangle_n. \end{aligned} \tag{2.10}$$

Here, $\alpha_n = e^{-i\omega_n x^{\min}} \Delta x$ and $\langle \text{FFT}[\mathbf{v}_m] \rangle_n$ denotes the n -th component of the discrete Fourier transform (DFT) of the vector $\mathbf{v}_m$, which is computed efficiently using the FFT algorithm. Similarly, the spatial domain prices can be computed from frequency domain prices via a discrete inverse transform:

$$\langle \mathbf{v}_m \rangle_n = \langle \text{FFT}^{-1} [\alpha^{-1} \cdot \hat{\mathbf{v}}_m] \rangle_n, \tag{2.11}$$

where $\alpha = [\alpha_0, \alpha_1, \dots, \alpha_{N-1}]$.

Combining the connections between frequency and spatial domains in equations (2.10) and

(2.11) with the transformed PIDE (2.9), a step backwards in time is computed by

$$\begin{aligned}
\mathbf{v}_{m-1} &= \text{FFT}^{-1} [\alpha^{-1} \cdot \hat{\mathbf{v}}_{m-1}] \\
&= \text{FFT}^{-1} [\alpha^{-1} \cdot \hat{\mathbf{v}}_m \cdot e^{\Psi(\cdot)\Delta t_m}] \\
&= \text{FFT}^{-1} [\alpha^{-1} \cdot \alpha \cdot \text{FFT}[\mathbf{v}_m] \cdot e^{\Psi(\cdot)\Delta t_m}] \\
&= \text{FFT}^{-1} [\text{FFT}[\mathbf{v}_m] \cdot e^{\Psi(\cdot)\Delta t_m}].
\end{aligned} \tag{2.12}$$

Notice that the coefficient α , which embeds information about the spatial boundary, cancels in the above equation and can be omitted during the numerical computation. An algorithm, analogous to the one-dimensional FST method in equation (2.12), can be obtained in a general multi-dimensional setting, where a step backwards in time is computed by

$$\mathbf{v}_{m-1} = \text{FFT}^{-1} [\text{FFT}[\mathbf{v}_m] \cdot e^{\Psi(\cdot)\Delta t_m}], \tag{2.13}$$

where $\text{FFT}[\cdot]$ is the multi-dimensional FFT transform and $\mathbf{v}_m$ is the d -dimensional array of option values at time t_m .

2.5 Numerical Properties

2.5.1 Stability

A potential source of instability of the FST method is that certain payoffs contain singularities in their Fourier transforms along the real axis. A simple shifting of $\omega \rightarrow \omega + i\epsilon$ avoids this problem, resulting in a slight modification of the time-stepping algorithm: $\mathbf{v}_{m-1} = \text{FFT}^{-1} [\text{FFT}[\tilde{\mathbf{v}}_m] \cdot e^{\tilde{\Psi}(\cdot)\Delta t_m}]$, where $\langle \tilde{\mathbf{v}}_m \rangle_n = e^{\epsilon x_n} \langle \mathbf{v}_m \rangle_n$ and $\tilde{\Psi}(\omega) = \Psi(\omega + i\epsilon)$. In the FST method, by truncating the payoff at very large/small spot prices, the singularity is pushed off the real axis. In all cases considered in this thesis, there is no need to treat the singularity specifically. Similarly, the CONV method of Lord, Fang, Bervoets, and Oosterlee (2008) is shown to be stable for a wide range of ϵ . Ultimately, $\epsilon = 0$ is used for their numerical experiments as it provides the best results.

Another typical source of instability in numerical methods is the choice of time step length. The FST method, by shifting to frequency space to perform the time step, has no time step restriction. Since the ODE (2.7) is solved exactly via equation (2.8), there is no error due to time discretization and the FST method (2.13) is valid for a time step Δt of any length. Furthermore, the characteristic exponent is available in closed form, hence the continuous solution (2.8) to the pricing PIDE and the FST method (2.13) have the same stability properties. Essentially, in the discrete case, the solution is sampled only at a finite number of frequencies, a subset of the solution domain in the continuous case. Since the characteristic function of any random

variable is bounded in absolute value by 1 for all ω , so is $e^{\Psi\Delta t}$. Therefore, the continuous and the discrete solutions are guaranteed to damp out and not to blow up, regardless of the step length.

Alternatively, finite difference schemes approximations, such as the forward Euler method

$$\hat{\mathbf{v}}_{m-1} = \hat{\mathbf{v}}_m \cdot (1 + \Psi(\cdot)\Delta t_m) \quad (2.14)$$

or the backward Euler method

$$\hat{\mathbf{v}}_{m-1} = \hat{\mathbf{v}}_m \cdot \frac{1}{1 - \Psi(\cdot)\Delta t_m} \quad (2.15)$$

can be used. However, such approximations are only first-order accurate and the forward Euler method has a stability restriction on the length of the time step. As such, the finite difference approximations do not offer any advantage over the explicit solution of the ODE and are not used in conjunction with the FST method.

In summary, the FST method (2.13) is valid for time steps of any length, has no error due to time discretization and the solution is guaranteed to not blow up.

2.5.2 Convergence

It is also of great interest to establish the convergence properties of the FST method. Here the estimation approach of d'Halluin, Forsyth, and Vetzal (2005) is extended to estimate the order as a function of number of space and time points independently. First, assume that the difference between the true option price v and its discrete approximation on a grid with N space points and M time points $\mathbf{v}^{[N, M]}$ is of the polynomial form

$$e^{[N, M]} \triangleq |\mathbf{v}^{[N, M]} - v| = c_n N^{-p_n} + c_m M^{-p_m},$$

where p_n and p_m are the space and time convergence rates, respectively, while c_n and c_m are convergence constants. Since the algorithm does not require time-stepping to value European options, the equation above can be simplified to depend only on N ,

$$e^{[N, \cdot]} \triangleq |\mathbf{v}^{[N, \cdot]} - v| = c_n N^{-p_n}.$$

The estimates of the option price $\mathbf{v}^{[N, \cdot]}$ on successively finer grids in space are used to establish the rate of convergence without requiring the true option value:

$$p_n = \log_2 \frac{|\mathbf{v}^{[N, \cdot]} - \mathbf{v}^{[2N, \cdot]}|}{|\mathbf{v}^{[2N, \cdot]} - \mathbf{v}^{[4N, \cdot]}|}. \quad (2.16)$$

In the numerical results reported throughout the thesis, the absolute changes in the numerator and the denominator are given in the table under the column ‘Change’, while the estimated

rate of convergence is given under the column ‘log₂Ratio’. Note that for multi-asset options, N refers to the number of points in each dimension (i.e., for a two-asset option there is a total of N^2 points on the grid) and the order of convergence is computed by doubling the number of points in each dimension.

For European options under various stock price models the results presented in the following section suggest that the FST method has second-order convergence in the space variable. It is also essential to verify that the FST method produces precise results and retains the second-order convergence for all spot prices, even far away from the center of the grid $\mathbf{x} = 0$. This is done by considering the method’s log errors $\log_{10} e^{[N,1]}$ for various N . For the FST method to retain its convergence order across a range of spot prices, the difference between the log errors remains constant and equal to the convergence order of the method (scaled by a constant):

$$\log_{10} e^{[N,\cdot]} - \log_{10} e^{[2N,\cdot]} = p_n \cdot \log_{10} 2.$$

For path-dependent options it is also necessary to establish convergence properties of the algorithm in the time variable. By holding N constant, the error becomes dependent only on M . Assuming

$$e^{[\cdot,M]} \triangleq |\mathbf{v}^{[\cdot,M]} - v| = c_m M^{-p_m},$$

the convergence order in the time variable p_m is then estimated by

$$p_m = \log_2 \frac{|\mathbf{v}^{[\cdot,M]} - \mathbf{v}^{[\cdot,2M]}|}{|\mathbf{v}^{[\cdot,2M]} - \mathbf{v}^{[\cdot,4M]}|}. \quad (2.17)$$

Time convergence results presented in Appendix C.2 overwhelmingly suggest that the straightforward application of the FST algorithm is order 1 in the time variable for path-dependent options.

A number of methods can be applied to improve convergence of the FST method both in space and in time. Richardson extrapolation can readily be used to increase the order 2 space convergence of the FST method to order 3. Given a sequence of option value estimates $\mathbf{v}^{[N,\cdot]}, \mathbf{v}^{[2N,\cdot]}, \mathbf{v}^{[4N,\cdot]}, \mathbf{v}^{[8N,\cdot]}, \dots$ that converges in space at the rate of p_n , the Richardson extrapolates $\mathbf{v}_R^{[2N,\cdot]}, \mathbf{v}_R^{[4N,\cdot]}, \mathbf{v}_R^{[8N,\cdot]}, \dots$ converge in space at the rate of at least $p_n + 1$, where

$$\mathbf{v}_R^{[2kN,\cdot]} = \frac{2^{p_n} \mathbf{v}^{[2kN,\cdot]} - \mathbf{v}^{[kN,\cdot]}}{2^{p_n} - 1}, \quad k = 1, 2, 4, 8, \dots \quad (2.18)$$

Similarly, Richardson extrapolates $\mathbf{v}_R^{[\cdot,2M]}, \mathbf{v}_R^{[\cdot,4M]}, \mathbf{v}_R^{[\cdot,8M]}, \dots$ converge in time at the rate of at least $p_m + 1$, where

$$\mathbf{v}_R^{[\cdot,2kM]} = \frac{2^{p_m} \mathbf{v}^{[\cdot,2kM]} - \mathbf{v}^{[\cdot,kM]}}{2^{p_m} - 1}, \quad k = 1, 2, 4, 8, \dots \quad (2.19)$$

Also, in Section 2.6.3, a penalty method for American options is developed to improve the convergence in time to order 2. Similarly, for barrier options, a method of images, discussed in Section 2.6.4, is used to improve the convergence rate of the FST method.

2.5.3 Precision

A CFT represents a general function by a sum of complex exponentials of various frequencies, which are periodic functions. Effectively, by truncating the solution domain to $[\mathbf{x}^{\min}, \mathbf{x}^{\max}]$ and working in frequency space, the solution outside of the domain is replaced by the periodic extension of the solution inside the domain. By choosing the domain large enough, the contribution of this error to the solution in the region of interest would be extremely small. However, the solution near the boundary (e.g. $\mathbf{S} = 0$) will be affected by the wrap-around effect and require special consideration. Various methods exist to tackle this issue, however, they are not specifically addressed in this thesis.

The pricing results are verified by comparing them to the prices found in the literature and/or computed via alternative methods, such as (semi) closed-form formulas and Monte-Carlo method. If a closed-form formula is not available, a very reliable estimate for European options can be found by evaluating the integral form in Carr and Madan (1999) using an adaptive Fourier quadrature method.

2.6 Applications to Option Pricing

In this section the precision and convergence properties of the FST method are examined by applying the method to pricing of options, such as single-asset European, American and barrier options, and multi-asset European and American spread and catastrophe options. In addition to the results presented, the FST method computes option prices for a range of spot prices, which is a significant advantage for pricing path-dependent options, such as the American options presented below, and other applications, such as computation of the Greeks.

The option and stock price models are specified below each table. Quoted reference price is the most precise result reported in the quoted paper. While it would be beneficial to extrapolate the results given in reference papers to obtain the theoretical limit of convergence, such analysis is hard to carry out in practice due to the limited precision of results typically reported. Where possible, reference results computed via alternative methods are provided.

The results reported throughout this work were obtained using code written in C++ in conjunction with the FFTW library for evaluation on FFTs. The code was run on a workstation with Intel Core 2 Duo E7200 2.53GHz CPU and 4GB of RAM. Option and model parameters can be found in Appendix B.

2.6.1 European Options

European options can only be exercised at maturity and, thus, their payoff depends only on the terminal value of the underlying asset at maturity. European options can be valued in a single time step, since (2.12) is a valid approximation for any Δt , with only truncation of spatial and frequency domains contributing to the error. In this case, given a payoff function φ , set $M = 1$, $\mathbf{v}_1 = \varphi(\mathbf{S}(0)e^{\mathbf{x}})$ and apply (2.12) to obtain $\mathbf{v}_0$. A variety of options can be priced in this manner by supplanting an appropriate payoff function $\varphi(\mathbf{S})$. For instance, in the single-asset case, the European and American call options have payoff $\varphi(S) = \max(S - K, 0)$, while the put options have payoff $\varphi(S) = \max(K - S, 0)$. Since the method does not require the analytic transform of option payoff, exotic options, such as digital call $\varphi(S) = \mathbb{1}_{\{S \geq K\}}$ and put $\varphi(S) = \mathbb{1}_{\{S < K\}}$ and asset-or-nothing call $\varphi(S) = S \cdot \mathbb{1}_{\{S \geq K\}}$ and put $\varphi(S) = S \cdot \mathbb{1}_{\{S < K\}}$ options can also be priced.

This approach is similar to Carr and Madan (1999), however, the analytic expression of the Fourier transformed option payoff is not required — clearly a great advantage for non-standard payoffs. Moreover, the FST approach is computationally more efficient than spatial PIDE solution-based methods since it does not require stepping in time.

Example 1: Single-asset European options

Table 2.2 presents pricing results for the European option *EUR-A* under the Merton jump-diffusion model *MJD-A*. The results (and further results in Table 6.2 and Tables C.1 - C.4, C.15 in Appendix C) overwhelmingly suggest that the FST method is precise and has second-order convergence in the space variable. Richardson extrapolation is used to improve the convergence in the space variable to 3.

N	No Extrapolation			Richardson Extrapolation			Time (msec.)
	Value	Change	$\log_2$ Ratio	Value	Change	$\log_2$ Ratio	
2048	18.00329705						0.994
4096	18.00354600	0.0002490		18.003628985			1.424
8192	18.00360820	0.0000622	2.0008	18.003628939	4.62×10^{-8}		2.811
16384	18.00362375	0.0000155	2.0004	18.003628933	5.84×10^{-9}	2.9840	5.778
32768	18.00362764	0.0000039	2.0002	18.003628932	7.34×10^{-10}	2.9921	11.572

Table 2.2: Pricing results for the European option *EUR-A* under the Merton jump-diffusion model *MJD-A*. Reference price 18.0034 and parameters from Andersen and Andreasen (2000). The reference price 18.00362936 is computed using a semi closed-form formula. The order of convergence is 3 in space with Richardson extrapolation (and 2 in space without extrapolation).

Figure 2.1: Errors for pricing the European option *EUR-F* under the Merton jump-diffusion model *MJD-B* (left) and the barrier option *CBR-A* under the Black-Scholes-Merton model *BSM-B* (right). The average rate of convergence (across all spot prices) is 2 in space for the former scenario and is 2 in space and 1 in time for the latter scenario.

As discussed in Section 2.5.2, for the method to retain the order of convergence across a range of spot prices, the difference between the log errors $\log_{10} e^{[N, \cdot]}$ on successively finer grids must remain constant and equal to the order of convergence (scaled by a constant). Figure 2.1 depicts the log errors of the FST method for pricing the European option *EUR-F* under the Merton jump-diffusion model *MJD-B*. For European options under the jump-diffusion models, a semi closed-form solution is available and $e^{[N, 1]}$ can be readily computed. The second-order convergence is retained over the range of spot prices as the difference between log errors in each successive refinement is $2 \cdot \log_{10} 2$. The curvature in the log errors is due to the change in convergence constant (or the absolute error) across the various spot prices.

Example 2: Two-asset spread options

An interesting class of multi-asset options are spread options — the option to exchange β_1 -units of one asset for β_2 -units of another asset. These options can be viewed as options on the difference (or spread) of two stock prices, hence the payoff at maturity is $\varphi(S_1, S_2) = \max(\beta_2 S_2 - \beta_1 S_1 - K, 0)$ for a call option and $\varphi(S_1, S_2) = \max(K - (\beta_2 S_2 - \beta_1 S_1), 0)$ for a put option.

Spread options do not admit a closed-form solution even for the BSM model if $K \neq 0$ ². Dempster and Hong (2000) present an FFT-based approach for valuation of spread options. Their approach involves breaking the region in which the option is in-the-money into a series of rectangular approximations. Unfortunately, they apply the method only to a pure diffusion

²See Carmona and Durrleman (2003) for a detailed discussion of spread options and various approximations

N	Value	Change	$\log_2 \text{Ratio}$	Time (sec.)
512 ²	13.71176205			0.212
1024 ²	13.71335323	1.59×10^{-3}		0.839
2048 ²	13.71353098	1.78×10^{-4}	3.1622	3.015
4096 ²	13.71352849	2.48×10^{-6}	6.1623	12.618
8192 ²	13.71352847	2.36×10^{-8}	6.7159	49.741

Table 2.3: Pricing results for the European spread option *ESPD* under the 2D Merton jump-diffusion model *MJD-E*. The reference price 13.714948858 is computed using Kirk's approximation formula. The order of convergence is at least 2 in space.

model with stochastic volatility and it seems difficult to extend this method to the Bermudan or barrier cases.

For the numerical experiments, a joint jump-diffusion with uncoupled idiosyncratic Merton-like jumps is assumed and results are compared with the Kirk (1995) approximation and its extension for jump-diffusions found in Carmona and Durrleman (2003). In this case, the Lévy density factors with $\nu_k(dy) = (\lambda_k / \sqrt{2\pi\tilde{\sigma}_k^2}) \exp\{-(y - \tilde{\mu}_k)^2 / 2\tilde{\sigma}_k^2\} dy$, for $k = 1, 2$, and the diffusive volatilities are σ_k with correlation ρ . Hence,

$$\begin{aligned} \Psi(\omega_1, \omega_2) = & i(\gamma_1 - \frac{\sigma_1^2}{2})\omega_1 + i(\gamma_2 - \frac{\sigma_2^2}{2})\omega_2 - \frac{\sigma_1^2\omega_1^2}{2} - \rho\sigma_1\sigma_2\omega_1\omega_2 - \frac{\sigma_2^2\omega_2^2}{2} \\ & + \lambda_1 (e^{i\tilde{\mu}_1\omega_1 - \tilde{\sigma}_1^2\omega_1^2/2} - 1) + \lambda_2 (e^{i\tilde{\mu}_2\omega_2 - \tilde{\sigma}_2^2\omega_2^2/2} - 1), \end{aligned} \quad (2.20)$$

where the drifts are fixed by risk-neutrality to be $\gamma_k = r - \lambda_k (e^{\tilde{\mu}_k + \tilde{\sigma}_k^2/2} - 1)$, for $k = 1, 2$.

Table 2.3 presents pricing results for the European spread option *ESPD* under the 2D Merton jump-diffusion model *MJD-E*. The results (and further results in Table C.17) suggest that the multi-dimensional FST method is at least second-order in space.

Example 3: Catastrophe equity put options

Catastrophe options pay the holder a function of total losses and the company's equity value. As a result, it is important to jointly model losses and equity, especially since large losses can cause significant drops in share value. Cox, Fairchild, and Pedersen (2004) and Jaimungal and Wang (2006) price European catastrophe options; Lin and Wang (2009) study the perpetual version of these options using ruin theory methods.

A catastrophe equity put (CatEPut) option is a European put option that can be exercised only if the losses exceed a predetermined level and has a payoff function $\varphi(S, L) = \mathbb{1}_{\{L > L^*\}}(K - S)_+$. Figure 2.2 depicts the payoff (as a function of stock price and loss level) of the CatEPut option *ECEP*.

Figure 2.2: Payoff (left) and value (right) of the catastrophe equity put option *ECEP* under the joint stock-loss model *JSL*.

Under the joint model for catastrophe loss and stock price model described by equation (1.7), the 2-dimensional Lévy density is $\nu(dy_1 \times dy_2) = f_L(y_2) \delta(y_1 + \chi y_2) dy_1 dy_2$ resulting in the characteristic function

$$\Psi(\omega_1, \omega_2) = i\gamma\omega_1 - \frac{1}{2}\sigma^2\omega_1^2 + \lambda \left[\left[1 - \frac{iv_l}{m_l}(-\chi\omega_1 + \omega_2) \right]^{-\frac{m_l^2}{v_l}} - 1 \right], \quad (2.21)$$

with the risk-neutral drift γ chosen by setting $\Psi(-i, 0) = r$.

Figure 2.2 depicts the value (as a function of stock price and loss level) of the CatEPut option *ECEP* and the pricing results for the option under the joint stock-loss model *JSL* are given in Table 6.3. It is intuitive that the European option price is a smoothed version of the payoff function. Due to the highly discontinuous nature of the payoff of CatEPut options, the FST method has only first-order convergence when applied to their pricing.

2.6.2 American Options

American options have the same structure as European options but can be exercised at any time prior to maturity. Upon exercise, the holder of the option receives a payoff that depends on the current value of the underlying asset. American options can be priced using a finite difference method either by solving a linear complementarity problem (see Wilmott, Howison, and Dewynne (1993), Dempster and Hutton (1997), Huang and Pang (1998), and Forsyth and Vetzal (2002)), or solving a free boundary value problem (see McKean (1965), Kim (1990), and Carr, Jarrow, and Myneni (1992)). Although the free boundary formulation for American options is an active area of research and can be potentially combined with the FST method, the linear complementarity formulation is easier to implement in the context of the FST method.

Figure 2.3: Schematic representation of the Fourier Space Time-stepping method. The boundary conditions (such as optimal exercise or barrier breach) are applied in real space while the time step is performed in Fourier space.

Since the value of an American option is always greater than or equal to the terminal payoff, the idea is to continuously enforce the condition $v(t, \mathbf{x}) \geq v(T, \mathbf{x})$. Numerically, this is enforced when boundary conditions are applied, resulting in the following algorithm:

$$\mathbf{v}_{m-1}^* = \text{FFT}^{-1} \left[\text{FFT} [\mathbf{v}_m] \cdot e^{\Psi(\cdot) \Delta t_m} \right], \quad (2.22a)$$

$$\mathbf{v}_{m-1} = \max \{ \mathbf{v}_{m-1}^*, \mathbf{v}_M \}, \quad (2.22b)$$

where $\mathbf{v}_{m-1}^*$ represents the holding value of the option and the max operator is applied componentwise. There is no convenient representation of the max operator in Fourier space; consequently, it is necessary to switch between real and Fourier spaces at each time step. Schematically, the algorithm is presented in Figure 2.3.

Example 4: Single-asset American options

Table 2.4 presents pricing results for the American option *AMR-A* under the CGMY model *CGMY-A*. The results (and further results Table 6.4 and Tables C.5, C.7, C.16 in Appendix C) suggest that the FST method for pricing of American options using equation (2.22) is second-order in space and first-order in time. Section 2.6.3 discusses an approach, based on the penalty method of Forsyth and Vetzal (2002), for attaining quadratic convergence when pricing American options with the FST method.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	128	9.22185444			0.011
4096	512	9.22447187	0.0026174		0.088
8192	2048	9.22520096	0.0007291	1.8440	0.717
16384	8192	9.22538213	0.0001812	2.0087	5.990
32768	32768	9.22542569	0.0000436	2.0565	49.538

Table 2.4: Pricing results for the American option *AMR-A* under the CGMY model *CGMY-A*. Reference price 9.2254803 and parameters from Forsyth, Wan, and Wang (2007). The order of convergence is 2 in space and 1 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512^2	64	13.92167479			1.054
1024^2	256	13.92603166	0.0043569		20.317
2048^2	1024	13.92683797	0.0008063	2.4339	326.589
4096^2	4096	13.92698521	0.0001472	2.4532	6524.372

Table 2.5: Pricing results for the American spread option *ASPD* under the 2D Merton jump-diffusion model *MJD-E*. The order of convergence is 2 in space and 1 in time.

Example 5: Two-asset American spread options

Table 2.5 presents pricing results for the American spread option *ASPD* under the 2D Merton jump-diffusion model *MJD-E*. The results (and further results in Table C.18) suggest that the FST method for pricing of American spread options is second-order in space and first-order in time. The penalty method can be used to improve the convergence as in the one-dimensional case.

For multi-asset options with early-exercise features, in addition to computing the price surface, the exercise boundary can be computed. Figure 2.4 depicts the exercise boundary (as a function of time to maturity and second stock price) of the American spread option *ASPD* under the 2D BSM model *BSM-C*. For a given time to maturity τ and stock price S_2 , the plotted stock price $S_1(\tau, S_2)$ is the price below which it is optimal to exercise the option. Naturally, as time approaches maturity, the exercise boundary tends to $S_1 = S_2 - K$.

Example 6: American double-trigger stop-loss options

A double-trigger stop-loss (DTSL) option is another type of catastrophe options discussed in the previous section. A DTSL option pays the holder the amount of losses incurred in excess of the attachment level L_a , tapering off at the detachment level L_d , with the payment being

Figure 2.4: Exercise boundary of the American spread option *ASPD* under the 2D Black-Scholes-Merton model *BSM-C*. The optimal exercise region is below the plotted curve $S_1(\tau, S_2)$.

conditional on the stock price being below a given strike price. The payoff function is given by $\varphi(S, L) = \mathbb{1}_{\{S < K\}}[(L - L_a)_+ - (L - L_d)_+]$. Figure 2.5 depicts the payoff (as a function of stock price and loss level) of the American DTSL option *ADTSL*.

