



City Research Online

City St George's, University of London

Citation: Kyriakou, I. (2010). Efficient valuation of exotic derivatives with path-dependence and early exercise features. (Unpublished Doctoral thesis, City University London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/23459/>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Efficient valuation of exotic derivatives with path-dependence and early-exercise features

by

Ioannis Kyriakou

A thesis submitted in partial fulfilment of the requirements for the
degree of
Doctor of Philosophy

City University, London
Sir John Cass Business School
Faculty of Actuarial Science and Insurance
November 2010

Acknowledgements

I would like to thank first my supervisors Dr Laura Ballotta and Professor Aleš Černý without whom, this research would not have been done. I thank also Dr Iqbal Owadally who motivated me towards the research path.

I am grateful to my friends, with particular reference to Demetris Kyprianou and my colleagues Nikolaos Papapostolou and Panos Pouliasis, for their constant emotional support and all the necessary distractions they provided during the research study period.

A special thank goes to Sir John Cass Business School and EPSRC for providing grant to support this research.

Last but not least, I would like to thank my family, especially my father Savvas and mother Despo, who unconditionally supported me, both emotionally and financially, throughout the course of this research. I am grateful to my grandmother Niki for her invaluable contribution throughout my entire education, even in her absence. I dedicate this thesis to my parents, grandparents Giannakis and Ersi, and late grandparents Tasos and Niki.

Declaration

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

I hereby confirm that the thesis is my own work, except for Chapter 4 which was co-authored with Professor Aleš Černý and for which I claim a share of 50%.

Abstract

The main objective of this thesis is to provide effective means for the valuation of popular financial derivative contracts with path-dependence and/or early-exercisable provisions. Starting from the risk-neutral valuation formula, the approach we propose is to sequentially compute convolutions of the value function of the contract at a monitoring date with the transition density between two dates, to provide the value function at the previous monitoring date, until the present date. A rigorous computational algorithm for the convolutions is then developed based on transformations to the Fourier domain.

In the first part of the thesis, we deal with arithmetic Asian options, which, due to the growing popularity they enjoy in the financial marketplace, have been researched significantly over the last two decades. Although few remarkable approaches have been proposed so far, these are restricted to the market assumptions imposed by the standard Black-Scholes-Merton paradigm. Others, although in theory applicable to Lévy models, are shown to suffer a non-monotone convergence when implemented numerically. To solve the Asian option pricing problem, we initially propose a flexible framework for independently distributed log-returns on the underlying asset. This allows us to generalize firstly in calculating the price sensitivities. Secondly, we consider an extension to non-Lévy stochastic volatility models. We highlight the benefits of the new scheme and, where relevant, benchmark its performance against an analytical approximation, control variate Monte Carlo strategies and existing forward convolution algorithms for the recovery of the density of the underlying average price.

In the second part of the thesis, we carry out an analysis on the rapidly growing market of convertible bonds (CBs). Despite the vast amount of research which has been undertaken about this instrument over the last thirty years, no pricing paradigm has been standardized yet. This is due to the need for proper modelling of the CBs' composite payout structure and the multifactor modelling arising in the CB valuation. Given the dimensional capacity of the convolution algorithm, we are now able to introduce a new jump diffusion structural approach in the CB literature, towards more realistic modelling of the default risk, and further include correlated stochastic interest rates. This aims at fixing dimensionality and convergence limitations which previously have been restricting the range of applicability of popular grid-based, lattice and Monte Carlo methods. The convolution scheme further permits flexible handling of real-world CB specifications; this allows us to properly model the firm's call policy and investigate its impact on the computed CB prices. We illustrate the performance of the numerical scheme and highlight the effects originated by the inclusion of jumps.