Since many catastrophe options are issued as Bermudan contracts, which can be exercised on a quarterly, monthly, weekly or daily basis, the FST method can be used to study the early-exercise premium. A slight modification of the American-styled options algorithm leads to an efficient pricing mechanism. Namely, the exercise policy is now chosen at each point in the $(S(t), L(t))$ plane independently. If $L(t)$ was not a separate observable, as it is in the usual jump-diffusion model case, then the exercise policy would be independent of $L(t)$. Figure 2.5 depicts the value (as a function of stock price and loss level) of the American DTSL option *ADTSL* and the pricing results under the joint stock-loss model *JSL* are given in Table 6.5. Again, the FST method has only order 1 convergence in space and order 1/2 convergence in time when applied pricing of American DTSL options.

The American price surfaces have smooth behavior when the stock price is above the trigger level of 100 while there is a distinct kink across this spot price level for losses above 5. Within this region, the American DTSL option *ADTSL* is in-the-money. The behavior along the $(S = 100, L > 5)$ reflects the discontinuous behavior of the intrinsic value of the option. These features are most easily explained and observed by investigating the optimal exercise policies. Figure 2.6 depicts the optimal exercise behavior at different points in time (exercise occurs in

Figure 2.5: Payoff (left) and value (right) of the American double-trigger stop-loss option *ADTSL* under the joint stock-loss model *JSL*.

the area above the exercise curve and to the left of the line $S = 100$). The exercise policy is governed by two competing factors. An investor is inclined to delay exercise to allow for the losses to accumulate and thus raise the option value at exercise. However, by waiting the investor is running the risk that the stock price will cross the threshold K and never venture below before maturity, thus rendering the option worthless. Since the latter risk is far greater than the former upside potential, the exercise boundary trends away from the threshold $S = K$ as maturity approaches. Furthermore, the exercise curve flattens at around $L = 38$ for all maturities when $S < 90$. This can be explained by the discrete nature of loss arrival — since the maximal payoff is achieved at $L = 40$, it is optimal to exercise at $L = 38$ rather than waiting and possibly exercising at the same loss level if no additional losses arrive. In a separate experiment, as the arrival rate of losses was increased, this upper threshold approached $L = 40$.

2.6.3 American Options with Penalty Method

As previously mentioned, the American option pricing problem can be expressed as a linear complementarity problem:

$$\begin{cases} (\partial_t + \mathcal{L}) v(t, \mathbf{x}) & \geq 0, \\ v(t, \mathbf{x}) - v(T, \mathbf{x}) & \geq 0, \end{cases} \quad (2.23)$$

where $v(T, \cdot)$ is the option payoff and at each point of solution domain either $(\partial_t + \mathcal{L}) v(t, \mathbf{x}) = 0$ or $v(t, \mathbf{x}) - v(T, \mathbf{x}) = 0$. The idea behind the penalty method is to replace problem (2.23) by

$$(\partial_t + \mathcal{L}) v(t, \mathbf{x}) + \xi P(v)(t, \mathbf{x}) = 0, \quad (2.24)$$

Figure 2.6: Exercise boundary of the American double-trigger stop-loss option *ADTSL* under the joint stock-loss model *JSL*.

where $P(v)(t, \mathbf{x}) = \max(v(T, \mathbf{x}) - v(t, \mathbf{x}), 0)$ and ξ is a penalty parameter. The Fourier transform is applied to the penalty equation to obtain an ODE with a potential-like term:

$$(\partial_t + \Psi(\omega)) \mathcal{F}[v](t, \omega) + \xi \mathcal{F}[P(v)](t, \omega) = 0. \quad (2.25)$$

This cannot be solved easily since $v(t, \omega)$ appears in a non-linear fashion. However, if the solution is viewed as a fixed point of an iteration scheme in which (2.25) is replaced by

$$(\partial_t + \Psi(\omega)) \mathcal{F}[v^{(k)}](t, \omega) + \xi \mathcal{F}[P(v^{(k-1)})](\omega) = 0, \quad (2.26)$$

then the ODE system can be solved explicitly for $v^{(k)}$, taking $v^{(k-1)}$ to be a known function (computed in the previous iteration). In this form, the transformed penalty term behaves as a source term at each iteration. The solution to the ODE (2.26) can be found by first considering solutions to the homogeneous equation (2.2) and a particular solution to the inhomogeneous equation (2.26). The solution to the homogeneous equation, $\mathcal{F}[v_H^{(k)}](t, \omega)$, is given by

$$\mathcal{F}[v_H^{(k)}](t_1, \omega) = C^{(k)}(\omega) \cdot e^{\Psi(\omega)(t_2 - t_1)}, \quad (2.27)$$

for some value $C^{(k)}(\omega)$ to be determined. A particular solution to the ODE (2.26) is given by

$$\mathcal{F}[v_P^{(k)}](t_1, \omega) = -\frac{\xi \mathcal{F}[P(v^{(k-1)})](\omega)}{\Psi(\omega)}. \quad (2.28)$$

The general solution is the sum of the homogeneous and particular solutions, $\mathcal{F}[v^{(k)}](t_1, \omega) = \mathcal{F}[v_H^{(k)}](t_1, \omega) + \mathcal{F}[v_P^{(k)}](t_1, \omega)$, subject to the initial value $\mathcal{F}[v](t_2, \omega)$. Now $C^{(k)}(\omega)$ can be

computed by letting $t_1 \rightarrow t_2$:

$$\begin{aligned} \mathcal{F}[v](t_2, \omega) &= C^{(k)}(\omega) \cdot e^{\Psi(\omega)(t_2-t_1)} - \frac{\xi \mathcal{F}[P(v^{(k-1)})](\omega)}{\Psi(\omega)} \\ \Rightarrow C^{(k)}(\omega) &= \mathcal{F}[v](t_2, \omega) + \frac{\xi \mathcal{F}[P(v^{(k-1)})](\omega)}{\Psi(\omega)}. \end{aligned} \quad (2.29)$$

Thus, the solution to the ODE (2.26) is

$$\begin{aligned} \mathcal{F}[v^{(k)}](t_1, \omega) &= \mathcal{F}[v](t_2, \omega) \cdot e^{\Psi(\omega)(t_2-t_1)} + \\ &\quad \xi \mathcal{F}[P(v^{(k-1)})](\omega) \cdot \left(\frac{e^{\Psi(\omega)(t_2-t_1)} - 1}{\Psi(\omega)} \right). \end{aligned} \quad (2.30)$$

Taking the inverse Fourier transform, the price is the fixed point of the iteration

$$\begin{aligned} v^{(k)}(t_1, \mathbf{x}) &= \mathcal{F}^{-1} \left[\mathcal{F}[v](t_2, \omega) \cdot e^{\Psi(\omega)(t_2-t_1)} \right](\mathbf{x}) + \\ &\quad \xi \mathcal{F}^{-1} \left[\mathcal{F}[P(v^{(k-1)})](\omega) \cdot \left(\frac{e^{\Psi(\omega)(t_2-t_1)} - 1}{\Psi(\omega)} \right) \right](\mathbf{x}). \end{aligned} \quad (2.31)$$

The iterative FST method can therefore be expressed as

$$\mathbf{v}_{m-1}^{(k)} = \mathbf{v}_{m-1} + \xi \text{FFT}^{-1} \left[\text{FFT} [P(\mathbf{v}_{m-1}^{(k-1)})] \cdot \left(\frac{e^{\Psi(\cdot)\Delta t_m} - 1}{\Psi(\cdot)} \right) \right], \quad (2.32)$$

where $\mathbf{v}_{m-1}^{(0)} = \mathbf{v}_{m-1}$ and $\mathbf{v}_{m-1}$ is computed using the usual time step in equation (2.12), which does not incorporate a penalty or optimal exercise, i.e., it is the holding value of the option. To avoid introducing bias into the explicit iteration, ξ is chosen so that $\xi \cdot \left(\frac{e^{\Psi(\cdot)\Delta t} - 1}{\Psi(\cdot)} \right) \rightarrow 1$ as $\Delta t \rightarrow 0$. In the small Δt limit, the scheme then corresponds to correcting the holding value with the explicit penalty. From the Taylor series of the exponential function one can obtain $\xi = 1/\Delta t$.

The numerical results indicate that only a single iteration of the penalty method is required to attain second-order for American options. No theoretical basis exists for this behavior and the convergence properties of the American penalty method warrant further research.

Example 4 continued: Single-asset American options

Table 2.6 presents pricing results for the American option *AMR-A* under the CGMY model *CGMY-A* using the FST penalty method. The results (and further results in Tables C.6, C.8 in Appendix C.1) suggest that the FST penalty method for pricing American options has second-order convergence in space and time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	128	9.22478538			0.027
4096	256	9.22523484	0.0004495		0.109
8192	512	9.22538196	0.0001471	1.6114	0.451
16384	1024	9.22542478	0.0000428	1.7808	1.869
32768	2048	9.22543516	0.0000104	2.0444	8.195

Table 2.6: Pricing results for the American option *AMR-A* under the CGMY model *CGMY-A* using the FST penalty method. Reference price 9.2254803 and parameters from Forsyth, Wan, and Wang (2007). The order of convergence is 2 in space and 2 in time.

2.6.4 Barrier Options

Barrier options are options for which the payoff depends on whether the underlying asset reaches a barrier level B . Many different types of barrier options exist. The two main classes are *knock-in* options (they come into existence if the barrier is reached) and *knock-out* options (they cease to exist if the barrier is reached). Depending on whether the barrier is below or above the initial spot price, barrier options can be classified as *down-and-in*, *down-and-out*, *up-and-in* or *up-and-out* options. Also, barrier options may pay a rebate when the barrier is reached.

Transform-based methods have been widely utilized to price barrier options. Kou and Petrella (2004) develop a numerical algorithm based on Laplace transforms which can be applied to all Lévy models. However, this method is computationally expensive and limited to single-barrier options. Fast Gauss transforms have been utilized by Broadie and Yamamoto (2005) to price discretely-monitored barrier options. Their method is very efficient for the Merton jump-diffusion model, however, it cannot be applied to the general class of Lévy models. Recently, Feng and Linetsky (2008) proposed an exponentially convergent method based on the fast Hilbert transform which can be applied to Lévy processes.

The FST method for barrier options is similar to the FST method for American options: both involve enforcement of constraints at each time step. Here, the up-and-out barrier option case is discussed; however, the results can be extended to other barrier option styles. In spatial coordinates, the barrier boundary condition forces

$$v(t, \mathbf{x}) = R \quad \text{for } \mathbf{x} \geq \mathbf{B}, t \leq T \quad (2.33)$$

where $\mathbf{B}$ is the knock-out barrier level (in log-space) and R is the rebate paid in the case of knock-out. In terms of the time-stepping algorithm

$$\mathbf{v}_{m-1} = \text{FFT}^{-1} \left[\text{FFT} [\mathbf{v}_m] \cdot e^{\Psi(\cdot)\Delta t_m} \right] \cdot \mathbb{1}_{\{\mathbf{x} < \mathbf{B}\}} + R \cdot \mathbb{1}_{\{\mathbf{x} \geq \mathbf{B}\}}. \quad (2.34)$$

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	128	0.25329897			0.017
4096	512	0.25396930	0.0006703		0.115
8192	2048	0.25414052	0.0001712	1.9690	0.653
16384	8192	0.25418296	0.0000424	2.0123	6.129
32768	32768	0.25419312	0.0000102	2.0638	46.227

Table 2.7: Pricing results for the barrier option *CBR-A* under the Black-Scholes-Merton model *BSM-A*. The reference price 0.2541963 is computed using a closed-form formula. The order of convergence is 2 in space and 1 in time.

For discretely-monitored barrier options, the time points are chosen to lie precisely on monitoring dates and equation (2.34) is applied. Since the FST method is exact between monitoring dates, the number of time steps required for the FST method is exactly the number of monitoring dates. Moreover, since only a single time step of the algorithm is required between monitoring dates, the FST method is significantly more efficient than finite difference schemes, which normally require several time steps between the monitoring dates. The numerical properties of discretely monitored barrier options are studied in the following chapter in the context of mean-reverting models.

For continuously-monitored barrier options, numerical experiments show that direct application of (2.34) results in slow convergence. To improve the convergence, constraint (2.33) is enforced numerically via the method of images (see e.g., Buchen (1996)) by truncating $v(t, \mathbf{x})$ at $\mathbf{x} = B$ and extending it to an odd function, i.e., setting $v(t, \mathbf{B} + \mathbf{y}) = 2R - v(t, \mathbf{B} - \mathbf{y})$ for $\mathbf{y} > 0$ and $v(t, \mathbf{B}) = R$ at each time step. For geometric Brownian motion this procedure does not introduce any bias into the solution of the equation for $\mathbf{x} \leq \mathbf{B}$ and improves the convergence of the FST algorithm. However, for jump-diffusion models the method of images (but not the original approach) introduces a bias, due to the augmentation of the option value function by reflection, and alternative methods, such as Richardson extrapolation, should be used to improve convergence.

Example 7: Barrier options

Table 2.7 presents pricing results for the barrier option *CBR-A* under the BSM model *BSM-A*. The results suggest that the FST method for pricing of barrier options using equation (2.34) with the method of images has second-order convergence in space and first-order convergence in time. Table C.9 in Appendix C.1 presents pricing results for the barrier option *CBR-B* under the Merton jump-diffusion model *MJD-D*. Here, the Richardson extrapolation, instead of the

method of images, is applied to improve the order of convergence in time to 1.

Figure 2.1 depicts the log errors of the FST method for pricing the barrier option *CBR-A* under the BSM model *BSM-B*. For barrier options under the BSM model, a closed-form solution is available and $e^{[N,1]}$ can be readily computed. The second-order convergence in space and first-order convergence in time is retained over the range of spot prices as the difference between log errors in each successive refinement is $2 \cdot \log_{10} 2$. Note that the number of time points is quadrupled in each refinement. The curvature in the log errors is due to the change in convergence constant (or the absolute error) across the various spot prices. Also note that there is no loss in neither the absolute error nor the convergence order near the boundary.

2.7 Applications to Hedging with Greeks

This section develops the greekFST method for computation of option Greeks (sensitivities of option value to changes in the price of the underlying spot price model parameters). The Greeks play a paramount role in risk management, where the goal is to minimize the exposure to risk, or hedge away the risk, of holding options. A single option or a whole portfolio of derivatives can be immunized from gains or losses by taking offsetting position in the underlying, with the size of the offsetting position determined by the value of the various Greeks.

While a number of methods exist for numerical pricing of options, the literature on numerical hedging is very limited. The main idea behind the greekFST method is to derive the PIDE satisfied by the Greeks and solve the PIDE by transforming it into an ODE in Fourier space. The solution relates the Fourier transform of option Greeks to the Fourier transform of option values via multiplication by a constant term that can be computed analytically. Furthermore, for several Greeks, such as Delta and Gamma, the solution can be found without deriving a new PIDE. This section presents an efficient numerical scheme for computing Greeks for European options, however, the analytic solution allows (but not discussed here) for straightforward extensions for path-dependent options as well.

Delta

Delta measures the sensitivity, or the rate of change, of the option price v with respect to the change in the price of the underlying S . The Fourier transform of Delta can be computed from the Fourier transform of option values via scaling:

$$\begin{aligned} \partial_{\mathbf{S}_k} v(t, \mathbf{x}) &= \partial_{\mathbf{x}_k} v(t, \mathbf{x}) e^{-\mathbf{x}_k} \\ &= \mathcal{F}^{-1} [i\boldsymbol{\omega}_k \cdot \mathcal{F}[v](t, \boldsymbol{\omega})](\mathbf{x}) / (\mathbf{S}_k(0)e^{\mathbf{x}_k}) . \end{aligned} \quad (2.35)$$

Note that k denotes the dimension of the underlying asset with respect to which the derivative is computed (and hence may be omitted for single-asset options).

Applying the theory developed in Chapter 2.4, the greekFST method for computing the Deltas at time t_{m-1} , given option values at time t_m , can be expressed in discrete space as

$$\Delta_{k,m-1} = \text{FFT}^{-1} [i\omega_k \cdot \hat{\mathbf{v}}_{m-1}] / (\mathbf{S}_k(0)e^{\mathbf{x}_k}) , \quad (2.36)$$

where $\hat{\mathbf{v}}_{m-1} = \text{FFT} [\mathbf{v}_m] \cdot e^{\Psi(\cdot)\Delta t_m}$

Gamma

Gamma measures the rate of change of Delta with respect to the change in the price of the underlying. Alternatively, Gamma is the second derivative of the option price with respect to the spot price. In multi-asset options, mixed Gamma is the mixed second derivative of option price with respect to two distinct spot prices. Similarly to Delta, the Fourier transform of Gamma can be computed from the Fourier transform of option values via scaling:

$$\begin{aligned} \partial_{\mathbf{S}_k}^2 v(t, \mathbf{x}) &= \left(-\partial_{\mathbf{x}_k} + \partial_{\mathbf{x}_k^2} \right) v(t, \mathbf{x}) / (\mathbf{S}_k(0)e^{\mathbf{x}_k})^2 \\ &= \mathcal{F}^{-1} [(-i\omega_k - \omega_k^2) \cdot \mathcal{F}[v](t, \omega)](\mathbf{x}) / (\mathbf{S}_k(0)e^{\mathbf{x}_k})^2 , \end{aligned} \quad (2.37)$$

and the greekFST method for computing Gammas is then given by

$$\Gamma_{k,m-1} = \text{FFT}^{-1} [-(i\omega_k + \omega_k^2) \cdot \hat{\mathbf{v}}_{m-1}] / (\mathbf{S}_k(0)e^{\mathbf{x}_k})^2 . \quad (2.38)$$

For the mixed Gamma, a similar scaling result is obtained:

$$\begin{aligned} \partial_{\mathbf{S}_k \mathbf{S}_l}^2 v(t, \mathbf{x}) &= \partial_{\mathbf{x}_k \mathbf{x}_l} v(t, \mathbf{x}) / (\mathbf{S}_k(0)e^{\mathbf{x}_k} \cdot \mathbf{S}_l(0)e^{\mathbf{x}_l}) \\ &= \mathcal{F}^{-1} [(i\omega_k)(i\omega_l) \cdot \mathcal{F}[v](t, \omega)](\mathbf{x}) / (\mathbf{S}_k(0)e^{\mathbf{x}_k} \cdot \mathbf{S}_l(0)e^{\mathbf{x}_l}) , \end{aligned} \quad (2.39)$$

with the greekFST method for mixed Gamma is easily shown to be

$$\Gamma_{k,l,m-1} = \text{FFT}^{-1} [-\omega_k \omega_l \cdot \hat{\mathbf{v}}_{m-1}] / (\mathbf{S}_k(0)e^{\mathbf{x}_k} \cdot \mathbf{S}_l(0)e^{\mathbf{x}_l}) . \quad (2.40)$$

Vega

Vega measures the sensitivity of the option value with respect to the change in the volatility of the underlying asset σ . Applying a derivative with respect to σ_k to the pricing PIDE, a PIDE satisfied by Vega is obtained:

$$\partial_{\sigma_k} \{(\partial_t + \mathcal{L}) v(t, \mathbf{x})\} = (\partial_t + \mathcal{L}) \partial_{\sigma_k} v(t, \mathbf{x}) + \mathcal{H}_{\sigma_k} v(t, \mathbf{x}) = 0 , \quad (2.41)$$

where $\mathcal{H}_{\sigma_k} = (\partial_{\sigma_k} \gamma)' \partial_{\mathbf{x}} + \partial_{\mathbf{x}}' (\partial_{\sigma_k} \Sigma) \partial_{\mathbf{x}}$. Applying the Fourier transform to the Vega PIDE yields an ODE for $\mathcal{F}[\partial_{\sigma_k} v](t, \omega)$ with a source term:

$$(\partial_t + \mathcal{L}) \mathcal{F}[\partial_{\sigma_k} v](t, \omega) + \mathcal{F}[\mathcal{H}_{\sigma_k}](\omega) \cdot \mathcal{F}[v](t, \omega) = \mathbf{0}, \quad (2.42)$$

which can be solved explicitly:

$$\partial_{\sigma_k} v(t, \mathbf{x}) = (T - t) \mathcal{F}^{-1} [\mathcal{F}[\mathcal{H}_{\sigma_k}](\omega) \cdot \mathcal{F}[v](t, \omega)](\mathbf{x}). \quad (2.43)$$

The greekFST method for computing Vegas is then given by

$$\nabla_{k,m-1} = \Delta t_m \text{FFT}^{-1} [\mathcal{F}[\mathcal{H}_{\sigma_k}](\omega) \cdot \hat{\mathbf{v}}_{m-1}], \quad (2.44)$$

where the $\mathcal{F}[\mathcal{H}_{\sigma_k}](\omega)$ term can be computed analytically. For example, in the case of the BSM model, $\mathcal{F}[\mathcal{H}_{\sigma_k}](\omega) = -(i\omega + \omega^2)\sigma$.

Volga

Volga (or Vomma) measures the rate of change of Vega with respect to the change in the volatility σ . Alternatively, Volga is the second derivative of the option value with respect to volatility. In multi-asset options, mixed Volga is the mixed second derivative of option price with respect to two distinct volatilities.

For the case of mixed Volga, assuming $k \neq l$ and applying the mixed derivative with respect to σ_k and σ_l to the pricing PIDE, a PIDE satisfied by the mixed Volga is obtained:

$$\partial_{\sigma_k \sigma_l}^2 \{(\partial_t + \mathcal{L}) v(t, \mathbf{x})\} = (\partial_t + \mathcal{L}) \partial_{\sigma_k}^2 v(t, \mathbf{x}) + \mathcal{H}_{\sigma_k \sigma_l} v(t, \mathbf{x}) = \mathbf{0}, \quad (2.45)$$

where $\mathcal{H}_{\sigma_k \sigma_l} = \partial_{\mathbf{x}}' (\partial_{\sigma_k \sigma_l}^2 \Sigma) \partial_{\mathbf{x}}$. Applying the Fourier transform to the mixed Volga PIDE yields an ODE with a source term, which can be solved explicitly:

$$\partial_{\sigma_k \sigma_l} v(t, \mathbf{x}) = (T - t) \mathcal{F}^{-1} [\mathcal{F}[\mathcal{H}_{\sigma_k \sigma_l}](\omega) \cdot \mathcal{F}[v](t, \omega)](\mathbf{x}). \quad (2.46)$$

The greekFST method for computing mixed Volga is then given by

$$\nabla_{k,l,m-1} = \Delta t_m \text{FFT}^{-1} [\mathcal{F}[\mathcal{H}_{\sigma_k \sigma_l}](\omega) \cdot \hat{\mathbf{v}}_{m-1}], \quad (2.47)$$

where the $\mathcal{F}[\mathcal{H}_{\sigma_k \sigma_l}](\omega)$ term can be computed analytically. For example, in the case of 2D BSM model, $\mathcal{F}[\mathcal{H}_{\sigma_k \sigma_l}](\omega) = -\rho \omega_1 \omega_2$.

For the case of Volga, assuming $k = l$ and applying twice the derivative with respect to σ_k to the pricing PIDE, a PIDE satisfied by Volga is obtained:

$$\partial_{\sigma_k}^2 \{(\partial_t + \mathcal{L}) v(t, \mathbf{x})\} = (\partial_t + \mathcal{L}) \partial_{\sigma_k}^2 v(t, \mathbf{x}) + \mathcal{H}_{\sigma_k^2} v(t, \mathbf{x}) = \mathbf{0}, \quad (2.48)$$

where $\mathcal{H}_{\sigma_k^2} = \mathcal{H}_{\sigma_k} \partial_{\sigma_k} + \partial_{\sigma_k} \mathcal{H}_{\sigma_k}$. Applying the Fourier transform to the Volga PIDE yields an ODE with a source term, which can be solved explicitly:

$$\partial_{\sigma_k}^2 v(t, \mathbf{x}) = (T - t) \mathcal{F}^{-1} \left[\mathcal{F} \left[\mathcal{H}_{\sigma_k^2} \right] (\boldsymbol{\omega}) \cdot \mathcal{F} [v](t, \boldsymbol{\omega}) \right] (\mathbf{x}). \quad (2.49)$$

The greekFST method for computing mixed Volga is then given by

$$\nabla_{k,m-1} = \text{FFT}^{-1} \left[\mathcal{F} \left[\mathcal{H}_{\sigma_k^2} \right] (\boldsymbol{\omega}) \cdot \hat{\mathbf{v}}_{m-1} \right], \quad (2.50)$$

where the $\mathcal{F} \left[\mathcal{H}_{\sigma_k^2} \right] (\boldsymbol{\omega})$ can be computed analytically by using the previously computed solution to the ODE satisfied by Vega. For example, in the case of the BSM model $\mathcal{F} \left[\mathcal{H}_{\sigma_k^2} \right] (\boldsymbol{\omega}) = -(i\omega + \omega^2) \Delta t_m + (i\omega + \omega^2)^2 \sigma^2 \Delta t_m^2$.

Theta

Theta measures the rate of change of option value with respect to the passage of time. Similarly to Delta and Gamma, the Fourier transform of Theta can be computed from the Fourier transform of option values via scaling:

$$\partial_t v(t, \mathbf{x}) = -\mathcal{L}v(t, \mathbf{x}) = \mathcal{F}^{-1} [-\Psi(\boldsymbol{\omega}) \cdot \mathcal{F} [v](t, \boldsymbol{\omega})] (\mathbf{x}). \quad (2.51)$$

The greekFST method for computing Theta is then given by

$$\boldsymbol{\Theta}_{m-1} = \text{FFT}^{-1} [-\Psi(\boldsymbol{\omega}) \cdot \hat{\mathbf{v}}_{m-1}]. \quad (2.52)$$

Rho

Rho measures the sensitivity of the option price with respect to the change in the interest rate. Applying derivative with respect to r to the pricing PIDE, a PIDE satisfied by Rho is obtained:

$$\partial_r \{ (\partial_t + \mathcal{L}) v(t, \mathbf{x}) \} = (\partial_t + \mathcal{L}) \partial_r v(t, \mathbf{x}) + \mathcal{H}_r v(t, \mathbf{x}) = \mathbf{0}, \quad (2.53)$$

where $\mathcal{H}_r = (\partial_r \gamma)' \partial_{\mathbf{x}}$. Similarly to the case of Vega and Volga, the Rho PIDE can be solved explicitly in Fourier space:

$$\partial_r v(t, \mathbf{x}) = (T - t) \mathcal{F}^{-1} [\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega}) \cdot \mathcal{F} [v](t, \boldsymbol{\omega})] (\mathbf{x}). \quad (2.54)$$

The greekFST method for computing Rho is then given by

$$\mathbf{P}_{m-1} = \Delta t_m \text{FFT}^{-1} [\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega}) \cdot \hat{\mathbf{v}}_{m-1}], \quad (2.55)$$

where the $\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega})$ term can be computed analytically. For example, in the case of the BSM model $\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega}) = i\omega$. Note that $v(\mathbf{x})$ is the discount adjusted option price; to obtain Rho for unadjusted option price $V(\mathbf{x})$, the source term $\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega})$ is modified to be $\mathcal{F} [\mathcal{H}_r] (\boldsymbol{\omega}) - 1$.

Figure 2.7: Error in option price (top, left), Delta (top, right), Gamma (middle, left), Rho (middle, right), Theta (bottom, left), and Vega (bottom, right) for the European option *EUR-E* under the Black-Scholes-Merton model *BSM-B*. The average rate of convergence (across all spot prices) for the option price and all Greeks is 2 in space.

Figure 2.7 depicts the option price, Delta, Gamma, Vega, Theta, and Rho errors for the European option *EUR-E* under the BSM model *BSM-B*. Notice that the dips in the error plots occur due to oscillation of the computed value around the true value. When the difference between the values changes sign due to such oscillation, the log of the absolute error tends to 0. The order of convergence of the greekFST method is 2 in space. Additional results are given in Figure C.1 in Appendix C.3, which depicts the option price and Greeks errors for the digital option *DIG-A* under the Merton jump-diffusion model *MJD-A*.

As in the computation of option prices using the FST method, option Greeks are computed for a range of spot prices. Also, notice that the computation of both the option values and the Greeks at time t_{m-1} requires the DFT of option values $\hat{\mathbf{v}}_{m-1}$. This DFT can be reused, so that the computation of both the option values and the k different Greeks requires only $2 + k$ FFT evaluations instead of $2k$.

Chapter 3

Mean-Reverting Fourier Space Time-stepping Method

3.1 Introduction

In this chapter, a multi-factor mean-reverting commodity model with Lévy drivers is introduced. The framework discussed here is quite flexible: it is able to capture many standard commodity models available in the literature and general enough to obtain new model specifications (e.g., multiple sources of jumps with differing mean-reversion scales). To illustrate the flexibility, the chapter explores both existing and new model specifications that can be obtained. Based on this model, an FST framework-based methodology is developed for valuing general commodity contingent claims, with motivation to price highly path-dependent American-style commodity options¹.

In this framework, the commodity spot price is driven by a mean-reverting jump-diffusion process. The discount-adjusted and log-transformed option price process satisfies a PIDE, which must be solved numerically in most cases. The solution to the PIDE produces option values for European options on a space-time grid. To value barrier options, additional boundary conditions along the barrier(s) supplement the PIDE. To value American- or Bermudan-style options, the PIDE is satisfied only in the continuation region and the optimal exercise point must be found at each time step.

With mean-reversion, the characteristic exponent of the commodity price process is state-dependent and the FST method cannot be readily applied. In this chapter, the FST framework is extended to handle mean-reverting processes and it is shown how the mrFST method can efficiently price European, Bermudan, American and barrier options. The mrFST method

¹See Eydeland and Wolyniec (2003) for an overview of commodity options

retains all of the advantages of the original FST method: European options can be priced using a single time step to obtain option values for a range of spot prices and Bermudan options do not require time-stepping between monitoring dates. Furthermore, the method can readily handle multi-asset options and easily extends to the regime-switching case, while maintaining the efficiency of the FST framework.

The outline of the remainder of this chapter is as follows. Section 3.2 presents the multi-factor mean-reverting model for the evolution of commodity spot prices. In Section 3.3 the option pricing PIDE is solved in Fourier space. Section 3.4 formulates the mrFST method based on the Fourier space solution and numerical results for pricing various options are given in Section 3.5.

3.2 Spot Price Model

Rather than modeling the n commodity spot prices $\mathbf{S}(t)$ themselves, their logarithms $\mathbf{X}(t)$ are modeled as a linear transformation of a set of d fundamental market factors $\mathbf{Y}(t)$ that are driven by the continuous-time counterpart of a VAR(1) model. In the discrete setting, the fundamental market factors evolve according to the time series model

$$\mathbf{Y}(t + \Delta t) = -\boldsymbol{\kappa}\mathbf{Y}(t)\Delta t + \boldsymbol{\epsilon}(t), \quad (3.1a)$$

$$\mathbf{X}(t) = \boldsymbol{\theta} + \boldsymbol{\Lambda}\mathbf{Y}(t). \quad (3.1b)$$

Here, $\boldsymbol{\kappa}$ is a $d \times d$ matrix, with eigenvalues that are less than 1 in magnitude, representing the mixing of the market factors, $\boldsymbol{\theta}$ is an n -dimensional vector representing the long-run means, $\boldsymbol{\Lambda}$ is a $d \times n$ matrix representing the linear transformation of the market factors into the observed log-prices, and $\boldsymbol{\epsilon}(0), \boldsymbol{\epsilon}(1) \dots$ are i.i.d. d -dimensional noise vectors (possibly heavy tailed) which, in the continuous-time model, are modeled via a Lévy process.

Notice that only when $n = d$ and $\boldsymbol{\Lambda}$ is invertible are the processes $\mathbf{Y}(t)$ available from the observed prices $\mathbf{X}(t)$. As such, if $\mathbf{Y}(t)$ are unobservable, the model is in general a hidden Markov model, and the initial values of the hidden processes $\mathbf{Y}(t)$ must be estimated through Kalman or particle filters. Consequently, to simplify the analysis, it is assumed that the processes $\mathbf{Y}(t)$ are either directly observable or obtainable through a filtering approach; however, it is not assumed they are tradable.

The continuous-time counterpart to the discrete model (3.1) is introduced by first defining the Lévy sources of risk $\mathbf{J}(t)$ which drive the market factors $\mathbf{Y}(t)$. Let $\mathbf{J}(t)$ denote a d -dimensional Lévy process with Lévy triple $(\boldsymbol{\gamma}, \boldsymbol{\Sigma}, \boldsymbol{\nu})$, where $\boldsymbol{\gamma}$ represents the vector of unadjusted drifts, $\boldsymbol{\Sigma}$ represents the variance-covariance matrix of the diffusions, and $\boldsymbol{\nu}$ is the multi-dimensional Lévy density. In this case, the process $\mathbf{J}(t)$ admits the following canon-

cal Lévy-Itô decomposition into its diffusion and jump components (see Section 2.2 for more details):

$$\begin{aligned}\mathbf{J}(t) &= \boldsymbol{\gamma} t + \mathbf{W}(t) + \mathbf{J}^l(t) + \lim_{\epsilon \searrow 0} \mathbf{J}^\epsilon(t), \\ \mathbf{J}^l(t) &= \int_0^t \int_{|\mathbf{z}| \geq 1} \mathbf{z} \mu(d\mathbf{z} \times ds), \\ \mathbf{J}^\epsilon(t) &= \int_0^t \int_{\epsilon \leq |\mathbf{z}| < 1} \mathbf{z} [\mu(d\mathbf{z} \times ds) - \nu(d\mathbf{z} \times ds)] .\end{aligned}$$

Here $\mathbf{W}(t)$ is a vector of correlated Brownian motions with covariance matrix $\boldsymbol{\Sigma}$, $\mu(d\mathbf{z} \times ds)$ is a Poisson random measure counting the number of jumps of size $\mathbf{z}$ occurring at time s , and $\nu(d\mathbf{z} \times ds) = \nu(d\mathbf{z}) ds$ is its compensator. Note that $\mathbf{J}^l(t)$ and $\mathbf{J}^\epsilon(t)$ carry the interpretation of large and small jumps, respectively.

The continuous-time counterpart to the discrete model (3.1) is then defined as follows:

$$d\mathbf{Y}(t) = -\boldsymbol{\kappa}\mathbf{Y}(t-)dt + d\mathbf{J}(t), \quad (3.2a)$$

$$\mathbf{X}(t) = \boldsymbol{\theta} + \boldsymbol{\Lambda}\mathbf{Y}(t), \quad (3.2b)$$

$$\mathbf{S}(t) = \mathbf{S}(0)e^{\mathbf{X}(t)}. \quad (3.2c)$$

Here, $\boldsymbol{\kappa}$ is a constant matrix representing the mixing of the market factors, $\boldsymbol{\theta}$ is a constant vector representing the long-run means, and, without loss of generality², the drift of the jump process is set to zero $\boldsymbol{\gamma} = \mathbf{0}$. This modeling framework is affine and is similar to that of Duffie, Pan, and Singleton (2000) for modeling interest rates and valuing quanto options among others. However, in that work, the methodology was restricted to valuing European-style claims and required the analytical valuation of the Fourier transform of the payoff function. Here, the framework is considered within the context of commodities, the analytic transform of payoff is not required, and path-dependent and early-exercise options can be valued easily.

Before discussing the valuation issues, a few specific examples are explored to illustrate the flexibility of this modeling framework.

Example 1: Mean-reverting jump-diffusion

Take $d = n = 1$, $\boldsymbol{\theta} = \theta$, $\boldsymbol{\kappa} = \kappa$, $\boldsymbol{\Sigma} = \sigma^2$, $\nu(d\mathbf{z}) = 0$ and $\boldsymbol{\Lambda} = 1$. This corresponds to the Gibson and Schwartz (1990) one-factor mean-reverting model:

$$d \ln S(t) = \kappa(\theta - \ln S(t)) dt + \sigma dW(t). \quad (3.3)$$

²The mean-level of $\mathbf{Y}(t)$ is fixed by both $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$. As such, there is a degeneracy which can be removed by either fixing $\boldsymbol{\theta}$ or $\boldsymbol{\gamma}$. Fixing $\boldsymbol{\gamma}$ provides simpler interpretations.

Also, by letting $\nu(d\mathbf{z}) = \lambda dF(z)$, where F is the normal distribution function, one obtains the Clewlow and Strickland (2000) model:

$$d \ln S(t) = \kappa(\theta - \ln S(t)) dt + \sigma dW(t) + dJ(t). \quad (3.4)$$

Other choices for the distribution function F , such as the double-exponential distribution function, can be used within this framework. See also Cartea and Figueroa (2005), where forward prices for a general distribution function $F(z)$ are derived.

Example 2: Mean-reverting jump-diffusion with decoupled jumps

Take $d = 2$, $n = 1$,

$$\boldsymbol{\theta} = (\theta), \quad \boldsymbol{\kappa} = \begin{pmatrix} \kappa_1 & 0 \\ 0 & \kappa_2 \end{pmatrix}, \quad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma^2 & 0 \\ 0 & 0 \end{pmatrix}, \quad \boldsymbol{\Lambda} = \begin{pmatrix} 1 & 1 \end{pmatrix}, \quad (3.5)$$

and $\nu(dZ_1 \times dZ_2) = \lambda \delta_{Z_1} dF(Z_2)$, where δ_z denotes the Dirac measure at 0 and $F(z)$ is a distribution function. This model corresponds to a mean-reverting jump-diffusion with different decay rates for the jumps and diffusion. In particular, log-prices have volatility of σ and mean-revert to level θ at a rate of κ_1 . Jumps arrive at a rate of λ , causing log-prices to jump with distribution function $F(z)$, and revert back to zero at a rate of κ_2 :

$$\ln S(t) = Y_1(t) + Y_2(t), \quad (3.6a)$$

$$dY_1(t) = \kappa_1(\theta - Y_1(t)) dt + \sigma dW(t), \quad (3.6b)$$

$$dY_2(t) = -\kappa_2 Y_2(t) dt + dJ(t), \quad (3.6c)$$

$$J(t) = \sum_{n=1}^{N(t)} j_n. \quad (3.6d)$$

Here, $j_1, j_2, \dots$ are i.i.d. with distribution function $F(z)$ and $N(t)$ is a Poisson process with activity λ . This model was proposed in Hikspoors and Jaimungal (2007) for electricity pricing as spikes are typically pulled back much faster than the diffusion components.

Example 3: Mean-reverting jump-diffusions with codependent jumps

Take $d = 2$, $n = 2$,

$$\boldsymbol{\theta} = \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix}, \quad \boldsymbol{\kappa} = \begin{pmatrix} \kappa_1 & 0 \\ 0 & \kappa_2 \end{pmatrix}, \quad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}, \quad \boldsymbol{\Lambda} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (3.7)$$

and $\nu(dz_1 \times dz_2) = dC(F_1(z_1), F_2(z_2))$, with $C(u, v)$ a copula and $F_1(z), F_2(z)$ two marginal distribution functions. Such a model corresponds to a two-dimensional jump-diffusion model

where the diffusions are correlated, and jumps may have codependent pieces. It is the presence of the copula function here which allows for either codependent jumps or independent jumps. This model can be used when two commodities are strongly dependent; however, it allows for jumps in one price without necessitating a jump in the other. Copulas are discussed in greater detail in Section 4.3.

Note the subtle, yet crucial difference between Examples 2 and 3. Both models are two-factor models where each factor is a mean-reverting process. However, in Example 2 the second factor has no diffusion, while in Example 3 there is an additional codependent jump component common to the two factors. The crucial difference between the two examples is that the two factors in Example 2 drive a single spot price process, making each driving factor unobservable, while in Example 3 each factor drives a distinct spot price process, hence each factor is observable.

Example 4: Mean-reverting GBM with mean-reverting reversion level

Take $d = 2$, $n = 1$,

$$\boldsymbol{\theta} = \begin{pmatrix} \theta \end{pmatrix}, \quad \boldsymbol{\kappa} = \begin{pmatrix} \kappa_1 & -\kappa_1 \\ 0 & \kappa_2 \end{pmatrix}, \quad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix}, \quad \boldsymbol{\Lambda} = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad (3.8)$$

and $\nu(dz_1 \times dz_2) = 0$. This corresponds to a two factor mean-reverting model. In this case, the log-prices mean-revert to a stochastic long-run mean which itself mean-reverts to a fixed level:

$$d \ln S(t) = \kappa_1(\Theta(t) - \ln S(t))dt + \sigma_1 dW_1(t), \quad (3.9a)$$

$$d\Theta(t) = \kappa_2(\theta - \Theta(t))dt + \sigma_2 dW_2(t), \quad (3.9b)$$

with $dW_1(t)dW_2(t) = \rho dt$. This model, similar to the stochastic convenience yield model of Schwartz (1997), was introduced in Barlow, Gusev, and Lai (2004) where the authors developed Kalman filter estimates for the hidden stochastic long run mean process $\Theta(t)$. The authors then fitted the model parameters to electricity price data. They found that the fit was not very good — this is no surprise as there are no spikes in model (3.9), while electricity data is known to have notoriously large spikes. On the other hand, the model is further analyzed in Hikspoors and Jaimungal (2007) in the context of valuation of oil derivatives. The authors find that the model calibrated well to oil spot and futures data.

3.3 PIDE Solution

Given the model (3.2), an expression for European option prices for all spots using a Fourier transform representation is derived in this section. Then, through this representation, Bermudan, American, barrier and other path-dependent options can be valued by applying the FST

methodology. The fundamental theorem of asset pricing, implies that the discount-adjusted price process $v(t, \mathbf{Y}(t)) \triangleq e^{r(T-t)} V(t, \mathbf{Y}(t))$ is a martingale under the risk-neutral measure $\mathbb{Q}$. Consequently, the option price must satisfy the PIDE

$$\begin{cases} (\partial_t + \mathcal{L}) v(t, \mathbf{y}) &= 0, \\ v(T, \mathbf{y}) &= \varphi(\mathbf{S}(0) e^{\boldsymbol{\theta} + \boldsymbol{\Lambda} \mathbf{y}}), \end{cases} \quad (3.10)$$

where $\mathcal{L}$ is the infinitesimal generator of the $\mathbf{Y}(t)$ process and acts on twice-differentiable functions $g(\mathbf{y})$ as follows:

$$\mathcal{L}g(\mathbf{y}) = (-\boldsymbol{\kappa} \mathbf{y}' \partial_{\mathbf{y}} + \frac{1}{2} \partial_{\mathbf{y}}' \boldsymbol{\Sigma} \partial_{\mathbf{y}}) g(\mathbf{y}) + \int_{\mathbb{R}^d / \{0\}} [g(\mathbf{y} + \mathbf{z}) - g(\mathbf{y}) - \mathbb{1}_{\{|\mathbf{z}| < 1\}} \mathbf{z}' \partial_{\mathbf{y}} g] \nu(d\mathbf{z}). \quad (3.11)$$

The difference of the above infinitesimal generator from the independent-increment exponential Lévy infinitesimal generator (2.3) is that the $\boldsymbol{\gamma}' \partial_{\mathbf{x}}$ term in (2.3) is replaced by $-\boldsymbol{\kappa} \mathbf{y}' \partial_{\mathbf{y}}$ in the mean-reverting model (as previously mentioned, the drift $\boldsymbol{\gamma}$ is set to $\mathbf{0}$). The $\mathbf{y}$ factor reflects the state-dependent nature of the mean-reverting model.

Transforming (3.10) into the Fourier domain $v(t, \mathbf{Y}) \mapsto \mathcal{F}[v](t, \boldsymbol{\omega})$ leads to a PDE in frequency space

$$\begin{cases} (\partial_t + \psi(\boldsymbol{\omega}) + \boldsymbol{\omega}' \boldsymbol{\kappa} \partial_{\boldsymbol{\omega}} + \text{Tr } \boldsymbol{\kappa}) \mathcal{F}[v](t, \boldsymbol{\omega}) &= 0, \\ \mathcal{F}[v](T, \boldsymbol{\omega}) &= \mathcal{F}[\varphi](\boldsymbol{\omega}), \end{cases} \quad (3.12)$$

where $\psi(\boldsymbol{\omega})$ is the characteristic exponent defined in (2.1). The additional two terms in (3.12), compared to the pricing ODE (2.7), arise from $\mathcal{F}[-\boldsymbol{\kappa} \mathbf{y}' \partial_{\mathbf{y}}] = \boldsymbol{\omega}' \boldsymbol{\kappa} \partial_{\boldsymbol{\omega}} + \text{Tr } \boldsymbol{\kappa}$. The PDE (3.12) is solved by first converting it into an ODE. By introducing a new coordinate system via frequency scaling

$$\mathcal{F}[\tilde{v}](t, \boldsymbol{\omega}) = \mathcal{F}[v](t, e^{-\boldsymbol{\kappa}'(T-t)} \boldsymbol{\omega}), \quad (3.13)$$

the PDE (3.12) reduces to an ODE in time

$$\begin{cases} (\partial_t + \psi(e^{-\boldsymbol{\kappa}'(T-t)} \boldsymbol{\omega}) + \text{Tr } \boldsymbol{\kappa}) \mathcal{F}[\tilde{v}](t, \boldsymbol{\omega}) &= 0, \\ \mathcal{F}[\tilde{v}](T, \boldsymbol{\omega}) &= \mathcal{F}[\tilde{\varphi}](\boldsymbol{\omega}). \end{cases} \quad (3.14)$$

This ODE is easily solved and, after changing coordinates back to the original ones and taking inverse Fourier transforms, the result on which the algorithm is based is obtained. Given the value of $v(\cdot, \mathbf{y})$ at time $t_2 \leq T$, the price of a European option at time $t_1 < t_2$, written on the vector of price processes defined by model (3.2) is

$$v(t_1, \mathbf{y}) = \mathcal{F}^{-1} \left[\mathcal{F}[v](t_2, e^{\boldsymbol{\kappa}'(t_2-t_1)} \boldsymbol{\omega}) \cdot e^{\Psi(t_2-t_1, \boldsymbol{\omega}) + (t_2-t_1) \text{Tr } \boldsymbol{\kappa}} \right] (\mathbf{y}), \quad (3.15)$$

where

$$\Psi(\Delta t, \boldsymbol{\omega}) = \int_0^{\Delta t} \psi(e^{\boldsymbol{\kappa}' u} \boldsymbol{\omega}) du, \quad \text{and} \quad (3.16)$$

$$\psi(\boldsymbol{\omega}) = -\frac{1}{2} \boldsymbol{\omega}' \boldsymbol{\Sigma} \boldsymbol{\omega} + \int_{\mathbb{R}^d / \{0\}} \left(e^{i \boldsymbol{\omega}' \mathbf{y}} - 1 - i \mathbb{1}_{\{|\mathbf{y}| < 1\}} \boldsymbol{\omega}' \mathbf{y} \right) \nu(d\mathbf{y}). \quad (3.17)$$

Notice that the function $\psi(\boldsymbol{\omega})$ is the characteristic exponent of the driving Lévy process, while the function $\Psi(\Delta t, \boldsymbol{\omega})$ accounts for the mean-reverting nature of the process through the exponential rescaling of frequencies. The function $\Psi(\Delta t, \boldsymbol{\omega})$ can be written in terms of the exponential integral function $Ei(x) \triangleq -\int_{-x}^{\infty} (e^{-t}/t) dt$ by integrating over time:

$$\begin{aligned} \Psi(\Delta t, \boldsymbol{\omega}) = & -\frac{1}{2}(\boldsymbol{\Gamma}\boldsymbol{\omega})'\tilde{\boldsymbol{\Sigma}}(\Delta t)(\boldsymbol{\Gamma}\boldsymbol{\omega}) + \int_{\mathbb{R}^d/\{0\}} \left\{ \boldsymbol{\kappa}^{-1} (\text{Ei}(-\boldsymbol{\omega}'\mathbf{y}) - \text{Ei}(-\boldsymbol{\omega}'e^{\boldsymbol{\kappa}\Delta t}\mathbf{y})) - \Delta t \right. \\ & \left. - \mathbb{1}_{\{|\mathbf{y}|<1\}} \boldsymbol{\omega}'\boldsymbol{\kappa}^{-1}(e^{\boldsymbol{\kappa}\Delta t} - \mathbf{I})\mathbf{y} \right\} \nu(d\mathbf{y}). \end{aligned} \quad (3.18)$$

where,

$$\langle \tilde{\boldsymbol{\Sigma}}(\Delta t) \rangle_{kl} = \frac{e^{(\lambda_k + \lambda_l)\Delta t} - 1}{\lambda_k + \lambda_l} \langle \boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma}' \rangle_{kl}, \quad (3.19)$$

and $\boldsymbol{\Gamma}$ is the matrix of orthonormalized eigenvectors of $\boldsymbol{\kappa}$ stacked column-wise with eigenvalues $\{\lambda_1, \dots, \lambda_d\}$.

On first glance, the appearance in expression (3.15) of the scaled frequency of the transform of option prices $\mathcal{F}[v](\cdot, e^{\boldsymbol{\kappa}'\Delta t}\boldsymbol{\omega})$ seems to pose problems. However, using the scaling property of Fourier transforms, the scaled option prices in frequency space can be obtained from the scaled option prices in real space:

$$\mathcal{F}[v](\cdot, e^{\boldsymbol{\kappa}'\Delta t}\boldsymbol{\omega}) = \mathcal{F}[\check{v}](\cdot, \boldsymbol{\omega}) e^{-\Delta t \text{Tr}\boldsymbol{\kappa}}, \quad (3.20)$$

where $\check{v}(\mathbf{y}) \triangleq v(\mathbf{y} e^{-\boldsymbol{\kappa}'\Delta t})$. Using equation (3.20) in equation (3.15) provides an alternate representation for European option prices (which is used in the derivation of the discrete method):

$$v(t_1, \mathbf{y}) = \mathcal{F}^{-1} \left[\mathcal{F}[\check{v}](t_2, \boldsymbol{\omega}) e^{\Psi(t_2-t_1, \boldsymbol{\omega})} \right] (\mathbf{y}). \quad (3.21)$$

Although this change in representation may appear trivial, it greatly reduces the complexity of the method by shifting the need to evaluate high-frequency modes, which requires *extrapolating* the transformed prices, to the need to evaluate small-scale spatial representation, which requires *interpolating* the prices in real space. Numerically, the former is significantly harder to do accurately than the latter.