Contents

1	Introduction	1
2	Fourier transforms	5
2.1	Introduction	5
2.2	Fourier transforms	6
2.2.1	Derivatives of a function and their Fourier transforms	6
2.2.2	Convolution	7
2.2.3	Inverse Fourier transform	7
2.2.4	Multi-dimensional Fourier transforms	8
2.3	Characteristic function	9
2.4	Discrete Fourier transforms	11
2.5	Fourier transforms on a grid	15
2.5.1	Multi-dimensional discrete approximations	17
3	Affine models and option pricing	19
3.1	Introduction	19
3.2	Lévy processes	20
3.3	Affine processes	21
3.4	Option pricing with Fourier transforms	23
4	A backward convolution algorithm for discretely sampled arithmetic Asian options	29
4.1	Introduction	29

4.2	Pricing approaches to arithmetic Asian options	31
4.3	Modelling on reduced state space	34
4.4	Pricing of Asian options by convolution	36
4.5	The backward price convolution algorithm	37
4.5.1	Numerical implementation	38
4.6	Numerical study	40
4.6.1	Models	40
4.6.2	Pricing in the Black-Scholes-Merton economy	41
4.6.3	Pricing in Lévy economies	44
4.6.4	Standard DFT versus fractional DFT	46
4.7	Concluding remarks	48
5	Computation of the Asian option price sensitivities	50
5.1	Introduction	50
5.2	Price sensitivities via direct differentiation	52
5.2.1	Distribution-based approach for price sensitivities in affine models	54
5.3	A convolution approach for the Asian option price sensitivities	56
5.4	Computation of the price sensitivities via Monte Carlo simulation	62
5.4.1	Likelihood ratio estimators coupled with Fourier transforms	64
5.5	Numerical study	66
5.5.1	Convolution versus Monte Carlo	67
5.6	Concluding remarks	68
6	Pricing Asian options under stochastic volatility	71
6.1	Introduction	71
6.2	Pricing approaches to Asian options under stochastic volatility	73
6.3	Market models	74
6.3.1	Laws of affine stochastic volatility models	76
6.4	Modelling on reduced state space with stochastic volatility	77
6.5	The backward price convolution algorithm	78
6.5.1	Numerical implementation	79

6.6	Asian option pricing via Monte Carlo simulation	82
6.6.1	Geometric Asian options	82
6.6.2	Simulation of the Heston model	85
6.7	Numerical study	88
6.7.1	Models	88
6.7.2	Pricing via convolution	88
6.7.3	Monte Carlo pricing	90
6.8	Concluding remarks	95
	Appendix 6.A–The truncated Gaussian and quadratic-exponential schemes	96
	Appendix 6.B–Spot measure change for Lévy and time-changed Lévy processes	99
	Appendix 6.C–Characterization of the log-return distribution conditional on the vari- ance at the endpoints of a time interval	101
7	Monte Carlo option pricing coupled with Fourier transformation	103
7.1	Introduction	103
7.2	Monte Carlo simulation coupled with Fourier transform	104
7.2.1	Numerical implementation	105
7.3	Market model setup and option pricing	107
7.4	The tempered stable framework	108
7.4.1	Properties	108
7.4.2	CGMY as time-changed Brownian motion	109
7.5	Numerical study	111
7.5.1	Distribution function tests	112
7.5.2	Simulation tests	113
7.6	Concluding remarks	118
	Appendix 7–Simulation of the CGMY process using change of measure	120
8	A backward convolution algorithm for convertible bonds in a jump diffusion setting with stochastic interest rates	121
8.1	Introduction	121
8.2	Market model	125

8.2.1	The firm value-interest rate setup	125
8.2.2	Stock versus firm value and real-world considerations	126
8.2.3	Calibration issues	129
8.3	Convertible bonds: contract features	131
8.3.1	The optimal call strategy	131
8.3.2	The payoff function and pricing considerations	133
8.4	The backward price convolution algorithm	134
8.4.1	The call payoff	136
8.4.2	Numerical implementation	137
8.5	Numerical study	139
8.5.1	Black-Scholes-Merton model	139
8.5.2	Jump diffusion setup	141
8.5.3	Effects of discrete coupon and dividend payments	143
8.5.4	Effects of call policy	145
8.6	Concluding remarks	147
	Appendix 8.A—Equivalent martingale measure changes: the t -forward measure	149
	Appendix 8.B—Characterization of the bivariate log-firm value-interest rate process under the forward measure	151
9	Concluding remarks	155
	References	159

Chapter 1

Introduction

In recent years there has been quite an increase in the popularity of path-dependent derivatives, so called since their payoffs are related to movements in the price of some underlying asset throughout the life (or part of the life) of the contract. Of particular interest to the traders are the Asian options, whose payout depends on the average value of some asset observed over a preset time window. Their appeal stems partly from the fact that the option payout does not depend on a single snapshot of the underlying asset's price, thus reducing the risk of market speculation. Also, the averaging has a smoothing effect on the fluctuating behaviour of the underlying asset, resulting in lower option prices.

The way the average is defined (geometric versus arithmetic; continuously versus discretely monitored asset values) plays a key role in the analytical tractability of the option; in contrast to the more prevalent arithmetic average, a closed analytical pricing formula is available for the geometric average under the standard Black-Scholes-Merton market model assumptions. This has given rise to a large amount of research over the last two decades towards the accurate and efficient calculation of the price and sensitivities of this instrument.

Common in the financial marketplace are also exotic structures with early-exercise features, i.e., contracts that may be exercised prior to their expiration. Convertible bonds fall into this category. These are corporate debt securities that offer the investors the right to forgo future coupon and/or principal payments in exchange for a predetermined number of shares. From the issuer's perspective, the key benefit of raising money by selling CBs is a reduced coupon level. Hence, CBs depend on variables related to the underlying firm value (or stock), the fixed

income part (interest rates and credit risk), and the interaction between these components; for this, they are frequently characterized as hybrid claims. So far, CBs have raised significant challenges for practitioners and academics. This is because no pricing technique, generally robust and easily adaptable to their complicated structure and the governing random factors, has been standardized yet.

This thesis is dedicated to the efficient valuation of the abovementioned exotic derivatives using techniques that strongly rely on numerical integration enhanced by Fourier transforms.

The next two chapters provide the building blocks for the chapters to follow. In particular, the notion of a Fourier transform and associated key results are delineated in Chapter 2. We discuss how continuous Fourier integrals can be approximated by discrete, truncated Fourier series expansions, to allow for fast and accurate computation via the fast Fourier transform (FFT) algorithm. Chapter 3 illustrates that Fourier transforms provide an important ground for pricing contingent claims on underlyings that are driven by affine models with closed-form characteristic functions; we consider three main applications of Fourier transforms in pricing European-type options (Heston (1993), Carr and Madan (1999), Fang and Oosterlee (2008a)). Applications in pricing exotic products, like Asian, Bermudan and American vanilla, barrier and lookback options (see Carverhill and Clewlow (1990), Eydeland (1994), Broadie and Yamamoto (2003), (2005), Lord et al. (2008), Fang and Oosterlee (2008b), Feng and Linetsky (2008), Feng and Lin (2009)), are also discussed.

Chapters 4, 5 and 6 contribute to the effective computation of the prices and sensitivities of discretely sampled arithmetic Asian options. Similar in spirit to backward pricing on a lattice, the method is based on backward recursive evaluation of the expected option payout via numerical integration, relying heavily on Fourier transformations. The core idea is to utilize the risk-neutral valuation integral formula within the so-called Carverhill-Clewlow-Hodges framework and, following any necessary changes of measure, recognize that this in fact forms a convolution. The convolution is then expressed in terms of Fourier integrals which are dealt with numerically by means of the FFT algorithm. The method only requires knowledge of the characteristic function identifying the joint law of the state variables involved, hence permitting its applicability within the class of affine models.

Under Lévy assumptions for the asset log-returns, it is shown in Chapter 4 that modelling

the option price straightaway provides substantial numerical improvement over existing forward convolutions (see Carverhill and Clewlow (1990), Benhamou (2002), Fusai and Meucci (2008)) aimed at recovering the density of the underlying average instead. Additional speed-accuracy comparisons with a control variate Monte Carlo strategy and the compelling analytical approximation by Lord (2006a) demonstrate the soundness of this approach.