Finally, for energy commodities it is also interest to value options on forward prices. Forward prices follow as an easy consequence of equation (3.21). Specifically, the T -maturity forward prices $\{F_1(t, T), \dots, F_n(t, T)\}$ based on the spot price model (3.2) are

$$\mathbf{F}_k(t, T) = \exp \left\{ \theta_k + \Psi(-i e^{-\boldsymbol{\kappa}'(T-t)} \boldsymbol{\lambda}_k, T-t) + e^{-\boldsymbol{\kappa}'(T-t)} \boldsymbol{\lambda}_k \mathbf{Y}(t) \right\} \quad (3.22)$$

where, $\langle \boldsymbol{\lambda}_k \rangle_l = \langle \boldsymbol{\Lambda} \rangle_{kl}$. A simple application of equation (3.21) with payoff function $\varphi_k(\mathbf{y}) = \mathbf{S}(0) e^{\boldsymbol{\theta} + \boldsymbol{\Lambda}\mathbf{y}}$ provides the result.

3.4 The Method

Let $\Omega = [-\mathbf{y}^{\min}, \mathbf{y}^{\max}]$ be a finite d -dimensional discrete space domain with N^d uniformly spaced points. The discretization of the frequency domain is then fixed as $\hat{\Omega} = [0, \boldsymbol{\omega}^{\max}]$, with the Nyquist condition $\Delta \mathbf{y} \cdot \Delta \boldsymbol{\omega} = 1/N$ being satisfied in each dimension. Numerical experiments show that choosing $\mathbf{y}^{\min}, \mathbf{y}^{\max} \in [4, 7]$ provides accurate results for most options and commodity price processes, although, for price processes having large spikes, larger values may be needed³. Let $t = t_0, t_1, \dots, t_M = T$, $\Delta t_m = t_m - t_{m-1}$ be a discretization of the time domain into M intervals and, as before, let $\mathbf{v}_m$ denote the discrete array of option values at time t_m .

The pricing equation (3.21) provides prices for a full range of spot values. Replacing the CFTs with DFTs, which in turn are computed using the FFT algorithm, leads to mrFST method for propagating a price one time step back:

$$\mathbf{v}_{m-1} = \text{FFT}^{-1} \left[\text{FFT} [\check{\mathbf{v}}_m] \cdot e^{\Psi(\Delta t_m, \cdot)} \right]. \quad (3.23)$$

Here, $\check{\mathbf{v}}_m$ represents the spatial rescaling of $\mathbf{v}_m$ as defined in equation (3.20).

European options can be priced using a single time step ($\Delta t = T - t$) of the mrFST method. Thus, two evaluations of the FFT algorithm and one scaling operation provides option prices for a range of spot prices. The mrFST algorithm can also be used to value Bermudan-style claims by taking one time step between each pair of exercise dates and enforcing the optimal exercise condition explicitly:

$$\mathbf{v}_{m-1} = \max \left\{ \text{FFT}^{-1} \left[\text{FFT} [\check{\mathbf{v}}_m] \cdot e^{\Psi(\Delta t_m, \cdot)} \right], \mathbf{v}_M \right\}, \quad (3.24)$$

where $\mathbf{v}_M$ is the exercise value of the option. Barrier options can be priced similarly by enforcing appropriate barrier conditions. For example, for an up-and-out barrier option, the price from one exercise date to the next is computed by

$$\mathbf{v}_{m-1} = \text{FFT}^{-1} \left[\text{FFT} [\check{\mathbf{v}}_m] \cdot e^{\Psi(\Delta t_m, \cdot)} \right] \cdot \mathbb{1}_{\{\boldsymbol{\theta} + \boldsymbol{\Lambda} \mathbf{y} < \mathbf{B}\}} + R \cdot \mathbb{1}_{\{\boldsymbol{\theta} + \boldsymbol{\Lambda} \mathbf{y} \geq \mathbf{B}\}}, \quad (3.25)$$

where R is the rebate being paid upon crossing of the barrier $\mathbf{B}$.

When there is no closed-form expression for the characteristic function, the calculation of $\Psi(\Delta t, \boldsymbol{\omega})$ can become computationally expensive. However, different contracts on the same underlying asset use the same characteristic function; consequently, $e^{\Psi(\Delta t, \boldsymbol{\omega})}$ could be precomputed and stored both for the valuation of path-dependent options, where the cardinality of the set $\{\Delta t_m\}$ is typically small, and for the valuation of a book of European options simultaneously.

Below the characteristic functions for the four examples, introduced in Section 3.2, are provided and associated computational aspects are discussed.

³See Chapter 2.4 for more details on grid selection for FST methods

Example 1: Mean-reverting jump-diffusion

$$\Psi(\Delta t, \omega) = -\frac{\omega^2 \sigma^2}{4\kappa} (e^{2\kappa \Delta t} - 1) + \int_0^{\Delta t} \psi(e^{\kappa u} \omega) du, \quad (3.26)$$

where ψ is the characteristic function of the jumps given by (3.16) with $\Sigma = \mathbf{0}$. For the Kou (2002) jump-diffusion model, in which jumps are double-exponentially distributed, the integral can be computed analytically

$$\int_0^{\Delta t} \psi(e^{\kappa u} \omega) du = \frac{\lambda}{\kappa} \left[\eta_p \ln \left(\frac{1 - i\omega\eta_+}{e^{-\kappa\Delta t} - i\omega\eta_+} \right) + (1 - \eta_p) \ln \left(\frac{1 + i\omega\eta_-}{e^{-\kappa\Delta t} + i\omega\eta_-} \right) \right] - \lambda \Delta t. \quad (3.27)$$

For the Merton (1976) jump-diffusion model, where jumps are normally distributed, the integral must be approximated numerically, for instance, using the composite quadrature rule

$$\int_0^{\Delta t} \psi(e^{\kappa u} \omega) du \approx \sum_{p=0}^P \zeta_p \psi(e^{\kappa u_p} \omega), \quad (3.28)$$

where u_p and ζ_p are the appropriate nodes and weights determined by the chosen quadrature rule and the number of subintervals P . Equation (3.28) gives accurate results, even for small number of integrand evaluations. For the Kou jump-diffusion model *KJDMR* model, the integral can be computed for all ω within an absolute error tolerance level of 10^{-3} by using a composite Simpson rule with 7 subintervals (and 9 subintervals for absolute error tolerance level of 10^{-4}). Choosing a larger number of subintervals has no significant effect on the resulting option prices in the numerical examples studied here.

Example 2: Mean-reverting jump-diffusion with decoupled jumps

$$\Psi(\Delta t, \omega_1, \omega_2) = -\frac{\omega_1^2 \sigma_1^2}{4\kappa_1} (e^{2\kappa_1 \Delta t} - 1) + \int_0^{\Delta t} \psi(e^{\kappa_2 u} \omega_2) du. \quad (3.29)$$

Again, the integral term depends on the jump distribution and is computed in the same fashion as in Example 1.

Example 3: Mean-reverting jump-diffusions with codependent jumps

$$\begin{aligned} \Psi(\Delta t, \omega_1, \omega_2) = & -\frac{\omega_1^2 \sigma_1^2}{4\kappa_1} (e^{2\kappa_1 \Delta t} - 1) - \frac{\omega_2^2 \sigma_2^2}{4\kappa_2} (e^{2\kappa_2 \Delta t} - 1) - \rho \omega_1 \omega_2 \sigma_1 \sigma_2 \frac{e^{(\kappa_1 + \kappa_2) \Delta t} - 1}{\kappa_1 + \kappa_2} \\ & + \int_0^{\Delta t} \psi_1(e^{\kappa_1 u} \omega_1) du + \int_0^{\Delta t} \psi_2(e^{\kappa_2 u} \omega_2) du + \int_0^{\Delta t} \psi_c(e^{\kappa_1 u} \omega_1, e^{\kappa_2 u} \omega_2) du. \end{aligned} \quad (3.30)$$

Here, ψ_1 and ψ_2 are the characteristic functions of the idiosyncratic jumps, given by equation (3.16) with $\Sigma = \mathbf{0}$, and $\psi_c(\omega_1, \omega_2)$ is the characteristic function of the codependent jumps:

$$\psi_c(\omega_1, \omega_2) = \lambda_c \left[\int_{\mathbb{R}^2} e^{i(\omega_1 x_1 + \omega_2 x_2)} f_1(x_1) f_2(x_2) c(F_1(x_1), F_2(x_2)) dx_1 dx_2 - 1 \right], \quad (3.31)$$

where $f_1(x_1), f_2(x_2)$ are the marginal density functions, $F_1(x_1), F_2(x_2)$ are the marginal cumulative distribution functions (cdf), $c(u_1, u_2)$ is the copula density function, and λ_c is the arrival rate of the codependent jumps. For more details on copulas see Section 4.3. The first two integrals in equation (3.30) are computed in the same manner as in Example 1 while the computation of the last integral depends on the chosen copula and marginal densities. For the special case of a Gaussian copula, with parameter ρ , and Gaussian marginal distribution of jumps, with means μ_1, μ_2 and variances σ_1^2, σ_2^2 , the codependent jumps have a bivariate Gaussian distribution and ψ is easily computed as

$$\psi_c(\omega_1, \omega_2) = \lambda_c \left(\exp \left\{ i\mu_1\omega_1 + i\mu_2\omega_2 - \frac{1}{2}\sigma_1^2\omega_1^2 - \rho\sigma_1\sigma_2\omega_1\omega_2 - \frac{1}{2}\sigma_2^2\omega_2^2 \right\} - 1 \right). \quad (3.32)$$

For other cases, the copula characteristic function (3.32) can be approximated numerically using adaptive quadrature, as in Example 1. A more efficient approach for computing such two-dimensional integrals is to utilize the FFT algorithm. Press, Teukolsky, Vetterling, and Flannery (1992) present the algorithm for computing of general one-dimensional Fourier integrals. Section 4.3 discusses its two-dimensional extension in the context of evaluating the characteristic exponent of a copula jump process.

Example 4: Mean-reverting GBM with mean-reverting reversion level

$$\Psi(\Delta t, \omega) = -\frac{1}{2}(\Gamma\omega)' \tilde{\Sigma}(\Delta t)(\Gamma\omega), \quad (3.33)$$

where $\tilde{\Sigma}(\cdot)$ is defined in equation (3.19).

3.5 Applications to Option Pricing

In this section, several numerical experiments are presented to investigate the convergence and precision of the mrFST method. Specifically, single-asset European, American and discrete barrier options, and multi-asset European and Bermudan spread options are priced under the four examples of spot price processes described in Section 3.2. As previously mentioned, the mrFST method provides option prices for a range of spot prices, which is a significant advantage for pricing path-dependent options, such as the Bermudan options presented below, and other applications, such as computation of the Greeks. Option and model parameters can be found in Appendix B.

The pricing results are verified by comparing them to closed-form formulas or Monte Carlo simulation results. Note that the number of time points used in the Monte Carlo simulation (and hence the time required to perform it) differ between the mean-reverting jump-diffusion models with log-normal and double-exponential jumps. Under the mean-reverting Merton jump-diffusion model, e.g., the *MJDMR* model, a closed-form expression exists for the terminal spot

Figure 3.1: Sample price paths under the mean-reverting Merton and Kou jump-diffusion models *MJDMR* and *KJDMR* (left) and the mean-reverting Kou jump-diffusion with decoupled jumps model *KJDMRD* (right).

price distribution. In such a case, the Monte Carlo method is quite efficient for pricing of European options. If instead, a mean-reverting Kou jump-diffusion model is used, e.g., the *KJDMR* model, the evolution of the spot price using many time steps must be computed and the simulation speed suffers. For the European options under the mean-reverting GBM with mean-reverting reversion level model, a closed-form price is provided.

Example 1: Mean-reverting jump-diffusion

In this example, two models are used in the pricing experiments. The first model, referred to as *MJDMR*, has log-normal jumps. It can be viewed as a mean-reverting extension of the Merton (1976) jump-diffusion model. The second model, referred to as *KJDMR*, has double-exponential jumps. It can be viewed as a mean-reverting extension of the Kou (2002) jump-diffusion model. Figure 3.1 shows sample paths of both models (both paths have identical Brownian increments at each time step). Notice the large upward jumps and rapid mean-reversion of the *KJDMR* model. For each of the two models three options are priced: European, American and discretely-monitored barrier options.

Pricing European option under mean-reverting jump-diffusion models (in theory) requires a single time step only, as outlined in Section 3.3. However, long option maturity combined with high mean-reversion rates can result in large scaling of the option value function by a factor of $e^{-\kappa\Delta t}$. In general, when pricing European options, the number of time steps is chosen to be $\lceil 2\kappa T \rceil$ to reduce the effect of extreme scaling. Such a choice guarantees that $e^{-\kappa\Delta t} > 0.6$ and the scaling remains moderate. Thus, pricing the European option *EUR-C* under the mean-reverting Merton jump-diffusion model *MJDMR* requires a single time step ($\kappa = 0.5, T = 0.5$)

N	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	10.20227645			6.312
4096	10.20274951	0.0004731		11.583
8192	10.20286764	0.0001181	2.0016	21.763
16384	10.20289715	0.0000295	2.0011	36.513
32768	10.20290453	0.0000074	2.0006	53.053

Table 3.1: Pricing results for the European option *EUR-C* under the mean-reverting Merton jump-diffusion model *MJDMR*. The reference price 10.20338594 is computed using Monte Carlo simulation (95% confidence interval width of 0.0086786 requires 20.235 seconds for 2^{25} sample paths with 1 time point). The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	256	5.15180139			0.109
4096	512	5.15127911	0.0005223		0.270
8192	1024	5.15114933	0.0001298	2.0089	1.124
16384	2048	5.15111672	0.0000326	1.9923	4.771
32768	4096	5.15110840	0.0000083	1.9709	20.588

Table 3.2: Pricing results for the American option *AMR-B* under the mean-reverting Merton jump-diffusion model *MJDMR* with Richardson extrapolation. The order of convergence is 2 in space and 2 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	252	1.97328658			0.073
4096	252	1.97534827	0.0020617		0.135
8192	252	1.97585589	0.0005076	2.0220	0.308
16384	252	1.97606794	0.0002121	1.2593	0.659
32768	252	1.97614506	0.0000771	1.4593	1.292

Table 3.3: Pricing results for the discrete barrier option *DBR-A* under the mean-reverting Merton jump-diffusion model *MJDMR*. The reference price 1.97628292 is computed using Monte Carlo simulation (95% confidence interval width of 0.0039546 requires 102.201 seconds for 2^{20} sample paths with 252 time points). The order of convergence is almost 2 in space.

and pricing the European option *EUR-B* under the mean-reverting Kou jump-diffusion model *KJDMR* requires 10 equal time steps ($\kappa = 5, T = 1$). The mrFST method attains second-order convergence in space when pricing European options, as shown in Table 3.1 and Table C.12 in Appendix C.4.

When pricing the American options *AMR-B* and *AMR-A* under the *MJDMR* and *KJDMR* models, respectively, small time steps are taken to monitor exercise behavior and therefore the scaling at each time step is small. By applying Richardson extrapolation, the mrFST method attains second-order convergence in time for pricing of American options, as shown in Table 3.2 and Table C.13 in Appendix C.4. Alternatively, an mrFST penalty method for pricing American options could be developed to improved the order of convergence but not discussed here.

For pricing the discrete barrier options *DBR-A* and *DBR-B* under the *MJDMR* and *KJDMR* models, respectively, the time steps are taken to equal the intervals between monitoring dates. Thus, for an option with daily monitoring and 1 year maturity, 252 time steps are required (equivalent to 252 trading days in a year). Weekends and holidays can also be accounted for by varying Δt at various time steps. The mrFST method has second-order convergence in space for pricing of the discrete barrier options *DBR-A* and *DBR-B*, as shown in Table 3.3 and Table C.14 in Appendix C.4.

Example 2: Mean-reverting jump-diffusion with decoupled jumps

The parameters for the mean-reverting Kou jump-diffusion with decoupled jumps model *KJDMRD* are

$$\theta = (\ln 50), \quad \kappa = \begin{pmatrix} 7.5 & 0 \\ 0 & 100 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} 1.0 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 1 & 1 \end{pmatrix}, \quad (3.34)$$

$r = 0.05$, and $\nu(d\mathbf{Z}) = \lambda dF(z_2)$, where $\lambda = 20.0$ and F is the double-exponential cdf with parameters $\eta_p = 0.99, \eta_+ = 0.4, \eta_- = 0.05$. Figure 3.1 provides a sample path for the *KJDMRD* model. Notice the extreme spikes and their quick reversion to the mean, common in electricity markets. However, the relatively low mean-reversion of the diffusion term allows for non-trivial diffusion structure. Under this model two options are priced: the European option *EUR-G* and the Bermudan option *BRM-A*.

The mrFST method has second-order convergence in space for both European and Bermudan options, as can be seen from Tables 3.4 and 3.5. Since the problems are two-dimensional, the computations are significantly slower than those in Example 1. For the *EUR-G* option, 16 time steps are taken to reduce the scaling required by the mrFST method. For the *BRM-A* option, 63 time steps are taken since the contract allows for daily exercise during the 3-month life of the options (total of 63 business days).

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512 ²	16	9.17654584			1.108
1024 ²	16	9.14611429	0.0304316		4.731
2048 ²	16	9.13874353	0.0073708	2.0457	19.812
4096 ²	16	9.13644634	0.0022972	1.6819	83.891

Table 3.4: Pricing results for the European option *EUR-G* under the mean-reverting Kou jump-diffusion with decoupled jumps model *KJDMRD*. The reference price 9.13188887 is computed using Monte Carlo simulation (95% confidence interval width of 0.0578199 requires 12.893 minutes for 2^{20} sample paths with 2^{10} time points). The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512 ²	63	53.33337341			3.725
1024 ²	63	53.09605240	0.2373210		16.227
2048 ²	63	53.01721019	0.0788422	1.5898	66.238
4096 ²	63	53.02131805	0.0041079	4.2625	298.261

Table 3.5: Pricing results for the Bermudan option *BRM-A* under the mean-reverting Kou jump-diffusion with decoupled jumps model *KJDMRD*. The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512 ²	2	31.02793888			1.710
1024 ²	2	31.01715659	0.0107823		6.589
2048 ²	2	31.01555052	0.0016061	2.7471	24.986
4096 ²	2	31.01540552	0.0001450	3.4694	98.691

Table 3.6: Pricing results for the European spread option *ESPD* under the 2D mean-reverting Kou jump-diffusion with Gaussian copula jumps model *KJDMRC*. The reference price 30.99793702 is computed using Monte Carlo simulation (95% confidence interval width of 0.126244 requires 18.086 minutes for 2^{20} sample paths with 2^{10} time points). The order of convergence is at least 2 in space.

Figure 3.2: Sample price paths under the 2D mean-reverting Kou jump-diffusion with Gaussian copula jumps model *KJDMRC* (left) and the geometric Brownian motion with mean-reverting reversion level model *GBMMRM* (right).

Example 3: Mean-reverting jump-diffusions with codependent jumps

The chosen parameters for the model are

$$\theta = \begin{pmatrix} \ln 92 \\ \ln 110 \end{pmatrix}, \quad \kappa = \begin{pmatrix} 0.5 & 0 \\ 0 & 0.75 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} 0.2^2 & 0.7 \cdot 0.06 \\ 0.7 \cdot 0.06 & 0.3^2 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (3.35)$$

$r = 0.05$, and $\nu(d\mathbf{Z}) = \lambda_1 dF_1(z_1) d(\mathbb{1}_{\{z_2 > 0\}}) + \lambda_2 d(\mathbb{1}_{\{z_1 > 0\}}) dF_2(z_2) + \lambda_c dC(F_{1c}(z_1), F_{2c}(z_2))$. The independent jumps have arrival rates of $\lambda_1 = 0.75$ and $\lambda_2 = 0.5$, respectively. The jumps in both dimensions have double-exponential distribution with parameters $\eta_p = 0.45, \eta_+ = 0.25, \eta_- = 0.125$ in the first dimension and $\eta_p = 0.55, \eta_+ = 0.3, \eta_- = 0.2$ in the second dimension. The copula jumps are driven by a Gaussian copula with correlation parameter 0.7 and have an arrival rate of $\lambda_c = 1.0$. The copula jumps have Gaussian distribution with means -0.1 and 0.1 , and variances 0.2^2 and 0.3^2 respectively. This model is referred to as *KJDMRC*. Figure 3.2 illustrates a sample path from the model. Notice the codependent structure of jumps in the two dimensions. As previously mentioned, the model is flexible enough to allow for independent jumps (see $t \approx 2.75$) and simultaneous jumps (see $t \approx 1.6$).

Under this model, two options are priced: the European spread option *ESPD* and the Bermudan spread option *BSPD*. Tables 3.6 and 3.7 show that the mrFST attains second-order convergence for European and Bermudan multi-asset options. For the *ESPD* option, 2 time steps are taken to reduce the effect of scaling. For the *BSPD* option, 12 time steps are required, as the option can only be exercised once a month over its lifetime of 1 year.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512^2	12	34.35252633			2.240
1024^2	12	34.30198167	0.0505447		9.070
2048^2	12	34.28961403	0.0123676	2.0310	35.712
4096^2	12	34.28671769	0.0028963	2.0943	147.489

Table 3.7: Pricing results for the Bermudan spread option *BSPD* under the 2D mean-reverting Kou jump-diffusion with Gaussian copula model *KJDMRC*. The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512^2	5	7.26020997			1.107
1024^2	5	7.23547685	0.0247331		4.752
2048^2	5	7.23117672	0.0043001	2.5240	18.677
4096^2	5	7.22989936	0.0012774	1.7512	75.446

Table 3.8: Pricing results for the European option *EUR-B* under the geometric Brownian motion with mean-reverting reversion level model *GBMMRM*. The reference price 7.22941336 is computed using a closed-form formula. The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
512^2	2	19.51361044			0.988
1024^2	2	19.50835432	0.0052561		4.074
2048^2	2	19.50704347	0.0013109	2.0035	16.059
4096^2	2	19.50671612	0.0003273	2.0016	64.532

Table 3.9: Pricing results for the European option *EUR-E* under the geometric Brownian motion with mean-reverting reversion level model *GBMMRM*. The reference price 19.50660534 is computed using a closed-form formula. The order of convergence is 2 in space.

Example 4: Mean-reverting GBM with mean-reverting reversion level

The chosen model parameters are

$$\theta = (\ln 75), \quad \kappa = \begin{pmatrix} 2.5 & -2.5 \\ 0 & 1.0 \end{pmatrix}, \quad \Sigma = \begin{pmatrix} 0.2^2 & 0.5 \cdot 0.2 \cdot 0.3 \\ 0.5 \cdot 0.2 \cdot 0.3 & 0.3^2 \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad (3.36)$$

$\Theta(0) = \ln 150$, and $r = 0.04$. This model is referred to as *GBMMRM* and a sample path under this model is given in Figure 3.2. Notice how the spot price process $S(t)$ reverts to the stochastic mean process $\exp \Theta(t)$. Under this model, the European options *EUR-B* and *EUR-E* are priced.

The results in Table 3.8 and 3.9 suggest that the mrFST method has second-order convergence in space. Note that 5 and 2 time steps, respectively, are required for *EUR-B* and *EUR-E* options respectively to reduce the effect of extreme scaling. While a closed-form solution exists for European options under the mean-reverting GBM with mean-reverting reversion level model, using the mrFST method allows for efficient pricing of path-dependent options and computation of Greeks.

Similarly to Example 1 in Section 2.6.1, the convergence order of the mrFST method across a range of spot prices is established by considering the differences between log errors $\log_{10} e^{[N, \cdot]}$ on successively finer grids. Figure 3.3 depicts the log errors of the mrFST method for pricing the European option *EUR-G* under the mean-reverting GBM model *GBMMR* and the European option *EUR-C* under the mean-reverting GBM with mean-reverting level model *GBMMRM*.

Figure 3.3: Errors for pricing the European option *EUR-G* under the mean-reverting GBM model *GBMMR* (left) and the European option *EUR-C* under the mean-reverting GBM with mean-reverting level model *GBMMRM* (right). The average rate of convergence (across all spot prices) for both scenarios is 2 in space.

The second-order convergence is retained over the range of spot prices as the difference between log errors in each successive refinement is $2 \cdot \log_{10} 2$. Again, the curvature in the log errors is due to the change in convergence constant across the various spot prices.

Chapter 4

Spot Price Model Extensions

4.1 Introduction

This chapter presents two spot price model extensions — regime-switching, which introduces regime changes into a stationary model via a Markov chain, and codependent jumps driven by a copula, which is crucial in modeling of assets driven by common factors.

In typical markets, jump models alone cannot match the implied volatility (IV) skew for long maturities; however, the observed market behavior can be captured by incorporating regime switches. This is the motivation to introduce a non-stationary extension of the multi-dimensional Lévy processes using regime changes. The regime changes are induced through a homogeneous continuous-time Markov chain. This allows the index(es) to exhibit stochastic volatility and/or stochastic correlation behavior which can be important for long-term options.

Copulas are a flexible tool to represent comovement between stock prices, assets or the underlying variables, and have become widely used in mathematical finance. Copula-based models are especially attractive in modeling of assets that respond in a codependent fashion to the changes in market economic conditions, supply-demand imbalance or information arrival in general.

The outline of the remainder of this chapter is as follows. Section 4.2 introduces the regime-switching extension of the multi-dimensional exponential Lévy model. The option pricing PIDE, arising under the non-stationary extension, is solved in Fourier space and the rsFST method for pricing under the regime-switching extension is derived. The effect of inducing regime-switching behavior into stationary models is explored through numerical examples. Section 4.3 develops a codependent-jumps model extension for two-asset models. An efficient FFT-based algorithm for computation of the characteristic function under such an extension is developed.

4.2 Regime-Switching

Regime-switching introduces regime changes into an otherwise stationary model and can be traced back to the early work of Lindgren (1978). Ever since the seminal work of Hamilton (1989, 1990) regime-switching has become a popular approach to incorporate non-stationary behavior into otherwise stationary models. The essential idea is to assume that the world switches between states representing, for example, moderate, low, and high volatility regimes.

Although regime-switching is popular for describing time series, little work has been done for option valuation. Two-state European options in log-normal models were studied in Naik (1993) while European options in a two-state VG model were studied by Konikov and Madan (2000). Albanese, Jaimungal, and Rubisov (2003) derive closed-form results for European and barrier options, and semi closed-form formulas for American options, in a special class of two-state VG models. Elliott and Osakwe (2006) propose an asset price model which is the exponential of a pure jump process with a multi-state Markov switching compensator. This section demonstrates that the FST framework can easily incorporate path-dependent options, such as barrier and American options, with multiple regimes and multiple assets in computational time proportional to the number of regimes. This work extends the method of Jackson, Jaimungal, and Surkov (2007) by developing a time-stepping technique without a time-step restriction.

Regime states can either be visible or hidden from market participants. If the states are hidden, then the initial probability of being in each state becomes part of the modeling assumptions and prices are provided by a weighted average of the conditional prices. If the states are visible, then the initial state is given a priori and no averaging is necessary. However, in both cases, at all future times, prices in all states are required to compute the conditional price one time step backwards. Let $\mathbb{K} \triangleq \{1, \dots, K\}$ denote the possible states of the world, and let $Z(t) \in \mathbb{K}$ denote the prevailing state of the world at time t . $Z(t)$ is assumed to be driven by a continuous-time Markov chain with generator $\mathbf{A}$, i.e., the transition probability from state k at time t_1 to state l at time t_2 is $P(Z(t_2) = l | Z(t_1) = k) = \langle \exp\{(t_2 - t_1)\mathbf{A}\} \rangle_{kl}$. The real matrix $\mathbf{A}$ satisfies the usual requirements: $A_{ll} = -\sum_{k \neq l} A_{lk}$ and $A_{lk} \geq 0 \forall k \neq l$.

For the numerical experiments in this section the transition rate between any two neighboring states is assumed to be λ . Under such a modeling assumption the Markov chain generator $\mathbf{A}$ is given by

$$\mathbf{A}(\lambda) = \begin{pmatrix} -\lambda & \lambda & 0 & 0 & 0 \\ \lambda & -2\lambda & \lambda & 0 & 0 \\ & \ddots & \ddots & \ddots & \\ 0 & 0 & \lambda & -2\lambda & \lambda \\ 0 & 0 & 0 & \lambda & -\lambda \end{pmatrix}. \quad (4.1)$$

However, it is important to note that the method derived in this section is general with respect to the structure of the Markov chain generator matrix $\mathbf{A}$.

Given that $Z(t) = k$, the joint stock price process $\mathbf{S}(t)$ is assumed to follow a d -dimensional exponential Lévy process with Lévy triple $(\gamma^{(k)}, \Sigma^{(k)}, \nu^{(k)})$. The drift vectors of each state are assumed prefixed at their risk-neutral levels of $\gamma_j^{(k)}$, such that $\Psi^{(k)}(-i\mathbf{1}_j) = r$ for each $j = 1, \dots, d$, where $\Psi^{(k)}(\omega)$ denotes the characteristic exponent of the respective Lévy processes and $\mathbf{1}_j$ is the vector with zeroes everywhere except a single entry of 1 at dimension j . This modeling assumption can succinctly be written $d\mathbf{X}(t) = d\mathbf{X}^{(Z(t))}(t)$, where $\mathbf{X}^{(k)}(t)$ is the k -th d -dimensional Lévy process. Unfortunately, this approach does not allow for correlation between the volatility level and the spot price.

Chourdakis (2005) investigates the $d = 1$ version of this framework and derives the characteristic function of the terminal stock price. The author calculates European option prices via FFT methods, however, resorts to numerical integration for the valuation of path-dependent options. Still, a quadrature evaluation for each grid point is required. Here, a slightly different approach is taken which makes use of a generalization of the FST framework and allows path-dependent options based on the regime-switching models to be valued efficiently.

Under the above assumptions, let $v^{(k)}(t, \mathbf{x})$ denote the discounted-adjusted and log-transformed price at time t conditional on the state $Z(t) = k$ and spot level $\mathbf{x}$. It is not difficult to show that European option prices satisfy the following system of PIDEs:

$$\begin{cases} (\partial_t + A_{kk} + \mathcal{L}^{(k)}) v^{(k)}(t, \mathbf{x}) + \sum_{l \neq k} A_{kl} v^{(l)}(t, \mathbf{x}) &= \mathbf{0}, \\ v^{(k)}(T, \mathbf{x}) &= \varphi(\mathbf{S}(0)e^{\mathbf{x}}) \end{cases} \quad (4.2)$$

for every $k \in \mathbb{K}$. Here, $\mathcal{L}^{(k)}$ represents the infinitesimal generator of the k -th d -dimensional Lévy process. It is possible in principle to apply any of the usual finite-difference schemes to this system of PIDEs to solve the problem. However, as discussed earlier, this is quite difficult due to the non-local integral terms and especially so for multi-dimensional problems. Instead, the FST methodology is applied to solve (4.2).

As before, applying Fourier transforms to (4.2) leads to a coupled system of ODEs indexed by the vector of frequencies ω :

$$\begin{cases} (\partial_t + A_{kk} + \Psi^{(k)}(\omega)) \mathcal{F}[v^{(k)}](t, \omega) + \sum_{l \neq k} A_{kl} \mathcal{F}[v^{(l)}](t, \omega) &= \mathbf{0}, \\ \mathcal{F}[v^{(k)}](T, \omega) &= \mathcal{F}[\varphi](\omega). \end{cases} \quad (4.3)$$

This can be rewritten in a compact matrix form

$$\begin{cases} (\partial_t + \Psi(\omega)) \mathcal{F}[\vec{v}](t, \omega) &= \mathbf{0}, \\ \mathcal{F}[\vec{v}](T, \omega) &= \mathcal{F}[\varphi](\omega) \vec{1}. \end{cases} \quad (4.4)$$

Here, $\vec{v}$ is the collection of $v^{(k)}$'s stacked into a column vector, $\vec{1}$ is a column vector of K 1's, and the elements of the matrix characteristic function Ψ are

$$\langle \Psi(\omega) \rangle_{kl} \triangleq \begin{cases} A_{kk} + \Psi^{(k)}(\omega), & k = l, \\ A_{kl}, & k \neq l. \end{cases}$$

Given the homogenous matrix form of the coupled system of ODEs (4.4), it is easy to see that the solution for the vector of the transformed prices is

$$\mathcal{F}[\vec{v}](t, \omega) = \exp\{\Psi(\omega)(T - t)\} \cdot \mathcal{F}[\varphi](\omega) \vec{1}. \quad (4.5)$$

For the European case, option prices may be obtained in a single step — even with regime-switching. Discretizing the spatial and frequency domains as before, the following scheme for European options is obtained:

$$\vec{v} = \text{FFT}^{-1} \left[e^{\Psi(T-t)} \cdot \text{FFT}[\varphi] \vec{1} \right]. \quad (4.6)$$

In a visible-state regime-switching model, the option price is the entry in the $Z(0)$ -th row of the vector $\vec{v}$ (recall that $Z(0)$ is the initial state of the world); while in a hidden-state regime-switching model, the option price is a weighted average of all prices $\vec{p} \cdot \vec{v}$, where $\vec{p}$ represents the vector of probabilities the investor assigns¹ to being in a given state.

For Bermudan options, the scheme must be modified to account for the early-exercise feature. Here, only the simple time-stepping algorithm is provided without penalty terms for the Bermudan put option. If the vector of prices $\vec{v}_m$ is known at time step m , then the conditional holding price $\vec{v}_{m-1}^*$ at time step $m - 1$ is

$$\vec{v}_{m-1}^* = \text{FFT}^{-1} \left[e^{\Psi \Delta t} \cdot \text{FFT}[\vec{v}_m] \right], \quad (4.7)$$

with $\mathbf{v}_M = \varphi(\mathbf{S}(0)e^{\mathbf{x}})$ and $\vec{v}_M = \mathbf{v}_M \vec{1}$.

If the states are visible, the exercise boundary must be obtained independently for each state and therefore the time step $m - 1$ prices are

$$[\vec{v}_{m-1}]^{(k)} = \max \left\{ [\vec{v}_{m-1}^*]^{(k)}, \mathbf{v}_M \right\} \quad (4.8)$$

for each $k \in \mathbb{K}$. This vector of prices then propagates backwards to time 0 via (4.7) and (4.8). Notice that although the exercise boundaries are determined individually for each state, all prices feed into the conditional price at the previous time step and consequently all boundaries affect the prices, albeit indirectly.

If the states are hidden, the exercise boundary is the same for every state. Furthermore, the distribution of the hidden states at the current time step given the initial distribution $\vec{p}$

¹Alternatively, the probabilities may be estimated using extended Kalman filter or particle filter methods.

Figure 4.1: Price (left) and exercise boundary (right) of the American option *AMR-D* under the regime-switching Black-Scholes-Merton model *BSM-B*. The curves in the stationary model are indicated by markers.

must be obtained and is based on investor beliefs or estimated through a filtering method. Since the states are driven by the generator matrix $\mathbf{A}$ the probability distribution is $\exp\{(m-1)\Delta t\mathbf{A}\}\vec{p}$. Consequently, the (single) unconditional holding price $\mathbf{v}_{m-1}^*$ for the hidden-state regime-switching model is

$$\mathbf{v}_{m-1}^* = \exp\{(m-1)\Delta t\mathbf{A}\}\vec{p} \cdot \vec{\mathbf{v}}_{m-1}^*. \quad (4.9)$$

The optimal exercise boundary is determined from the above price by comparing it to the intrinsic value. The unconditional price at time step $m-1$ is then

$$\mathbf{v}_{m-1} = \max\{\mathbf{v}_{m-1}^*, \mathbf{v}_M\}. \quad (4.10)$$

This is the price in all states at time step $m-1$, thus the vector of prices which propagates back to time 0 via (4.7), (4.9), and (4.10) is $\mathbf{v}_{m-1}\vec{1}$.

Example 1: American options

Figures 4.1 depicts the prices and exercise boundaries of the American option *AMR-D* under the BSM model *BSM-B* with stationary and regime-switching volatility. The regime-switching parameters are as follows: volatility states $\vec{\sigma} = 0.15, 0.25, 0.4$, initial probabilities $\vec{p} = 0.25, 0.5, 0.25$, and transition matrix $\mathbf{A}(0.75)$. In the stationary model, as expected, the price of the option is larger while the exercise boundary is lower for high volatility states as compared to the low volatility states. In the visible-state regime-switching model, the price and exercise boundary curves are pulled towards the mean since the ‘stochasticity’ of the volatility term averages out the overall contribution of various volatility levels. In the hidden-state

Figure 4.2: Implied volatility smiles in regime-switching vs. stationary volatility Merton jump-diffusion models. The curves in the stationary models are indicated by markers.

regime-switching model, the single option price and exercise boundary are computed as described above. Note that the prices and the exercise boundaries can erroneously be thought of as the visible-state prices and exercise boundaries weighted by the transition probabilities — they are not and not defined as such.

Example 2: Implied volatility surfaces

As mentioned in Section 1.1, regime-switching exponential Lévy models have numerous advantages for modeling of equity and currency processes. Specifically, incorporating stochastic volatility into the equity spot price model allows the long-term smile of the IV surface evident in the markets to be matched. In currency markets, the skew of the IV surface is non-stationary, motivating the development of stochastic skew models. This example demonstrates the advantages of using regime-switching volatility and/or skew models, as compared to stationary ones, in fitting of IV surfaces.

Figure 4.2 depicts the IV smiles for the 4-state regime-switching and stationary volatility models at various terms to maturity ($\tau = 0.25, 0.5, 1.0$) obtained using the Merton jump-diffusion model *MJD-D*. The regime-switching parameters are as follows: volatility states $\vec{\sigma} = 0.1, 0.2, 0.3, 0.5$, initial probabilities $\vec{p} = 0.25, 0.45, 0.2, 0.1$, and transition matrix $\mathbf{A}(1)$. In the stationary case, the volatility smile flattens out for the terms to maturity greater than 3 months. Introducing the regime-switching volatility into the models allows significantly more

Figure 4.3: Implied volatility smiles in regime-switching vs. stationary volatility and skew Variance Gamma models. The curves in the stationary models are indicated by markers.

pronounced smiles that persist over all terms to maturity to be obtained. Matching such smiles with stationary models would require an extremely high jump activity rate. Moreover, stationary jump models have difficulties generating IV smiles that do not flatten out for longer terms to maturity.

Figure 4.3 depicts the IV smiles for the 3-state regime-switching and stationary volatility and skew models at various terms to maturity ($\tau = 0.25, 0.5, 1.0$) obtained using the Variance Gamma model *VG-B*. To introduce skew into the IV curves, the stationary model is modified to have the skew $\theta = -0.2$ while in the regime-switching model the parameters are as follows: skew states $\vec{\theta} = -0.3, 0.0, 0.3$, volatility states $\vec{\sigma} = 0.225, 0.125, 0.225$, initial probabilities $\vec{p} = 0.7, 0.2, 0.1$, and transition matrix $\mathbf{A}(1.25)$.

In the stationary case, the volatility smile flattens out for the terms to maturity greater than 3 months, although not as fast as in the jump-diffusion model described above. Introducing the regime-switching volatility into the models allows more pronounced smiles to be obtained while incorporating the desired stylistic feature of the market, i.e., stochastic skew.

4.3 Codependent Jumps via Copulas

Copulas provide a convenient tool to model codependent asset prices or market factors. For example, a major oil pipe disruption during the winter months has a similar upwards effect on both the price of heating oil and natural gas, or revelation of financial difficulties of an automo-

bile company can adversely affect its share price as well as the share prices of its competitors. Capturing such dependency is critical for accurate pricing of multi-asset options with codependent underlyings. This section extends the FST framework methods to handle codependent jumps with the motivation of pricing multi-asset options in currency and commodity markets.

A brief overview of copulas is given below². Copulas formulate a multivariate distribution by transforming the marginal distributions into uniform distributions and imposing a dependence structure on the uniform random variables. As such, a copula is a joint distribution function of standard uniform random variables. The theorem of Sklar (1959) states that if $F(x_1, x_2)$ is a joint cumulative distribution with marginal cumulative distributions $F_1(x_1), F_2(x_2)$ that are continuous, then there exists a copula $C(u, v)$ such that

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)). \quad (4.11)$$

Essentially, copulas allow the marginal behavior of $F_1(x_1)$ and $F_2(x_2)$ to be factored out from the overall dependence structure $F(x_1, x_2)$. A wide array of copulas have been proposed in the literature and Table 4.1 presents joint cdf's and pdf's for some of the most commonly used copulas. The dependence structure of a copula can be visualized by drawing multiple multivariate random samples from its distribution. Figure 4.4 depicts such dependence.

Differentiating the joint cdf $F(x_1, x_2)$ with respect to x_1 and x_2 one obtains the joint pdf $f(x_1, x_2)$:

$$\begin{aligned} f(x_1, x_2) &= \partial_{x_1 x_2} F(x_1, x_2) \\ &= \partial_{x_1} F_1(x_1) \partial_{x_2} F_2(x_2) \partial_{u_1 u_2} C(F_1(x_1), F_2(x_2)) \\ &= f_1(x_1) f_2(x_2) c(F_1(x_1), F_2(x_2)) \end{aligned}$$

where $f_1(x_1), f_2(x_2)$ are the marginal pdf's and $c(u_1, u_2)$ is the copula density. The characteristic function of a compound Poisson process driven by codependent jumps with density f can be readily computed from the Lévy-Khintchine representation:

$$\begin{aligned} \psi_c(\omega_1, \omega_2) &= \lambda_c \left[\int_{\mathbb{R}^2} e^{i(\omega_1 x_1 + \omega_2 x_2)} f(x_1, x_2) dx_1 dx_2 - 1 \right] \\ &= \lambda_c \left[\int_{\mathbb{R}^2} e^{i(\omega_1 x_1 + \omega_2 x_2)} f_1(x_1) f_2(x_2) c(F_1(x_1), F_2(x_2)) dx_1 dx_2 - 1 \right] \end{aligned} \quad (4.12)$$

As previously mentioned in Section 3.4, for the special case of a Gaussian copula and Gaussian marginal jumps distributions, the codependent jump distribution is a bivariate Gaussian

²See Cherubini, Luciano, and Vecchiato (2004) for a detailed treatment of copulas with applications to finance

Copula	Cumulative Distribution $C(u_1, u_2)$	Probability Density $c(u_1, u_2)$
Gaussian	$\Phi_\rho(\Phi^{-1}(u_1), \Phi^{-1}(u_2))$	$\frac{1}{\sqrt{1-\rho^2}} \exp \left\{ -\frac{1}{2}(u_1^2\rho^2 - 2\rho u_1 u_2 + u_2^2\rho^2)/(1-\rho^2) \right\}$
Student's t	$t_{\rho,n}(t_n^{-1}(u_1), t_n^{-1}(u_2))$	$\rho^{-\frac{1}{2}} \frac{\Gamma(\frac{n+2}{2})\Gamma(\frac{n}{2})}{\Gamma(\frac{n+1}{2})^2} \left(\frac{n+(u_1^2-2\rho u_1 u_2+u_2^2)/(1-\rho^2)}{(n+u_1^2)(n+u_2^2)} \right)^{-(n+2)/2}$
Clayton	$\max \left((u_1^{-\theta} + u_2^{-\theta} - 1)^{-1/\theta}, 0 \right)$	$(1+\theta)(u_1 u_2)^{-\theta-1} (u_1^{-\theta} + u_2^{-\theta} - 1)^{-1/\theta-2}$
Frank	$-\frac{1}{\theta} \ln \left(1 + \frac{(e^{-\theta u_1} - 1)(e^{-\theta u_2} - 1)}{e^{-\theta} - 1} \right)$	$\frac{e^{(1+u_1+u_2)\theta} (e^\theta - 1)\theta}{e^{(u_1+u_2)\theta} - e^\theta (e^{\theta u_1} + e^{\theta u_2} - 1)}$

Table 4.1: Cumulative distribution and probability density functions for various copulas. Φ_ρ and $t_{\rho,n}$ are the joint cumulative distribution functions of the bi-variate standard normal and Student's t distributions, respectively, with correlation ρ .

Figure 4.4: Random samples from four commonly used copulas: Gaussian copula with $\rho = -0.85$, Student's t copula with $\rho = -0.75, n = 2$, Clayton copula with $\theta = 4$, and Frank copula with $\theta = 10$.

and ψ_c can be computed in closed form. In the general case, such as non-Gaussian copula or non-Gaussian jumps, the copula characteristic function (4.12) has to be approximated numerically. Fortunately, such an integral can be computed efficiently using the FFT algorithm. The FST algorithm requires the characteristic exponent to be evaluated on a grid $\hat{\Omega} = [0, \omega_1^{\max}] \times [0, \omega_2^{\max}]$, with $\omega_{1m} = m \cdot \Delta_{\omega_1}$ and $\omega_{2n} = n \cdot \Delta_{\omega_2}$. The integral is approximated via

$$\begin{aligned}
 \int_{\mathbb{R}^2} e^{i(\omega_{1m}x_1 + \omega_{2n}x_2)} f(x_1, x_2) dx_1 dx_2 &\approx \Delta_{x_1} \Delta_{x_2} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} e^{i(\omega_{1m}x_{1k} + \omega_{2n}x_{2l})} f_{kl} \\
 &= \Delta_{x_1} \Delta_{x_2} e^{i(\omega_{1m}x_1^{\min} + \omega_{2n}x_2^{\min})} \sum_{k=0}^{N-1} \sum_{l=0}^{N-1} e^{i(mk + nl)/N} f_{kl} \\
 &= \alpha_{mn} \langle \text{FFT}^{-1} [\mathbf{f}] \rangle_{mn}, \tag{4.13}
 \end{aligned}$$

where $\Omega = [x_1^{\min}, x_1^{\max}] \times [x_2^{\min}, x_2^{\max}]$ is the truncated integration domain with $x_{1k} = x_1^{\min} + k\Delta_{x_1}$, $x_{2l} = x_2^{\min} + l\Delta_{x_2}$; also $f_{kl} = f(x_{1k}, x_{2l})$ and $\alpha_{mn} = \Delta_{x_1} \Delta_{x_2} e^{i(\omega_{1m}x_1^{\min} + \omega_{2n}x_2^{\min})}$. The two grids in real and frequency space are related via $\Delta_{\omega_1} \cdot \Delta_{x_1} = \Delta_{\omega_2} \cdot \Delta_{x_2} = 1/N$. The advantage of the above method is that only one evaluation of the FFT algorithm is required to obtain

Figure 4.5: Approximation error for the characteristic exponent ψ_c of Gaussian copula jumps model.

$\psi_c(\omega_{1m}, \omega_{2n})$ for all n, m .

Figure 4.5 shows the real and imaginary components of the error in evaluating the characteristic exponent ψ_c of copula jumps in the model *KJDMRC*. The integral was computed using $N = 1024$. In practice the same N and $\hat{\Omega}$ should be used to compute the integral as in the pricing FST method, so that the integral is computed on the same frequency grid and no interpolation or extrapolation is required. Note that this FFT-based integration of the copula characteristic function can be used in conjunction with any of the FST framework-based methods.

Although the precision of the aforementioned approach is sufficient for precise pricing of options, the method can be further improved by incorporating high-order quadrature. Press, Teukolsky, Vetterling, and Flannery (1992) outline such extension for the one-dimensional case; the two-dimensional extension is left for future research.

Chapter 5

Exotic Options

5.1 Introduction

This chapter applies the FST framework to price two exotic options — shout options, which offer greater degree of protection compared to vanilla options by allowing the strike price to be reset, and swing options, which offer constrained flexibility in terms of amount and timing of commodity delivery.

Vanilla options (normal options without special features) offer investors a means of protection and/or speculation in volatile markets. For instance, put options offer protection against a decline in a prescribed stock price or index from the level decided upon at inception. Shout options offer a greater degree of flexibility (and thus greater degree of protection), as compared to vanilla options, and allow for the protection level to be set at a later date, possibly several times. The option holder, at each exercise opportunity, must decide whether to set the strike price to the current prevailing spot price or wait for a possibly more opportune moment in the future to possibly reset the strike price. Option prices and optimal exercise policy can be computed by recursing on the number of remaining exercise opportunities.

Swing options, common in commodity markets, allow the holder of the option to modify the amount and timing of the commodity delivered. By modifying the supply today, the investor forgoes an opportunity to modify the supply at a future time, which can potentially have a higher payoff. Naturally, the increased flexibility, as compared to a simple forward agreement, comes at a price. Again, option prices and optimal consumption policy of swing options can be computed via a dynamic programming algorithm.

The outline of the remainder of this chapter is as follows. Section 5.2 introduces shout options and develops a recursive algorithm, based on the FST framework, for their valuation. A numerical example demonstrates the behavior of shout option prices and optimal exercise policies. Section 5.3 introduces swing options and develops a dynamic programming algorithm,

also based on the FST framework, for their valuation. The behavior of swing option prices is examined through a numerical example.

5.2 Shout Options

A shout option is a European option which gives the investor the right to reset the strike level of the option to the prevailing spot price, possibly multiple times. The act of resetting the strike level is called *shouting* and the time of the shout is called the *shout time*. Clearly, the holder of a shout call (put) option shouts only when the prevailing spot price is below (above) the strike price set initially or at a previous shout time. Consequently, the investor receives an at-the-money (strike price is equal to the prevailing spot price) option with one fewer shout opportunities remaining. Hence, the value of a shout option is greater or equal to the value of a vanilla European at-the-money option.

The problem of shout option valuation has been studied by Thomas (1993), Cheuk and Vorst (1997), Windcliff, Forsyth, and Vetzal (2001), Dai, Kwok, and Wu (2004), and Dai and Kwok (2008). The above methods are limited to pricing of shout options under the BSM model. In this chapter the analysis is extended to price shout options under exponential Lévy models and mean-reverting jump-diffusion models. As previously mentioned in Section 1.1, such models account for the unique characteristics of equity and commodity markets, such as volatility smile/skew and reversion to the long-run mean.