In light of the need for accurate price sensitivities for risk management purposes, but also as a measure of the pricing error resulting from potentially inappropriate parameter values, Chapter 5 generalizes the pricing methodology introduced in Chapter 4 in computing the sensitivities with respect to any parameter of interest. Furthermore, standard Monte Carlo techniques for the estimation of the sensitivities (see Broadie and Glasserman (1996)) are revisited and, after suitably adapting to non-Gaussian Lévy log-returns, we run numerical experiments for comparison with the backward convolution technique.

The contribution of Chapter 6 is twofold. Firstly, we extend the valuation scheme of Chapter 4 to two dimensions to allow for non-Lévy log-returns with stochastic volatility. Secondly, we derive the exact distribution law of the discrete log-geometric average and, subsequently, obtain the price of the geometric Asian option in terms of a Fourier transform. We then set up an effective control variate Monte Carlo strategy and use this as a benchmark to the price convolution method.

The need to generate exact sample trajectories for underlyings driven by exponential Lévy models, for the purpose of computing the prices and sensitivities of Asian options in Chapters 4 and 5, has motivated our work in Chapter 7. Building on an idea of Broadie and Kaya (2006) originally implemented for Heston's stochastic volatility model, a Monte Carlo scheme is set up and tailored here to Lévy models, coupled with Fourier-inversion of the associated characteristic functions to recover the implicitly known cumulative distribution functions to sample from. As an example, we consider the family of tempered stable models, which allow for processes of finite or infinite activity and variation, and whose simulation so far has proved problematic. In particular, the CGMY subclass, named after Carr et al. (2002), is investigated against two alternative simulation methods by Madan and Yor (2008) and Poiriot and Tankov (2006) in pricing European-type vanilla and Asian options.

Finally, Chapter 8 focuses on the valuation of convertible bonds. CBs traditionally encom-

pass the holder's right for conversion to the issuing firm's stock prior to or at maturity, and features like with-notice premature redemption by the issuer (call-back option), discrete coupons paid prior to conversion and dividends on the issuing firm's stock received post conversion. Such a specification endows the contract with strong early-exercise features and path-dependence, limiting the use of Monte Carlo (see Lvov et al. (2004), Ammann et al. (2008)) and lattice techniques (see Goldman Sachs (1994), Ho and Pfeffer (1996), Takahashi et al. (2001), Davis and Lischka (2002)) for valuation purposes. Firstly, this chapter proposes a step-by-step convolution approach which operates on backward propagation from the CB maturity, while allowing for the abovementioned provisions at the relevant time points, to provide, eventually, the CB price at inception. Secondly, we consider a four-factor model which comprises stochastic interest rates, and further describes the firm value evolution by a diffusion augmented with jumps, subject to random arrival and size, to the effective modelling of the credit risk. To the best of our knowledge, such a setup has not been implemented previously in the CB literature (and also other instruments with early-exercise provisions) due to dimensionality issues affecting standard numerical schemes for partial differential equations (see Brennan and Schwartz (1977), (1980), Carayannopoulos (1996), Tsiveriotis and Fernandes (1998), Zvan et al. (1998), (2001), Takahashi et al. (2001), Barone-Adesi et al. (2003), Bermúdez and Webber (2004)). The proposed convolution procedure is shown to handle flexibly the dimensionality imposed by the chosen market model, while remaining convergent and precise. The effect of the jump diffusion structural approach, as well as the effects of coupons and dividends paid, and the impact of a varying call policy on the computed prices are explored.

Chapter 2

Fourier transforms

2.1 Introduction

This chapter introduces the concept of a Fourier transform in L^1 and provides key theory and results for use in the chapters to follow. Throughout this thesis, we demonstrate that Fourier transforms offer a valuation framework for contingent claims which is generally applicable to underlyings that are driven by affine processes with characteristic functions known in closed form. In Section 2.3, we define the characteristic function of a random variable and interpret this as the Fourier transform of its distribution. Recognizing that many probability distributions are only known through their characteristic functions¹, we illustrate how to retrieve the density and distribution functions via inversion of the characteristic function.