Let $V^{(k)}(t, \mathbf{S})$ denote the price of a k -shout option at time t and spot price $\mathbf{S}$. Also, let $\dot{V}^{(k)}(t, \mathbf{S})$ denote the price of a k -shout option at time t and spot price $\mathbf{S}$ where the initial strike level *has* been set to $\mathbf{S}(0)$. Note that the initial strike level for a k -shout option $V^{(k)}$ is *not* set initially. Also, let $E(\mathbf{S}, \mathbf{K}, \tau)$ be the price of a European option with spot price $\mathbf{S}$, strike price $\mathbf{K}$ and time to maturity τ . Below, a recursive relation is derived for the shout option price $V^{(k)}$ using shout options prices $\dot{V}^{(l)}$ for $l = 1, \dots, k - 1$.

A single-shout option is equivalent to a European option except that the strike level is not set initially but rather set to the prevailing spot price at some future shout time. If the investor shouts at maturity, then the shout option is worthless as the shout strike price is equal to the spot price (the option also expires worthless if a shout has not been made). If the investor shouts prior to maturity, then he receives an at-the-money European option. Thus, over a small time interval Δt , the value of a single-shout option $V^{(1)}$ is the larger of its continuation value over the time interval and a European option E struck at the current spot price. Similarly, a single-shout option with initial strike set, $\dot{V}^{(1)}$, has value which is the larger of *its* continuation value and the at-the-money European option E . This can be expressed via the recursive relation for 1-shout options:

$$V^{(1)}(t, \mathbf{S}) = \max \left\{ e^{-r\Delta t} \mathbb{E}_t^{\mathbb{Q}} \left[V^{(1)}(t + \Delta t, \mathbf{S}) \right], E(\mathbf{S}, \mathbf{S}, T - t) \right\}, \quad (5.1a)$$

$$\dot{V}^{(1)}(t, \mathbf{S}) = \max \left\{ e^{-r\Delta t} \mathbb{E}_t^{\mathbb{Q}} \left[\dot{V}^{(1)}(t + \Delta t, \mathbf{S}) \right], E(\mathbf{S}, \mathbf{S}, T - t) \right\}. \quad (5.1b)$$

A multi-shout option gives the investor the right to reset the strike level to the prevailing spot price multiple times. At shout time, the holder of a k -shout option receives a $k-1$ -shout option with the strike level set to the prevailing spot price. Thus, over a small time interval Δt , the value of a k -shout option $V^{(k)}$ is the larger of its continuation value over the time interval and a $k-1$ -shout option $\dot{V}^{(k-1)}$ with strike price being the prevailing spot price at the shout time. Similarly, a k -shout option with initial strike set $\dot{V}^{(k)}$ has value which is the larger of its continuation value and a $k-1$ -shout option $\dot{V}^{(k)}$ with strike price being set to the prevailing spot price at the shout time. Thus, the following recursive relations for k -shout options is obtained:

$$V^{(k)}(t, \mathbf{S}) = \max \left\{ e^{-r\Delta t} \mathbb{E}_t^{\mathbb{Q}} \left[V^{(k)}(t + \Delta t, \mathbf{S}) \right], \dot{V}^{(k-1)}(t, \mathbf{S}(0)) \cdot \mathbf{S}/\mathbf{S}(0) \right\}, \quad (5.2a)$$

$$\dot{V}^{(k)}(t, \mathbf{S}) = \max \left\{ e^{-r\Delta t} \mathbb{E}_t^{\mathbb{Q}} \left[\dot{V}^{(k)}(t + \Delta t, \mathbf{S}) \right], \dot{V}^{(k-1)}(t, \mathbf{S}(0)) \cdot \mathbf{S}/\mathbf{S}(0) \right\}. \quad (5.2b)$$

Above, one European at-the-money option with spot and strike price $\mathbf{S}$ is equivalent to $\mathbf{S}/\mathbf{S}(0)$ European at-the-money options with spot and strike price $\mathbf{S}(0)$, i.e. $\dot{V}^{(*)}(t, \mathbf{S}(0)) \cdot \mathbf{S}/\mathbf{S}(0)$. This similarity reduction is required as $\dot{V}^{(*)}$ denotes the price of a shout option with the strike price set to $\mathbf{S}(0)$ at time $t = 0$.

At maturity, a shout option without initial strike being set expires worthless and a shout option with initial strike set has the payoff at expiry of a European option. Thus, the terminal conditions for $l = 1, \dots, k$ are

$$V^{(l)}(T, \mathbf{S}) = 0, \quad (5.3a)$$

$$\dot{V}^{(l)}(T, \mathbf{S}) = E(\mathbf{S}, \mathbf{S}(0), 0). \quad (5.3b)$$

Armed with these recursive relations for the k -shout option, the numerical algorithm is quite straightforward. At each time step, the $2k$ vectors of option values are propagated backwards in time (using the appropriate FST method) to compute the necessary expectations. Then, equations (5.1) and (5.2) are applied recursively to yield $1, \dots, k$ -shout option prices, with and without initial strike set.

Figure 5.1 depicts the prices and exercise boundaries of the multi-shout option SHT under the Variance Gamma model $VG-A$. One immediately notices that the value of k -shout option $V^{(k)}$, indicated by curves without markers, increases linearly with the stock price, i.e., $V^{(k)}(\cdot, \alpha \mathbf{S}) = \alpha V^{(k)}(\cdot, \mathbf{S})$. Since the initial strike is not set, the value of the option depends

Figure 5.1: Prices (left) and exercise boundaries (right) of the multi-shout options *SHT* under the Variance Gamma model *VG-A*. Curves for the option with initial strike set are indicated by markers.

only on the spot price $\mathbf{S}$. Similarly to the price of a vanilla European at-the-money option, the linear scaling property is necessary to avoid arbitrage. On the other hand, the value of a k -shout option with initial strike set $\dot{V}^{(k)}$, indicated by curves with markers, depends on the spot price in non-linear fashion. Obviously, shout options with the initial strike price set should command a premium, hence $V^{(k)}(\cdot, \mathbf{S}) \leq \dot{V}^{(k)}(\cdot, \mathbf{S})$. Alternatively, this can easily be shown from the recursive equations above. The size of the premium $\dot{V}^{(k)}(\cdot, \mathbf{S}) - V^{(k)}(\cdot, \mathbf{S})$ is approximately equal to the value of the vanilla European option $E(\mathbf{S}, \mathbf{S}(0), T)$.

The optimal exercise strategy for a k -shout option with initial strike set is as follows: once the spot price crosses the exercise boundary, the strike price is set to the prevailing spot price and the option has $k - 1$ shouts remaining. For the k -shout option (without the initial strike set), the optimal exercise boundary is a vertical line and the exercise time does not depend on the prevailing spot price. In the example considered in Figure 5.1, 5-, 4- and 3-shout options must be shouted immediately while for the 2- and 1-shout options the optimal shout times are $t \approx 3.5$, marked by the vertical dashed non-marked line, and $t \approx 8.5$, marked by the vertical solid non-marked line, respectively. Once the first shout has been made, the optimal shout time is determined from the boundaries of shout options with initial strike, which are denoted by marked curves.

Figure 5.2 depicts the optimal shout times of the single-shout option *SHT* for a range of skews (left) and volatilities (right). The Variance Gamma model *VG-A* is taken as the base case. Higher kurtosis of the spot price process corresponds to earlier optimal shout times. Similar effect is achieved by increasing the magnitude of the skew in either direction. The excess kurtosis introduced by the stochastic time change along with the highly positive or negative

Figure 5.2: Optimal shout times of a single-shout option *SHT* for a range of model parameters. The Variance Gamma model *VG-A* is taken as the base case.

skewness implies that the probability that the option ends in the money relative to the current spot price is higher than in the case of low kurtosis and low skewness. Hence, it is advantageous to shout earlier. Interestingly, shout option with highly positive skew have much earlier optimal shout time than option with highly negative skew. The effect of volatility on the optimal shout time is more involved. For low volatilities ($\sigma < 0.15$) high kurtosis leads to earlier shout times while for high volatilities ($\sigma > 0.25$) the reverse is true (note that the skewness in this case is $\gamma = -0.28113$). For low volatilities this behavior is consistent with the plot of optimal shout times for various skews, where similar observations are made for extremely negative skews. For high volatilities, the reversal in optimal shout behavior may be attributed to the fact that high kurtosis along with highly negative skew imply that the probability that the option ends in the money is lower as compared to the low kurtosis case.

5.3 Swing Options

Swing options provide constrained flexibility with respect to volume and timing of energy delivery. As an example, consider the following contract: The holder of the option agrees to purchase 100MW of electricity per hour at a cost of \$45/MWh over a period of 1 month. At the start of each hour, the holder has the right to increase power consumption (*swing up*) to 110MW for that hour or decrease to 90MW (*swing down*) at the same price. The total number of such changes is limited to 50. There are two essential components to the swing option: a pure forward agreement to deliver a fixed amount of energy over a period of time and the variational or *swing* component, which is the right to change consumption at the option holder's choosing.

The problem of pricing swing options has been studied extensively in the literature¹. The early work of Pilipović and Wengler (1998) discusses special cases that can be solved with straightforward procedures. Jaillet, Ronn, and Tompaidis (2004) develop a binomial forest methodology for pricing of swing options under a one-factor mean-reverting stochastic process for energy prices with seasonal effects. Ibáñez (2004) prices swing options using the Monte-Carlo method combined with the dynamic programming approach while Figueroa (2006) extends the least-squares Monte-Carlo approach for pricing of American options to pricing of swing options. Kluge (2006) extends the binomial forest approach by pricing on a finite grid. Several books on commodity markets, such as Pilipović (1997), Clewlow and Strickland (2000) and Eydeland and Wolyniec (2003) discuss swing options as well.

Let $V(t, \mathbf{S}, Q)$ denote the value of the swing option at time t , spot price $\mathbf{S}$, having exercised the right to change consumption Q times prior to t . At each swing opportunity, the holder has a choice to change their consumption q times (or by q units). For example, when commodity prices are high, the investor may increase the supply from the option seller and sell the excess supply into the market. On the other hand, when commodity prices are low, the investor may decrease the supply from the option seller and buy the remainder of the supply in the market (for a lower price than the one charged by the option seller). In summary, the choices available to the option holder are: do nothing ($q = 0$), increase consumption, or swing up ($q > 0$), and decrease consumption, or swing down ($q < 0$). The cash-flow function $\Phi(\mathbf{S}, q)$ captures the immediate monetary benefit of such a change in consumption. The cash flow function may also include a penalty term to enforce additional limits on Q or may be as simple as the value of selling the extra supply into the market: $\Phi(t, \mathbf{S}, q) = q(\mathbf{S} - K)$.

The value of the swing option can then be expressed as a solution to the following dynamic programming equation:

$$V(t, \mathbf{S}, Q) = \max_q \left\{ \Phi(t, \mathbf{S}, q) + e^{-r\Delta t} \mathbb{E}_t^Q [V(t + \Delta t, \mathbf{S}, Q + q)] \right\}, \quad (5.4)$$

where $Q(t)$ is the total number of consumption changes prior to time t and is defined as either the total number of swings $Q(t) = \sum_{t_k \leq t} |q(t_k)|$ or the difference between the up-swings and down-swings $Q(t) = \sum_{t_k \leq t} q(t_k)$. Nonetheless, the total number of swings is bounded $Q^{\min} \leq Q(t) \leq Q^{\max}$ and both choices can be handled by the dynamic programming equation above.

While the structure of the dynamic programming equation may be slightly different for more exotic types of swing contracts, the critical part of the algorithm is the computation of the expectation in equation (5.4). Depending on the spot price model, the FST framework-based methods, outlined in the previous chapters, can be used to carry out that essential computation.

¹See Ware (2005) for an overview of swing option valuation

Figure 5.3: Effect of mean-reversion level $\theta = \ln 75$ (left) and $\theta = \ln 125$ (right), and various mean-reversion speeds on the value of the swing option $SWNG$ under the mean-reverting Merton jump-diffusion model $MJDMR$.

By using such methods, one can easily incorporate jumps, and, if two commodities are tied together in the swing contract, codependence. Furthermore, there is a huge efficiency gain over traditional Monte Carlo methods or multi-node forest methods (stacks of trees). For the FST methods (as with some other methods, such as finite difference methods) one must keep a stack of option values $V(\cdot, \cdot, Q)$ for different levels of total consumption. These prices are then propagated back in time using the FST/mrFST algorithm. Furthermore, forward and backward FFT transforms can be applied efficiently to the entire stack by utilizing the multi-data FFT transform available in most FFT packages (multi-core architectures can be especially effective for such applications).

Figure 5.3 depicts the effect of mean-reversion levels $\theta = \ln 75$ and $\theta = \ln 125$, and various mean-reversion speeds on the value of the swing option $SWNG$ under the mean-reverting Merton jump-diffusion model $MJDMR$. Under the parameters chosen, the swing option should be especially valuable when the spot price moves away significantly from the strike price in either direction and this is manifested in the U-shaped value curve. For low spot prices, all swing option values converge to the $-3(S - K)$ line — it is optimal to swing down 3 times immediately, profiting from the ability to buy the commodity cheaply. For low mean-reversion level, high mean-reversion rate makes the swing option less valuable — there is less volatility in the spot price and the probability of the spot price reaching very high or very low levels is small. Similar behavior is observed for high mean-reversion level at high spot prices. On the other hand, for high mean-reversion level and low spot prices, high mean-reversion rate makes the swing option more valuable as there is higher probability the price of the commodity will rise.

Chapter 6

Graphics Processing Units

6.1 Introduction

Over the past decade, the computing needs in a typical financial institution setting have increased significantly. Pricing, calibration and risk management tasks must now be delegated to dedicated computational servers to be completed in a sufficiently short time. This increase in computing demands has been brought about by the tremendous growth in both the scope and size of problems being addressed and the complexity of models being used. Increasingly, jump-diffusion and exponential Lévy models are used instead of the classical BSM model to correct for the observed implied volatility (or skew) and term structure. The ever-increasing complexity and scope of problems being tackled in the area of computational finance brings about a need for efficient pricing architectures that are powerful and flexible. This chapter discusses the application of graphics cards to option pricing and shows that they can provide a significant increase in performance over standard CPUs when pricing path-dependent, single- and multi-asset options.

CPU clusters have been successfully used by Gerbessiotis (2004), Sak, Özekici, and Boduroğlu (2007) and others in parallel computation of options under complex models. However, over the last several years, graphics processing units (GPUs) have evolved from mere dedicated graphics rendering devices to computing ‘workhorses’. The fact that GPUs are designed with data processing in mind, rather than data caching and flow control, together with their highly parallel, throughput oriented structure and focus on individual thread performance, makes them more effective than typical CPUs in compute-intensive and highly parallel applications. The literature on utilizing GPUs in option pricing is quite sparse and limited to the implementation of the BSM formula, Monte Carlo simulation by Podlozhnyuk (2007) and the binomial lattice pricing method by Kolb and Phar (2005). This chapter shows that GPUs can also be effectively leveraged for pricing exotic, path-dependent and multi-dimensional options, where the under-

lying stock price indices are modeled using jump-diffusion and exponential Lévy processes by accelerating the FFT computation in the FST framework-based methods. Additionally, the parallel structure of GPUs is leveraged to price multiple options simultaneously to further increase computational throughput of FST methods.

The outline of the remainder of this chapter is as follows. Section 6.2 shows that GPUs provide a highly efficient alternative to CPUs for computing FFTs. Section 6.3 develops numerical algorithms for utilizing GPUs to price single- and multi-asset options using the FST method. The computational speedup attained by using GPUs is illustrated through pricing experiments. Lastly, section 6.4 shows how the parallel architecture of a GPU can be further utilized to price multiple options concurrently. Parts of this chapter have been published in Surkov (2008)¹.

6.2 FFT Computation on GPUs

GPUs were originally developed to perform several graphics primitive operations, such as texture mapping and polygon rendering. However, over the past several years the functionality of such cards increased tremendously to allow for their use in general scientific and business computing. GPUs have evolved into cheap, powerful and highly parallel processing units that rival traditional CPUs in computationally intensive applications.

Figure 6.1 shows the performance of high-end CPUs, GPUs and Cell BE, with the reported values being the peak theoretical FLOPS (floating point operations per second). The results reported are from manufacturers' specifications and, for the GPUs, FLOPS refers to the number of floating point operations that can be performed by shader cores. The GPUs denoted by $\star$ are dual systems (two GPUs located on a single card) and the performance is reported for a single GPU. Note that comparing different architectures using FLOPS as a benchmark is quite tenuous. First, many operations in GPUs are not performed on the shader cores which makes defining FLOPS consistently very difficult. Second, the reported numbers are for *peak theoretical* rather than *sustained* throughput, the latter being more relevant for large scale scientific computations. Lastly, differences in power consumption among the different architectures (GPUs typically having a significantly higher power consumption as compared to the CPUs and Cell BE) can skew the results of the comparison significantly. The purpose of Figure 6.1 is to highlight that GPUs have significantly evolved and offer an attractive architecture for carrying out intensive scientific computations in either standalone manner or as a coprocessor to the CPU. While early GPUs were at a disadvantage relative to CPUs, due to limited available memory, the current generation of GPUs has a comparable memory capacity. With the typical cards supporting around 1GB of memory, a GPU can be used to address high-dimensional

¹© 2008 IEEE

Figure 6.1: Peak theoretical performance of various high-end Central and Graphics Processing Units, and Cell Broadband Engine for single-precision computing

problems in the same manner as a workstation computer.

A significant bottleneck in utilizing GPUs for any type of computing is the transfer of data to and from the card. Thus, it is of paramount importance to reduce data traffic when designing numerical algorithms that utilize GPUs. In this section, to assess the performance of a GPU in pricing options with the FST method, computational times required to perform FFTs of various sizes and dimensions on a CPU and a GPU are compared. Also, the total round-trip time to compute an FFT, which includes the data transfer time, is measured. The experiments were conducted on a NVIDIA GeForce 9800 GX2 video card with 1GB of memory, running on a workstation with an Intel Core 2 Duo E7200 2.53GHz CPU and 4GB of RAM. The FFTW library of Frigo and Johnson (2005), which provides a flexible C interface and is one of the fastest FFT algorithm implementations currently available, was used to execute FFTs on the CPU. The NVIDIA CUFFT library provides an interface modeled after FFTW and was used to execute FFTs on the GPU.

Table 6.1 summarizes the timing results for executing one- and two-dimensional FFTs of various sizes on the CPU and the GPU. ‘CPU time’ measures the computational time for a combination of forward and backward, out-of-place, complex-to-complex FFTs on the CPU. ‘GPU time’ measures the time to perform the same combination of FFTs on the GPU where the data is *not* moved to or from the device. ‘GPU round-trip time’ measures the same combination of FFTs but with data uploaded to the device before and downloaded from the device after the FFT evaluation. Note that, while NVIDIA GeForce 9800 GX2 video card is a dual card, only one GPU was used for the computation.

Transform size	CPU time (msec.)	GPU time round-trip (msec.)	GPU time (msec.)
4096	0.11	0.21	0.11
8192	0.33	0.28	0.14
16384	0.65	0.37	0.18
32768	1.33	0.66	0.25
512 ²	14.0	4.09	0.94
1024 ²	95.7	15.4	4.08
2048 ²	453	69.5	26.7

Table 6.1: Fast Fourier Transform execution performance on the Intel Core 2 Duo E7200 2.53 GHz CPU and the NVIDIA GeForce 9800 GX2 GPU. Only one core for the CPU and one card for the GPU are utilized.

As evident from the results presented in Table 6.1, the GPU is more efficient than the CPU at evaluating FFTs for all sizes considered. As CPUs are optimized for latency and GPUs are optimized for high throughput, the computational times for small one-dimensional transforms on the CPU and the GPU are comparable. However, for two-dimensional and large one-dimensional transforms, the GPUs are significantly faster. The GPU achieves a speedup factor of approximately 5 for one-dimensional transforms and 17 for two-dimensional transforms. If data transfer is taken into account, the advantage of GPUs is reduced by a factor of 2 for one-dimensional transforms and 3 for two-dimensional transforms.

Note that although the results obtained are quite impressive, current state-of-the-art GPUs, such as the NVIDIA GeForce GTX 200 series cards and the ATI Radeon 4800 series cards, have become available on the market and are capable of even faster computations. Further advances in the performance of GPU architectures will result in their improved performance and bigger advantage compared to corresponding CPU-based methods.

6.3 Applications to Option Pricing

In this section, the FST method for pricing European and American options, referred to as FST-GPU, is discussed. In addition, results for timing tests are presented to compare the efficiency of the FST-GPU method and FST method on a CPU, referred to as FST-CPU. As illustrated by the results of the previous section, memory transfer is a critical issue when designing the option pricing algorithms for GPUs. From the results of timing tests one would expect the FST-GPU to be marginally more efficient than the FST-CPU for pricing of standard European options (where typically only 8192 space points for single-asset problems and 2048²

space points for two-asset problems are required to achieve accuracy of 1/10 of a cent) since a full memory round-trip is required for only two FFT evaluations. For American options, on the other hand, one can expect a greater efficiency gain for the FST-GPU method as it does not require a memory round-trip between every time step. The degree of efficiency also depends on the length of the FFT evaluation as a share of the overall computational time.

Algorithm 1: FST-GPU algorithm for pricing European options.

Input: Option payoff $\mathbf{v}_1$, characteristic exponent Ψ

Output: Option values $\mathbf{v}_0$

Upload $\mathbf{v}_1$, $e^{\Psi \Delta t}$ to GPU

$\mathbf{v}_0 \leftarrow \text{FFT}^{-1} [\text{FFT} [\mathbf{v}_1] \cdot e^{\Psi \Delta t}]$

Download $\mathbf{v}_0$ from GPU

return $\mathbf{v}_0$

The FST-GPU algorithm for pricing of European options is outlined in Algorithm 1 and is naturally derived from equation (2.12). For performing pricing with N space points, the algorithm must upload N floating point values for the option payoff and $N/2 + 1$ complex floating point values for the characteristic factor $e^{\Psi \Delta t}$ (since option values are real, half the complex values are redundant due to Hermitian symmetry) and download N floating point values for $\mathbf{v}_0$ to the host. If the option value is required only at a specific spot price then only one floating-point value has to be downloaded. In addition to the memory transfer, one forward and one inverse FFT evaluation are required.

In the two-asset case, option payoff $\mathbf{v}_1$ constitutes a matrix of values and Ψ is the corresponding characteristic exponent matrix with the same dimensions. Similarly, $\text{FFT}[\cdot]$ and $\text{FFT}^{-1}[\cdot]$ refer to the two-dimensional forward and inverse FFT algorithms, respectively. For pricing with $N \times N$ space points, the algorithm must upload N^2 floating point values for the option payoff and $N \cdot (N/2 + 1)$ complex floating point values for the characteristic factor $e^{\Psi \Delta t}$ (again, due to Hermitian symmetry). Also, N^2 floating point values are downloaded to the host (only one floating point value may be downloaded if the entire price surface is not needed). As in the single-asset case, pricing of European options requires the execution of one forward and one inverse two-dimensional FFT.

All computations in this section were done in single precision, as opposed to double precision in the rest of the thesis. Also, the timing results for pricing multi-asset options with FST-GPU method on grid sizes larger and including 4096^2 are not available. In the numerical experiments 2-dimensional transforms of such sizes would not fit into memory and cause program crashes².

²The excessive memory usage of CUFFT library has been reported by several developers. See for instance <http://forums.nvidia.com/index.php?showtopic=38931>.

N	Value	Change	$\log_2$ Ratio	CPU Time (msec.)	GPU Time (msec.)
2048	7.28155746			1.167	1.178
4096	7.27979297	0.0017645		1.734	1.932
8192	7.28005799	0.0002650	2.7351	3.353	3.501
16384	7.28011385	0.0000559	2.2463	6.601	6.519
32768	7.28012262	0.0000088	2.6710	13.234	12.642

Table 6.2: Pricing results for the European option *EUR-B* under the Kou jump-diffusion model *KJD-A*. The reference price 7.27993383 is computed using the Fourier quadrature method. The order of convergence is 2 in space.

N	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
512 ²	1.92890266			0.187	0.191
1024 ²	1.92652784	0.0023748		0.749	0.720
2048 ²	1.92550786	0.0010200	1.2193	2.972	2.816
4096 ²	1.92500700	0.0005009	1.0260	12.123	N/A
8192 ²	1.92477518	0.0002318	1.1115	50.010	N/A

Table 6.3: Pricing results for the European catastrophe equity put option *ECEP* under the joint stock-loss model *JSL*. The order of convergence is 1 in space.

Example 1: European options

To test the performance of the FST-GPU algorithm in the single-asset case, the European option *EUR-B* under the Kou jump-diffusion model *KJD-A* and the European option *EUR-D* under the CGMY model *CGMY-B* are priced. The convergence and timing results are given in Table 6.2 and Table C.15 in Appendix C.5. To test the performance of the FST-GPU method for two-asset path-independent options, the European CatEPut option *ECEP* under the joint stock-loss model *JSL* and the European spread option *ESPD* under the 2D BSM model *BSM-C* were priced. The convergence and timing results are presented in Table 6.3 and Table C.17 in Appendix C.5.

The timing results for pricing single-asset European options in Table 6.2 and Table C.15 in Appendix C.5 suggest that a GPU offers no significant advantage over a CPU in pricing of path-independent options. The result is directly linked to the fact that the evaluation of the characteristic function constitutes a significant share of the overall work. In these experiments, the computation is performed on the CPU for both methods (so that it can be carried out in

double precision), rendering the advantage of FST-GPU insignificant. Delegating the evaluation of the characteristic function to a GPU should allow the FST-GPU method to achieve a computational speedup of approximately 5, as demonstrated by the FFT computation results presented in Table 6.1.

Similar to the results for the one-dimensional European case, the FST-GPU and the FST-CPU methods produce comparable results in the two-asset case as demonstrated by the timing results in Table 6.3 and Table C.17 in Appendix C.5. Due to the fixed overhead associated with each memory transfer, transforms of large size, from computational point of view, are relatively more efficient than small transforms. Thus, as opposed to the single-asset case, the large size of the problem has increased the advantage of the multi-asset FST-GPU over the FST-CPU, albeit by a small margin.

Algorithm 2: FST-GPU algorithm for pricing American options.

Input: Option payoff $\mathbf{v}_M$, characteristic exponent Ψ

Output: Option values $\mathbf{v}_0$

Upload $\mathbf{v}_M, e^{\Psi \Delta t}$ to GPU

for $n \leftarrow M$ **to** 1 **do**

$\tilde{\mathbf{v}}_n \leftarrow \text{FFT}^{-1} [\text{FFT} [\mathbf{v}_n] \cdot e^{\Psi \Delta t}]$

$\mathbf{v}_{n-1} = \max\{\tilde{\mathbf{v}}_n, \mathbf{v}_M\}$

end

Download $\mathbf{v}_0$ from GPU

return $\mathbf{v}_0$

The FST-GPU algorithm for pricing American options extends Algorithm 1 by incorporating the time-stepping equation (2.22) and is given in Algorithm 2. When M time steps are used, M forward and inverse FFTs of size N are executed and $M \cdot N$ evaluations of the max function are required. Yet, the algorithm requires the same amount of memory transfer as in the European case. Thus, as M increases, the evaluation of the payoff and characteristic functions and the memory transfer overhead become a less significant factor in the performance of FST-GPU.

Example 2: American options

To test the performance of the FST-GPU method for American options, the American option *AMR-A* under the Merton jump-diffusion model *VG-B* and the American option *AMR-B* under the Variance Gamma model *MJD-A* are priced. The convergence and timing results are given in Tables 6.4 and C.16, respectively. As examples of the multi-asset path-dependent options, the American double-trigger stop-loss option *ADTSL* under the joint stock-loss model *JSL* and the American spread option *ASPD* under the 2D BSM model *BSM-C* were priced. The convergence

N	M	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
2048	128	8.01846275			0.009	0.017
4096	512	8.01147970	0.0069831		0.077	0.075
8192	2048	8.01337394	0.0018942	1.8822	0.656	0.343
16384	8192	8.01402855	0.0006546	1.5329	4.711	1.720
32768	32768	8.01391362	0.0001149	2.5098	47.216	10.601

Table 6.4: Pricing results for the American option *AMR-A* under the Variance Gamma model *VG-B*. The order of convergence is 2 in space and 1 in time.

N	M	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
512 ²	64	2.53730898			1.097	0.258
1024 ²	256	2.67880465	0.1414957		20.087	1.746
2048 ²	1024	2.74424933	0.0654447	1.1124	326.903	31.246
4096 ²	4096	2.77366599	0.0294167	1.1536	6539.073	N/A

Table 6.5: Pricing results for the American double-trigger stop-loss option *ADTSL* under the joint stock-loss model *JSL*. The order of convergence is 1 in space and 1/2 in time.

and timing results for the two test cases are presented in Tables 6.5 and C.18 respectively.

As expected, the FST-GPU method outperforms the FST-CPU method for larger problems due to the substantial decrease in the fraction of the overall computational time taken by the computation of the payoff and characteristic functions and memory transfer. For single-asset American options, the FST-GPU method is nearly 5 times faster for the largest problem tested — almost the same speedup as the one attained for the pure FFT evaluation. For two-asset options, the FST-GPU method outperforms FST-CPU method by a factor of 10 for the largest problem tested. This is significantly less than the speedup of 17 attained by the pure two-dimensional FFT evaluation and may be attributed to the increased use of shared memory by the GPU on large-size computations.

6.4 Applications to Parallel Option Pricing

The inherently parallel, multi-processor structure of the GPUs makes them especially attractive for the parallel solution of multiple problems. Parallelizing multiple computations allows to streamline the memory transfer and saturate the processor cores with work, thus increasing the overall throughput of the algorithm. As this section shows, multiple single-asset options can be

Figure 6.2: Timing results for batched Fast Fourier Transform computation.

priced efficiently by performing the valuations concurrently.

The CUFFT library provides a convenient interface for evaluating multiple one-dimensional FFTs concurrently and thus fully exploiting the parallel architecture of the GPU. Results in Figure 6.2 show that by batching the FFTs together, one can achieve a higher throughput on the GPU. For FFTs of size less than 4096 points, the GPU achieves an increase in throughput of 100% when the batch size is increased from 4 to 64. However, the marginal return of parallelizing the FFTs diminishes as the size of the FFT increases. In fact, for FFTs of size greater or equal to 16384 points there is no improvement in throughput across the various batch sizes. Given the large size of the FFT and the relatively small number of processors on the NVIDIA GeForce 9800 GX2, there is little idling of the processors and thus little benefit to parallelization. More powerful GPUs with more processors should benefit from parallelization even at such large transform sizes.

The increase in FFT throughput due to batching translates into faster pricing of multiple

Batch size	GPU time (msec.)	Options/sec.
1	3.51	284
4	6.12	653
16	18.39	870
64	66.93	956
256	261.52	979

Table 6.6: Timing results for parallel pricing of European options.

options, as shown in Table 6.6. In this example, European options are priced with 8192 space points and various batch sizes, ranging from 1 to 256. Parallelization of the FST-GPU method increases the throughput of the algorithm from 284 options per second to 979 options per second — an increase of almost 350%. While most of the performance gains came from parallelization of FFTs, streamlining of memory data transfer contributed to the improvement as well. For path-dependent options, where the computation of FFTs takes up the majority of the computational time, the performance gain due to parallelization of the FST-GPU method would be even higher.

Chapter 7

Conclusions

7.1 Summary of Research

This thesis develops a Fourier transform-based framework, called Fourier Space Time-stepping, for computing the evolution of option prices in time. The framework allows pricing of various options, including single- and multi-asset European, American, and barrier options, under various asset price models, including independent-increment, mean-reverting and regime-switching exponential Lévy models. The framework is generic in the sense that various options under different spot price processes can be priced by supplying the appropriate payoff function and characteristic exponent without further modifications to the numerical algorithm. The following summarizes the major features of the methods based on this framework.

- **Precision, Speed and Convergence**

The FST framework-based methods have been tested under a variety of scenarios and produce accurate results, as verified by closed-form solutions, alternative numerical methods and results in the literature computed by experts in the field using different methods. The methods are fast, requiring only two FFT evaluations per time step, and attain second-order convergence in the space variable. Path-dependent options have first-order convergence while for American options second-order convergence in time is attained through the implementation of the penalty method.

- **Efficient handling of path-independent and discretely-monitored derivatives**

For path-independent options, the FST methods require only one time step, computing option values and Greeks for a range of spot prices in a fraction of a second. For discretely-monitored derivatives, no time-stepping is required between the monitoring dates, i.e., the number of time steps is equal to the number of monitoring dates.

- **Ability to handle path-dependent and multi-asset derivatives**

The FST framework allows derivative prices (simultaneously for a range of spot prices) to be computed backwards in time over a time step of any length. As such, path-dependent options can be priced by choosing an appropriate time step and customizing the algorithm to take into account the unique features of the contract at each time step. Moreover, the framework is inherently multi-dimensional, allowing multi-asset options to be priced by utilizing the multi-dimensional FFT algorithm.

- **Generic handling of various spot price models and option payoffs**

The FST framework-based methods can accommodate a wide class of spot price models — single- and multi-asset, including independent-increment, mean-reverting, and regime-switching models. In the multi-asset case, the diffusion component can be correlated and jumps can be driven by a copula. Supplying different characteristic functions allows pricing under different stock price models (within the particular ‘class’ of models) with no further modifications to the algorithm. Furthermore, different options can be priced by supplying the appropriate payoff function, even where the Fourier transform of option payoff cannot be computed analytically.

- **Utilization of multi-core architectures**

With the FFT algorithm being inherently parallelizable, the FST framework-based methods are able to efficiently leverage multi-core computing architectures. Utilizing GPUs, for instance, delivers significant computational speedups for valuation of options (especially path-dependent options and in concurrent pricing) as compared to CPU-based pricing. For example, pricing of single-asset American options is sped up by a factor of 5 and two-asset American options by a factor of 10. Additional performance gains can be achieved by parallelizing the simultaneous pricing of multiple options.

7.2 Future Work

A number of open problems that can be potentially tackled using the FST methodology are mentioned below.

- **Pricing of multi-asset derivatives**

Currently, memory constraints limit the application of the FST method to two-dimensional problems, making pricing of multi-asset options, such as basket options, highly impractical. One potential approach is to utilize sparse grids — a technique to represent and integrate high dimensional functions.

- **Pricing under stochastic volatility models**

In Chapter 4.2 regime-switching models are used to generate stochastic volatility behavior

in stationary volatility models. While such an approach allows to better fit the long-term smile present in the markets, it does not capture the correlation between the spot price and the volatility level, as can be done in stochastic volatility models. Thus, applying the FST approach to stochastic volatility models remains a relevant open problem.

- **Error and convergence analysis**

The error and convergence properties of the FST and mrFST methods are studied in this work numerically. It is of great interest to establish analytic error and convergence rate bounds for these methods. Such analysis would also contribute to a more robust grid selection algorithm.

- **Improved precision for quadrature with FFTs**

The DFT approximates the Fourier integral using essentially the trapezoidal rule. To improve the accuracy of the integration, Simpson rule can be incorporated by premultiplying the vector v_m by the composite Simpson rule weights. Although, initial results indicate no advantage to usage of high-order quadrature methods, further investigation may yield positive results. Another area of application of high-order quadrature rule in conjunction with FFT is to improve precision of characteristic function evaluation for copula spot processes.

- **Pricing of interest rate options**

The one- and two-factor Hull-White and Vasicek models for the short interest rate are similar to the mean-reverting GBM models discussed in Chapter 3. Hence, developing a Fourier transform-based method for valuation of path-dependent interest rate derivatives is a viable research direction.

- **Pricing of real options**

Real options have become a valuable tool in making capital budgeting decision in corporate finance. As with financial options, the decision by a corporation to make a business decision, such as capital investment or sale of assets, is a right and not an obligation. Moreover, financial option pricing techniques are readily applicable to valuation of real options in corporate finance. The application of FST method to valuation of real options is another open research area.

- **Efficient utilization of high-performance architectures**

The results in Chapter 6 indicate that GPUs are a promising architecture for developing high-performance numerical algorithms. Efficient utilization of GPUs and alternative high-performance architectures for pricing of financial derivatives using various numerical algorithms could potentially yield a number of challenging research problems.

Appendix A

Acronyms and Notation

A.1 Acronyms

BSM	Black-Scholes-Merton (model)
CatEPut	Catastrophe equity put (option)
cdf	Cumulative distribution function
Cell BE	Cell Broadband Engine
CGMY	Carr-Geman-Madan-Yor (model)
CFT	Continuous Fourier transform
DFT	Discrete Fourier transform
DTSL	Double-trigger stop-loss (option)
FFT	Fast Fourier transform
GBM	Geometric Brownian motion
GPU	Graphics processing unit
i.i.d.	Independently and identically distributed
IV	Implied volatility
ODE	Ordinary differential equation
PDE	Partial differential equation
pdf	Probability density function
PIDE	Partial integro-differential equation
VG	Variance Gamma (model)

A.2 Notation

Stochastic Processes:

$X(t), \mathbf{X}(t)$	Log-spot price
$S(t), \mathbf{S}(t)$	Spot price
$W(t), \mathbf{W}(t)$	Brownian motion
$N(t), \mathbf{N}(t)$	Poisson process governing the arrival of jumps or losses
$J(t), \mathbf{J}(t)$	Jump process
$L(t), \mathbf{L}(t)$	Loss process
$\Theta(t)$	Reversion level process

Model Parameters:

r	Risk free interest rate
$\gamma, \boldsymbol{\gamma}$	Brownian motion drift
$\sigma, \boldsymbol{\Sigma}$	Brownian motion volatility and variance-covariance matrix
ρ	Correlation of Brownian motions
$\tilde{\nu}$	Poisson random measure
$\nu, \boldsymbol{\nu}$	Lévy density
v	Stochastic volatility level
$\kappa, \boldsymbol{\kappa}$	Mean-reversion speed
$\theta, \boldsymbol{\theta}$	Mean-reversion level
λ	Jump arrival rate
$\tilde{\mu}, \tilde{\sigma}$	Merton jump-diffusion model
η_p, η_+, η_-	Kou jump-diffusion model
μ	Variance Gamma model
C, G, M, Y	Carr-Geman-Madan-Yor model
χ, m_l, v_l	Joint catastrophe loss - stock price model

Option Parameters:

$\varphi(\mathbf{S})$	Payoff function
K	Strike price
T	Time to maturity
$\mathbf{B}$	Barrier
R	Rebate
β_1, β_2	Spread option
L^*	Catastrophe equity put option loss threshold

L_a, L_d Double-trigger stop-loss option attachment and detachment level

Option Value:

$V(t, \mathbf{x})$ Option value function

$v(t, \mathbf{x})$ Discount-adjusted, log-transformed option value function

$\mathbf{V}_m$ Option value on a discrete grid at time t_m

$\mathbf{v}_m$ Discount-adjusted, log-transformed option value on a discrete grid at time t_m

$\mathbf{v}^{(k)}$ Discount-adjusted, log-transformed option value on a discrete grid at iteration k for iterative methods

$\mathbf{v}^{(k)}$ Discount-adjusted, log-transformed option value on a discrete grid at regime k under regime-switching model

$\mathbf{v}^{[N,M]}$ Discount-adjusted, log-transformed option value on a discrete grid as a function of N space points and M time points

$\vec{\mathbf{v}}$ Collection of $\mathbf{v}^{(\cdot)}$ stacked into a column vector

$\Omega, \hat{\Omega}$ Discrete real and frequency space grids

Other:

$\mathcal{F}[\star](\omega), \hat{\star}$ Continuous Fourier transform of $\star$

$\mathcal{F}^{-1}[\star](\mathbf{x})$ Continuous inverse Fourier transform of $\star$

$\mathcal{L}$ Infinitesimal generator

$\mathcal{D}, \mathcal{J}$ Diffusion and integral (jump) components of the infinitesimal generator

$\mathbf{D}, \mathbf{J}$ Matrices associated with discretization of $\mathcal{D}$ and $\mathcal{J}$

$\mathbf{I}$ Identity matrix

$\Psi(\omega)$ Characteristic exponent

$\mathbb{Q}$ Risk-neutral pricing measure

$\mathbb{P}$ Real-world pricing measure

$\mathbb{E}_t^\star$ Expectation under $\star$ measure given information at t

$f_\star$ Probability density function of $\star$

α_d, α_j Finite-difference scheme parameter

$\langle \star \rangle_n$ The n -th component of vector $\star$

$\langle \star \rangle_{nm}$ The (n, m) -th component of matrix $\star$

$\tilde{\star}, \check{\star}$ Scaling of $\star$ (context dependent)

p_n, p_m Space and time convergence order

c_n, c_m Space and time convergence constant

Appendix B

Option and Model Parameters

Single-asset Options:

<i>EUR-A</i>	European put ($S = 100, K = 100, T = 10$)
<i>EUR-B</i>	European call ($S = 100, K = 110, T = 1$)
<i>EUR-C</i>	European call ($S = 100, K = 100, T = 0.25$)
<i>EUR-D</i>	European call ($S = 1, K = 1, T = 0.2$)
<i>EUR-E</i>	European call ($S = 100, K = 100, T = 0.46575$)
<i>EUR-F</i>	European put ($S = 10, K = 10, T = 0.25$)
<i>EUR-G</i>	European call ($S = 40, K = 50, T = 0.25$)
<i>DIG-A</i>	Digital call ($S = 100, K = 100, T = 0.5$)
<i>AMR-A</i>	American put ($S = 90, K = 98, T = 0.25$)
<i>AMR-B</i>	American put ($S = 100, K = 100, T = 0.25$)
<i>AMR-C</i>	American put ($S = 1369.41, K = 1200, T = 0.56164$)
<i>AMR-D</i>	American put ($S = 100, K = 95, T = 0.75$)
<i>BRM-A</i>	Bermudan daily-monitored put ($S = 40, K = 50, T = 0.25$)
<i>CBR-A</i>	Up-and-out barrier call ($S = 100, K = 100, B = 110, T = 1$)
<i>CBR-B</i>	Down-and-out barrier call ($S = 100, K = 110, B = 85, T = 1, R = 1$)
<i>DBR-A</i>	Discrete daily-monitored down-and-out barrier put ($S = 100, K = 105, T = 1, B = 90, R = 3$)
<i>DBR-B</i>	Discrete daily-monitored up-and-out barrier call ($S = 100, K = 100, T = 0.5, B = 115, R = 0.5$)
<i>SHT</i>	Shout put ($S = 100, T = 15, n = 5$)
<i>SWNG</i>	Swing ($S = 100, K = 100, T = 0.25, Q^{\min} = -3, Q^{\max} = 3$, Q is the difference between <i>up-swings</i> and <i>down-swings</i>)

Multi-asset Options:

<i>ESPD</i>	European spread call ($\beta_1 = 1, S_1 = 100, \beta_2 = 1, S_2 = 100, K = 2, T = 1$)
<i>BSPD</i>	Bermudan monthly-monitored spread call (same as above)
<i>ASPD</i>	American spread call (same as above)
<i>ECEP</i>	European catastrophe equity put ($S = 100, K = 100, L^* = 10, T = 5.0$)
<i>ADTSL</i>	American double-trigger stop-loss ($S = 100, K = 100, L_a = 5, L_d = 40, T = 7$)

Independent-Increment Models:

<i>BSM-A</i>	Black-Scholes-Merton ($\sigma = 0.15, r = 0.05, q = 0.02$)
<i>BSM-B</i>	Black-Scholes-Merton ($\sigma = 0.25, r = 0.03, q = 0.01$)
<i>BSM-C</i>	2D Black-Scholes-Merton ($\sigma_1 = 0.45, \sigma_2 = 0.15, \rho = 0.5, r = 0.05, q_1 = q_2 = 0.01$)
<i>MJD-A</i>	Merton jump-diffusion ($\sigma = 0.15, \lambda = 0.1, \tilde{\mu} = -1.08, \tilde{\sigma} = 0.4, r = 0.05, q = 0.02$)
<i>MJD-B</i>	Merton jump-diffusion ($\sigma = 0.15, \lambda = 0.1, \tilde{\mu} = 0.92, \tilde{\sigma} = 0.4, r = 0.05$)
<i>MJD-C</i>	Merton jump-diffusion ($\sigma = 0.15, \lambda = 0.1, \tilde{\mu} = -0.9, \tilde{\sigma} = 0.45, r = 0.05$)
<i>MJD-D</i>	Merton jump-diffusion ($\sigma = 0.25, \lambda = 2.0, \tilde{\mu} = 0.0, \tilde{\sigma} = 0.1, r = 0.05$)
<i>MJD-E</i>	2D Merton jump-diffusion ($\sigma_1 = 0.1, \lambda_1 = 0.25, \tilde{\mu}_1 = -0.13, \tilde{\sigma}_1 = 0.37, \sigma_2 = 0.2, \lambda_2 = 0.5, \tilde{\mu}_2 = 0.11, \tilde{\sigma}_2 = 0.41, \rho = 0.5, r = 0.1, q_1 = q_2 = 0.05$.)
<i>KJD-A</i>	Kou jump-diffusion ($\sigma = 0.2, \lambda = 0.2, \eta_p = 0.5, \eta_+ = 1/3, \eta_- = 1/2, r = 0$)
<i>VG-A</i>	Variance Gamma ($\gamma = -0.28113, \sigma = 0.19071, \mu = 0.49083, r = 0.0549, q = 0.011$)
<i>VG-B</i>	Variance Gamma ($\gamma = -0.22898, \sigma = 0.20722, \mu = 0.50215, r = 0.0541, q = 0.012$)
<i>CGMY-A</i>	CGMY ($C = 0.42, G = 4.37, M = 191.2, Y = 1.0102, r = 0.06$)
<i>CGMY-B</i>	CGMY ($C = 1.0, G = 8.8, M = 9.2, Y = 1.8, r = 0.1$)
<i>JSL</i>	Joint stock-loss ($\sigma = 0.15, \lambda = 1, m_l = 2, v_l = 5, \chi = 0.005, r = 0.05$)

Mean-Reverting Models:

<i>GBMMR</i>	Mean-reverting GBM model ($\sigma = 0.5, \theta = \ln 50, \kappa = 1, r = 0.04$)
<i>GBMMRM</i>	Mean-reverting GBM model with stochastic mean (defined on page 67)
<i>MJDMR</i>	Mean-reverting Merton jump-diffusion ($\sigma = 0.25, \lambda = 2, \tilde{\mu} = 0.1, \tilde{\sigma} = 0.2, \theta = \ln 90, \kappa = 0.5, r = 0.05$)
<i>KJDMR</i>	Mean-reverting Kou jump-diffusion ($\sigma = 0.3, \lambda = 4, \eta_p = 0.95, \eta_+ = 0.3, \eta_- = 0.1, \theta = \ln 92.0, \kappa = 5, r = 0.06$)
<i>KJDMRD</i>	Mean-reverting Kou jump-diffusion with decoupled jumps (defined on page 63)
<i>KJDMRC</i>	Mean-reverting Kou jump-diffusion with codependent copula jumps (defined on page 65)

Appendix C

Supplementary Results

In this section additional numerical results are presented to demonstrate the convergence and precision properties of the FST framework-based methods. Sections C.1 and C.2 present additional pricing and convergence results for the FST method that supplement Chapter 2. The former section further establishes the precision and convergence of the FST method while the latter section focuses on computing the time convergence of the method. Section C.3 provides a figure of option price and Greeks errors for an additional pricing scenario in Section 2.7. Section C.4 gives further pricing results for the mrFST method to supplement Chapter 3. Finally, Section C.5 provides pricing and timing results for the FST-GPU method in addition to the ones given in Chapter 6.

C.1 FST Pricing Results

N	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	4.38735869			0.590
4096	4.39027496	0.0029163		1.326
8192	4.39100310	0.0007281	2.0018	2.897
16384	4.39118505	0.0001820	2.0007	5.710
32768	4.39123053	0.0000455	2.0003	11.522

Table C.1: Pricing results for the European option *EUR-C* under the Merton jump-diffusion model *MJD-C*. Reference price 4.391243 and parameters from d’Halluin, Forsyth, and Vetzal (2005). The order of convergence is 2 in space.

N	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	0.042570525			0.465
4096	0.042628518	5.80×10^{-5}		1.151
8192	0.042642990	1.45×10^{-5}	2.0026	2.379
16384	0.042646605	3.62×10^{-6}	2.0009	4.706
32768	0.042647509	9.04×10^{-7}	1.9999	9.515

Table C.2: Pricing results for the European option *EUR-D* under the Kou jump-diffusion model *KJD-A*. Reference price 0.0426761 and parameters from Almendral and Oosterlee (2005). The reference price 0.0426478 is computed using a semi closed-form formula. The order of convergence is 2 in space.

N	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	7.49444358			0.620
4096	7.49618757	0.0017440		1.096
8192	7.49633753	0.0001500	3.5397	2.285
16384	7.49638178	0.0000442	1.7609	4.689
32768	7.49639296	0.0000112	1.9845	9.427

Table C.3: Pricing results for the European option *EUR-E* under the Variance Gamma model *VG-A*. Parameters from Hirs and Madan (2004). The reference price 7.49639670 is computed using the Fourier quadrature method. The order of convergence is 2 in space.

N	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	4.38983113			1.492
4096	4.38984017	9.04×10^{-6}		2.817
8192	4.38984243	2.26×10^{-6}	2.0008	5.671
16384	4.38984299	5.64×10^{-7}	2.0004	11.377
32768	4.38984313	1.41×10^{-7}	2.0002	23.014

Table C.4: Pricing results for the European option *EUR-F* under the CGMY model *CGMY-B*. Reference price of 4.3714972 and parameters from Forsyth, Wan, and Wang (2007). The reference price 4.38984331 is computed using the Fourier quadrature method. The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (s)
2048	128	3.23945333			0.011
4096	512	3.24080513	0.0013518		0.094
8192	2048	3.24114185	0.0003367	2.0053	0.753
16384	8192	3.24122597	0.0000841	2.0011	5.919
32768	32768	3.24124692	0.0000210	2.0050	49.513

Table C.5: Pricing results for the American option *AMR-B* under the Merton jump-diffusion model *MJD-C*. Reference price 3.2412435 and parameters from d’Halluin, Forsyth, and Labahn (2003). The order of convergence is 2 in space and 1 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (s)
2048	128	3.24037711			0.028
4096	256	3.24102477	0.0006477		0.116
8192	512	3.24119307	0.0001683	1.9442	0.469
16384	1024	3.24123826	0.0000452	1.8970	1.943
32768	2048	3.24125040	0.0000121	1.8957	8.276

Table C.6: Pricing results for the American option *AMR-B* under the Merton jump-diffusion model *MJD-C* using the FST penalty method. Reference price 3.2412435 and parameters from d’Halluin, Forsyth, and Labahn (2003). The order of convergence is 2 in space and approximately 2 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	128	35.47514840			0.011
4096	512	35.49042468	0.0152763		0.096
8192	2048	35.49263417	0.0022095	2.7895	0.752
16384	8192	35.49290881	0.0002746	3.0081	5.960
32768	32768	35.49286657	0.0000422	2.7010	49.241

Table C.7: Pricing results for the American option *AMR-C* under the Variance Gamma model *VG-B*. Reference price 35.5301 and parameters from Hirsa and Madan (2004). The order of convergence is 2 in space and at least 1 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec)
2048	128	35.49985425			0.028
4096	256	35.49075271	0.0091015		0.116
8192	512	35.49340795	0.0026552	1.7773	0.481
16384	1024	35.49278252	0.0006254	2.0859	2.014
32768	2048	35.49261794	0.0001646	1.9261	8.200

Table C.8: Pricing results for the American option *AMR-C* under the Variance Gamma model *VG-B* using the FST penalty method. Reference price 35.5301 and parameters from Hirsa and Madan (2004). The order of convergence is 2 in space and 2 in time.

N	M	No Extrapolation			Richardson Extrapolation			Time (sec.)
		Value	Change	$\log_2$ Ratio	Value	Change	$\log_2$ Ratio	
2048	128	9.135559597						0.015
4096	512	9.072063216	0.0634964		9.00856684			0.111
8192	2048	9.038987131	0.0330761	0.9409	9.00591105	0.0026558		0.660
16384	8192	9.022144722	0.0168424	0.9737	9.00530231	0.0006087	2.1253	6.239
32768	32768	9.013649661	0.0084951	0.9874	9.00515460	0.0001477	2.0430	47.936

Table C.9: Pricing results for the barrier option *CBR-B* under the Merton jump-diffusion model *MJD-D*. Reference price 9.013 and parameters from Metwally and Atiya (2003). The order of convergence is 2 in space and 1 in time with Richardson extrapolation (and 1/2 in time without extrapolation).

C.2 FST Time Convergence Results

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
4096	256	4.536390832			0.067
4096	512	4.536915681	0.0005248		0.127
4096	1024	4.537178488	0.0002628	0.9979	0.242
4096	2048	4.537309690	0.0001312	1.0022	0.410
8192	256	4.536438057			0.138
8192	512	4.536961144	0.0005231		0.293
8192	1024	4.537224131	0.0002630	0.9921	0.555
8192	2048	4.537355290	0.0001312	1.0037	0.787
16384	256	4.536448484			0.238
16384	512	4.536971785	0.0005233		0.412
16384	1024	4.537234408	0.0002626	0.9946	0.728
16384	2048	4.537366137	0.0001317	0.9954	1.241

Table C.10: Pricing results for the American option $AMR-D$ under the Merton jump-diffusion model $MJD-C$. The order of convergence is 1 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
4096	256	47.097252960			0.045
4096	512	47.106248530	0.0089956		0.082
4096	1024	47.110929806	0.0046813	0.9423	0.151
4096	2048	47.113495967	0.0025662	0.8673	0.302
8192	256	47.097437718			0.099
8192	512	47.106412070	0.0089744		0.194
8192	1024	47.110940177	0.0045281	0.9869	0.357
8192	2048	47.113296390	0.0023562	0.9424	0.743
16384	256	47.097431921			0.201
16384	512	47.106407276	0.0089754		0.413
16384	1024	47.110938607	0.0045313	0.9860	0.814
16384	2048	47.113217736	0.0022791	0.9915	1.507

Table C.11: Pricing results for the American option $AMR-C$ under the CGMY model $CGMY-A$. The order of convergence is 1 in time.

C.3 FST Greeks Results

Figure C.1: Error in option price (top, left), Delta (top, right), Gamma (middle, left), Rho (middle, right), Theta (bottom, left), and Vega (bottom, right) for the digital option *DIG-A* under the Merton jump-diffusion model *MJD-A*. The average rate of convergence (across all spot prices) for the option price and all Greeks is 2 in space.

C.4 mrFST Pricing Results

N	M	Value	Change	$\log_2$ Ratio	Time (msec.)
2048	10	24.31867556			4.344
4096	10	24.31792551	0.0007500		8.512
8192	10	24.31769489	0.0002306	1.7015	17.553
16384	10	24.31763689	0.0000580	1.9912	35.218
32768	10	24.31762061	0.0000163	1.8334	72.281

Table C.12: Pricing results for the European option *EUR-B* under the mean-reverting Kou jump-diffusion model *KJDMR*. The reference price 24.28749709 is computed using Monte Carlo simulation (95% confidence interval width of 0.1494062 requires 445.367 seconds for 2^{20} sample paths with 2^{10} time points). The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	256	9.19726790			0.070
4096	512	9.19438807	2.88×10^{-3}		0.275
8192	1024	9.19648217	2.09×10^{-4}	0.4597	1.151
16384	2048	9.19642824	5.39×10^{-5}	5.2791	4.918
32768	4096	9.19642770	5.39×10^{-7}	6.6457	20.834

Table C.13: Pricing results for the American option *AMR-A* under the mean-reverting Kou jump-diffusion model *KJDMR* with Richardson extrapolation. The order of convergence is 2 in space and 2 in time.

N	M	Value	Change	$\log_2$ Ratio	Time (sec.)
2048	126	0.58232029			0.033
4096	126	0.58303190	0.0007116		0.067
8192	126	0.58244527	0.0005866	0.2786	0.156
16384	126	0.58232817	0.0001171	2.3248	0.318
32768	126	0.58231749	0.0000107	3.4541	0.674

Table C.14: Pricing results for the discrete barrier option *DBR-B* under the mean-reverting Kou jump-diffusion model *KJDMR*. The reference price 0.58289924 is computed using Monte Carlo simulation (95% confidence interval width of 0.0028937 requires 56.099 seconds for 2^{20} sample paths with 126 time points). The order of convergence is at least 2 in space.

C.5 GPU Pricing Results

N	Value	Change	$\log_2$ Ratio	CPU Time (msec.)	GPU Time (msec.)
2048	0.41935721			2.358	2.607
4096	0.41935826	1.05×10^{-6}		4.711	4.864
8192	0.41935852	2.62×10^{-7}	2.0017	9.291	9.449
16384	0.41935858	6.55×10^{-8}	2.0023	18.838	18.798
32768	0.41935860	1.63×10^{-8}	2.0041	38.024	37.303

Table C.15: Pricing results for the European option *EUR-D* under the CGMY model *CGMY-B*. The reference price 0.41935843 is computed using the Fourier quadrature method. The order of convergence is 2 in space.

N	M	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
2048	128	3.50287779			0.010	0.017
4096	512	3.50401998	0.0011422		0.078	0.075
8192	2048	3.50430760	0.0002876	1.9896	0.656	0.344
16384	8192	3.50437906	0.0000715	2.0090	6.267	1.718
32768	32768	3.50439685	0.0000178	2.0056	47.031	10.601

Table C.16: Pricing results for the American option *AMR-B* under the Merton jump-diffusion model *MJD-A*. The order of convergence is 2 in space and 1 in time.

N	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
512^2	14.46794939			0.157	0.134
1024^2	14.46916508	1.22×10^{-3}		0.571	0.466
2048^2	14.46924603	1.63×10^{-5}	3.9087	2.155	1.732
4096^2	14.46924541	6.14×10^{-7}	7.0433	8.585	N/A
8192^2	14.46924541	2.08×10^{-9}	8.2071	35.576	N/A

Table C.17: Pricing results for the European spread option *ESPD* under the 2D Black-Scholes-Merton model *BSM-C*. The reference price 14.47616356 is computed using Kirk's approximation formula. The order of convergence is at least 2 in space.

N	M	Value	Change	$\log_2$ Ratio	CPU Time (sec.)	GPU Time (sec.)
512^2	64	14.48581963			0.986	0.179
1024^2	256	14.48730320	0.0014836		19.786	1.431
2048^2	1024	14.48739276	0.0000896	4.0500	326.536	29.031
4096^2	4096	14.48738724	0.0000055	4.0191	6541.295	N/A

Table C.18: Pricing results for the American spread option *ASPD* under the 2-dimensional Black-Scholes-Merton model *BSM-C*. The order of convergence is a least 2 in space and 2 in time.

Bibliography

- Albanese, C., S. Jaimungal, and D. Rubisov (2003). A two-state jump model. *Quantitative Finance* 3, 145–154.
- Albanese, C., S. Jaimungal, and D. H. Rubisov (2001). The model of lines for option pricing with jumps. Working paper, available at <http://ssrn.com/abstract=256845>.
- Almendral, A. and C. W. Oosterlee (2005). Numerical valuation of options with jumps in the underlying. *Applied Numerical Mathematics* 53, 1–18.
- Andersen, L. and J. Andreasen (2000). Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research* 4, 231–262.
- Andricopoulos, A., M. Widdicks, P. Duck, and D. Newton (2003). Universal option valuation using quadrature. *Journal of Financial Economics* 67(3), 443–471.
- Ascher, U. M., S. J. Ruuth, and B. T. Wetton (1997). Implicit-explicit methods for time-dependent PDEs. *SIAM Journal on Numerical Analysis* 32(3), 797–823.
- Bachelier, L. (1900). Théorie de la speculation. *Annales Scientifiques de L'École Normale Supérieure* 17, 21–86.
- Barlow, M., Y. Gusev, and M. Lai (2004). Calibration of multifactor models in electricity markets. *International Journal of Theoretical and Applied Finance* 7(2), 101–120.
- Barndorff-Nielsen, O. E. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics* 24, 1–13.
- Bates, D. S. (1991). The crash of '87: Was it expected? the evidence from options markets. *The Journal of Finance* 46(3), 1009–1044.
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* 9(1), 69–107.
- Black, F. and M. S. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81(3), 637–654.

- Boyarchenko, S. I. and S. Z. Levendorskii (2002). *Non-Gaussian Merton-Black-Scholes Theory*, Volume 9 of *Advanced Series on Statistical Science and Applied Probability*. World Scientific Publishing Co.
- Boyle, P., M. Broadie, and P. Glasserman (1997). Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control* 21(8-9), 1267–1321.
- Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics* 4, 323–338.
- Boyle, P. P. (1986). Option valuation using a three jump process. *International Options Journal* 3, 7–12.
- Briani, M., R. Natalini, and G. Russo (2007). Implicit-explicit numerical schemes for jump-diffusion processes. *CALCOLO* 44(1), 33–57.
- Broadie, M. and Y. Yamamoto (2005). A double-exponential fast Gauss transform algorithm for pricing discrete path-dependent options. *Operations Research* 53(5), 764–779.
- Buchen, P. W. (1996). Pricing European barrier options. Technical Report 25, The University of Sydney.
- Carmona, R. and V. Durrleman (2003). Pricing and hedging spread options. *SIAM Review* 45(4), 627–685.
- Carr, P., H. Geman, D. B. Madan, and M. Yor (2002). The fine structure of asset returns: An empirical investigation. *Journal of Business* 75(2), 305–332.
- Carr, P., R. Jarrow, and R. Myneni (1992). Alternative characterizations of American put options. *Mathematical Finance* 2, 87–106.
- Carr, P. and D. B. Madan (1999). Option valuation using the fast Fourier transform. *The Journal of Computational Finance* 2(4), 61–73.
- Carr, P. and L. Wu (2007). Stochastic skew in currency options. *Journal of Financial Economics* 86, 213–247.
- Cartea, A. and M. Figueroa (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 22(4), 313–335.
- Cherubini, U., E. Luciano, and W. Vecchiato (2004). *Copula Methods in Finance*. John Wiley & Sons, Inc.
- Cheuk, T. and T. Vorst (1997). Shout floors. *Net Exposure* 2.
- Chiarella, C., N. El-Hassan, and A. Kucera (1999). Evaluation of American option prices in a path integral framework using Fourier-Hermite series expansions. *Journal of Economic Dynamics and Control* 23, 1387–1424.

- Chiarella, C. and A. Ziogas (2005). Pricing American options on jump-diffusion processes using Fourier Hermite series expansions. Research Paper Series 145, Quantitative Finance Research Centre.
- Chourdakis, K. (2005). Switching Lévy models in continuous time: Finite distributions and option pricing. Working paper, available at <http://ssrn.com/abstract=838924>.
- Clelow, L. and C. Strickland (2000). *Energy Derivatives: Pricing and Risk Management*. Lacima Publications.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. Chapman & Hall / CRC.
- Cont, R. and E. Voltchkova (2005). A finite difference scheme for option pricing in jump diffusion and exponential Lévy models. *SIAM Journal on Numerical Analysis* 43(4), 1596–1626.
- Cortazar, G. and E. Schwartz (1994). The valuation of commodity contingent claims. *The Journal of Derivatives* 1, 27–29.
- Cox, J. C., S. A. Ross, and M. Rubinstein (1979). Option pricing: A simplified approach. *Journal of Financial Economics* 7(3), 229–263.
- Cox, S. H., J. R. Fairchild, and H. W. Pedersen (2004). Valuation of structured risk management products. *Insurance: Mathematics and Economics* 34, 259–272.
- Dai, M. and Y. K. Kwok (2008). Optimal multiple stopping models of reload options and shout options. *Journal of Economic Dynamics and Control* 32, 2269–2290.
- Dai, M., Y. K. Kwok, and L. Wu (2004). Optimal shouting policies of options with strike reset right. *Mathematical Finance* 14(3), 383–401.
- Dempster, M. and J. Hutton (1997). Fast numerical valuation of American, exotic and complex options. *Applied Mathematical Finance* 4, 1–20.
- Dempster, M. A. and S. Hong (2000). Spread option valuation and the fast Fourier transform. Technical Report 26, Cambridge University.
- Derman, E. and I. Kani (1994, January). The volatility smile and its implied tree. Quantitative Strategies Research Notes, Goldman Sachs.
- d’Halluin, Y., P. A. Forsyth, and G. Labahn (2003). A penalty method for American options with jump diffusion processes. *Numerische Mathematik* 97(2), 321–352.
- d’Halluin, Y., P. A. Forsyth, and K. R. Vetzal (2005). Robust numerical methods for contingent claims under jump diffusion processes. *IMA Journal of Numerical Analysis* 25, 87–112.

- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Eberlein, E. and K. Prause (2002). The generalized hyperbolic model: Financial derivatives and risk measures. In H. Geman, D. Madan, S. Pliska, and T. Vorst (Eds.), *Mathematical Finance - Bachelier Congress 2000*, pp. 245–267. Springer.
- Elliott, R. and C. Osakwe (2006). Option pricing for pure jump processes with Markov switching compensators. *Finance and Stochastics* 10, 250–275.
- Eydeland, A. and K. Wolyniec (2003). *Energy and Power Risk Management*. John Wiley & Sons, Inc.
- Fang, F. and C. W. Oosterlee (2008). A novel pricing method for European options based on Fourier-Cosine series expansions. *SIAM Journal on Scientific Computing* 31(2), 826–848.
- Feng, L. and V. Linetsky (2008). Pricing discretely monitored barrier options and defaultable bonds in Lévy process models: A fast Hilbert transform approach. *Mathematical Finance* 18(3), 337–384.
- Figuerola, M. G. (2006). Pricing multiple interruptible-swing contracts. Working paper, available at <http://www.ems.bbk.ac.uk/research/wp/PDF/BWPEF0606.pdf>.
- Forsyth, P. A. and K. R. Vetzal (2002). Quadratic convergence of a penalty method for valuing American options. *SIAM Journal on Scientific Computing* 23, 2096–2123.
- Forsyth, P. A., J. W. Wan, and I. R. Wang (2007). Robust numerical valuation of European and American options under the CGMY process. *Journal of Computational Finance* 10(4), 31–69.
- Frigo, M. and S. G. Johnson (2005). The design and implementation of FFTW3. *Proceedings of the IEEE* 93(2), 216–231. Special issue on "Program Generation, Optimization, and Platform Adaptation".
- Geman, H. (2005). *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, Inc.
- Gerbessiotis, A. V. (2004). Architecture independent parallel binomial tree option price valuations. *Parallel Computing* 30, 301–316.
- Gibson, R. and E. Schwartz (1990). Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance* 45, 959 – 976.
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.

- Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics* 45, 39–70.
- Harrison, J. and S. Pliska (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications* 11(1), 215–260.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies* 6(2), 327–343.
- Hikspoors, S. and S. Jaimungal (2007). Energy spot price models and spread options pricing. *International Journal of Theoretical and Applied Finance* 10(7), 1111–1135.
- Hirsa, A. and D. B. Madan (2004). Pricing American options under Variance Gamma. *The Journal of Computational Finance* 7(2), 63–80.
- Huang, J. and J. Pang (1998). Option pricing and linear complementarity. *Journal of Computational Finance* 2, 31–60.
- Huang, J.-Z. and L. Wu (2004). Specification analysis of option pricing models based on time-changed Lévy processes. *The Journal of Finance* 59, 1405–1440.
- Hull, J. and A. White (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance* 42(2), 281–300.
- Hull, J. C. (2005). *Options, Futures, and Other Derivatives* (6th ed.). Prentice Hall.
- Ibáñez, A. (2004). Valuation by simulation of contingent claims with multiple early exercise opportunities. *Mathematical Finance* 14(2), 223–248.
- Jackson, K. R., S. Jaimungal, and V. Surkov (2007). Option pricing with regime switching Lévy processes using Fourier space time-stepping. In *Proceeding of the Fourth IASTED International Conference on Financial Engineering and Applications*, pp. 92–97.
- Jackson, K. R., S. Jaimungal, and V. Surkov (2008). Fourier space time-stepping for option pricing with Lévy models. *Journal of Computational Finance* 12(2), 1–28.
- Jaillet, P., E. I. Ronn, and S. Tompaidis (2004). Valuation of commodity-based swing options. *Management Science* 50(7), 909–921.
- Jaimungal, S. and T. Wang (2006). Catastrophe options with stochastic interest rates and compound Poisson losses. *Insurance: Mathematics and Economics* 38(3), 469–483.
- Kim, I. (1990). Analytic valuation of American options. *Review of Financial Studies* 3, 547–572.
- Kirk, E. (1995). *Correlation in the Energy Markets*. Risk Publications and Enron.
- Kluge, T. (2006). *Pricing Swing Options and other Electricity Derivatives*. Ph. D. thesis, University of Oxford.

- Kolb, C. and M. Phar (2005). Option pricing on the GPU. *GPU Gems 2*, 719–731.
- Konikov, M. and D. Madan (2000). Option pricing using Variance-Gamma Markov chains. *Review of Derivatives Research* 5, 81–15.
- Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science* 48(8), 1086–1101.
- Kou, S. G. and G. Petrella (2004). Numerical pricing of discrete barrier and lookback options via Laplace transforms. *Journal of Computational Finance* 8, 1–37.
- Lewis, A. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. Working paper, available at <http://ssrn.com/abstract=282110>.
- Lin, S. and T. Wang (2009). Pricing perpetual American catastrophe put options: A penalty function approach. *Insurance: Mathematics and Economics* 44(2), 287–295.
- Lindgren, G. (1978). Markov regime models for mixed distributions and switching regressions. *Scandinavian Journal of Statistics* 5, 81–91.
- Liu, H. and J. Zou (2006). Some new additive Runge-Kutta methods and their applications. *Journal of Computational and Applied Mathematics* 190, 74–98.
- Lord, R., F. Fang, F. Bervoets, and C. W. Oosterlee (2008). A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes. *SIAM Journal on Scientific Computing* 30(4), 1678–1705.
- Madan, D. and E. Seneta (1990). The Variance Gamma (V.G.) model for share market returns. *Journal of Business* 63(4), 511–524.
- Madan, D. B., F. Milne, and H. Shefrin (1989). The multinomial option pricing model and its Brownian and Poisson limits. *The Review of Financial Studies* 2(2), 251–265.
- Matache, A. M., P. Nitsche, and C. Schwab (2005). Wavelet Galerkin pricing of American options on Lévy driven assets. *Quantitative Finance* 5(4), 403 – 424.
- Matache, A. M., T. von Petersdorff, and C. Schwab (2004). Fast deterministic pricing of options on Lévy driven assets. *Modélisation Mathématique et Analyse Numérique* 38(1), 37–71.
- McKean, H. J. (1965). Appendix: A free boundary value problem for the heat equation arising from a problem in mathematical economics. *Industrial Management Review* 6, 32–39.
- Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science* 4(1), 141–183.

- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Metwally, S. A. and A. F. Atiya (2003). Fast Monte Carlo valuation of barrier options for jump diffusion processes. In *Proc. Computational Intelligence for Financial Engineering*, pp. 101–107.
- Naik, V. (1993). Option valuation and hedging strategies with jumps in the volatility of assets returns. *Journal of Finance* 48(5), 1969–1984.
- Osborne, M. (1959). Brownian motion in the stock market. *Operations Research* 7(2), 145–173.
- O’Sullivan, C. (2005). Path dependent option pricing under Lévy processes. In *EFA 2005 Moscow Meetings*.
- Papapantoleon, A. (2005). An introduction to Lévy processes with applications in finance. Lecture notes, available at <http://arxiv.org/abs/0804.0482>.
- Pareschi, L. and G. Russo (2000). Implicit-explicit Runge-Kutta schemes for stiff systems of differential equations. *Advances in the Theory of Computational Mathematics* 3, 269–288.
- Pareschi, L. and G. Russo (2005). Implicit-explicit Runge-Kutta schemes and applications to hyperbolic systems with relaxation. *Journal of Scientific Computing* 25(1), 129–155.
- Pilipović, D. (1997). *Energy Risk: Valuing and Managing Energy Derivatives*. McGraw-Hill.
- Pilipović, D. and J. Wengler (1998). Getting into the swing. *Energy Power Risk Management* 2(10).
- Podlozhnyuk, V. (2007). Black-Scholes option pricing. Part of CUDA SDK documentation.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery (1992). *Numerical Recipes in C: The art of scientific computing*. Cambridge University Press.
- Raible, S. (2000). *Lévy processes in finance: theory, numerics, and empirical facts*. Ph. D. thesis, University of Freiburg.
- Reiner, E. (2001). Convolution methods for path-dependent options. Presentation slides from Financial Mathematics conference at IPAM, UCLA, available at http://www.ipam.ucla.edu/publications/fm2001/fm2001_4272.pdf.
- Sak, H., S. Özekici, and I. Boduroğlu (2007). Parallel computing in Asian option pricing. *Parallel Computing* 33, 92–108.
- Sato, K.-I. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge Studies in Advanced Mathematics. Cambridge University Press.

- 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. (1987). Option pricing when the variance changes randomly: Theory, estimation and an application. *The Journal of Financial and Quantitative Analysis* 22(4), 419–438.
- Sklar, A. (1959). Fonctions de repartition à n dimensions et leurs marges. *Publications of the Institute of Statistics* 8, 229–231.
- Stein, E. M. and J. C. Stein (1991). Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies* 4(4), 727–52.
- Strauss, W. A. (1992). *Partial Differential Equations: An Introduction*. John Wiley & Sons, Inc.
- Surkov, V. (2008). Parallel option pricing with Fourier Space Time-stepping method on Graphics Processing Units. In *Proceedings of the First Workshop on Parallel and Distributed Computing in Finance*.
- Taylor, M. E. (1997). *Partial Differential Equations*, Volume I-III. Springer-Verlag.
- Thomas, B. (1993). Something to shout about. *Risk* 6(5), 56–58.
- Ware (2005). Swing options in a mean-reverting world. Presentation slides from Stochastic Calculus and its applications to Quantitative Finance and Electrical Engineering conference, University of Calgary, available at <http://finance.math.ucalgary.ca/papers/WareElliottConf2005.pdf>.
- Wilmott, P., S. Howison, and J. Dewynne (1993). *Option Pricing: Mathematical Models and Computation*. Oxford Financial Press.
- Wilmott, P., S. Howison, and J. Dewynne (1995). *The Mathematics of Financial Derivatives*. Cambridge University Press.
- Windcliff, H. A., P. A. Forsyth, and K. R. Vetzal (2001). Shout options: a framework for pricing contracts which can be modified by the investor. *Journal of Computational and Applied Mathematics* 134, 213–241.
- Yang, C. C., M. Wang, and X. Chen (2008). Catastrophe effects on stock markets and catastrophe risk securitization. *The Journal of Risk Finance* 9(3), 232–243.