In Section 2.4, we introduce the notion of a discrete Fourier transform and analyze efficient techniques for its computation. We use these techniques in Section 2.5 in order to approximate continuous Fourier transforms by discrete transforms on a grid.

¹Density functions may exist in analytical form for certain Lévy distributions, however they appear rather complicated involving special functions which render their computation cumbersome and slow. This phenomenon is even more pronounced in the case of distribution functions. On the contrary, characteristic functions are usually available in much simpler closed forms.

2.2 Fourier transforms

The simplest class of functions for which the Fourier transform can be introduced is the Lebesgue class L^1 on $\mathbb{R}$.

Definition 1 Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be an absolutely integrable function. The Fourier transform $\mathcal{F}(f) : \mathbb{R} \rightarrow \mathbb{C}$ is given by

$$\mathcal{F}(f)(u) = \int_{\mathbb{R}} e^{iux} f(x) dx.$$

The following proposition on Fourier transforms holds.

Proposition 2 Assume real constants x_0, u_0 . Let f be an absolutely integrable function with Fourier transform $\mathcal{F}(f)$. Then

1. $h_1(x) := f(x + x_0)$ has Fourier transform $\mathcal{F}(h_1)(u) = e^{-iux_0} \mathcal{F}(f)(u)$;
2. $h_2(x) := e^{iu_0x} h_1(x)$ has shifted Fourier transform $\mathcal{F}(h_2)(u - u_0) = \mathcal{F}(h_1)(u)$.

Proof. See Goldberg ((1961), Theorem 3C). ■

From Theorem 2, the equality

$$\mathcal{F}(h_2)(u - u_0) = e^{-iux_0} \mathcal{F}(f)(u)$$

holds trivially, resulting in

$$\mathcal{F}(f)(u) = e^{iu_0x_0} \int_{\mathbb{R}} e^{i(u-u_0)x} e^{iu_0x} f(x + x_0) dx = e^{i(u-u_0)x_0} \int_{\mathbb{R}} e^{i(u-u_0)(x-x_0)} e^{iu_0x} f(x) dx, \quad (2.1)$$

where the second equality follows by a change of variable. Result (2.1) will be revisited in Section 2.5.

2.2.1 Derivatives of a function and their Fourier transforms

The following theorem allows us to express the Fourier transforms of the derivatives of a function in terms of the Fourier transform of the original function. We will use this property to deduce useful expressions for the price sensitivities of contingent claims in Chapter 5.

Theorem 3 *Let f and $\tilde{f}$ be absolutely integrable functions with Fourier transforms $\mathcal{F}(f)$ and $\mathcal{F}(\tilde{f})$ respectively. If $\mathcal{F}(\tilde{f})(u) = (-iu)^n \mathcal{F}(f)(u)$, $n \in \mathbb{N}^*$, then the derivatives*

$$f^{(k)}(x) := \frac{\partial^k f(x)}{\partial x^k}, \quad k = 1, \dots, n$$

are absolutely integrable, and

$$\tilde{f} = f^{(n)}.$$

Proof. See Bochner and Chandrasekharan ((1949), Theorems 15-17). ■

2.2.2 Convolution

The notion of a convolution will form a main premise for the development of the pricing algorithms in Chapters 4, 5, 6, 8.

Definition 4 *The convolution of two functions f_1, f_2 is given by*

$$(f_1 * f_2)(x) := \int_{\mathbb{R}} f_1(x') f_2(x - x') dx' = \int_{\mathbb{R}} f_2(x') f_1(x - x') dx'.$$

A key result about convolutions states that the Fourier transform of a convolution is equal to the product of the Fourier transforms of the functions being convolved. We present this in the following theorem.

Theorem 5 *If f_1, f_2 are absolutely integrable functions, then their convolution $(f_1 * f_2)(x)$ is absolutely integrable and has Fourier transform*

$$\mathcal{F}(f_1 * f_2) = \mathcal{F}(f_1) \mathcal{F}(f_2).$$

Proof. See Bochner and Chandrasekharan ((1949), Theorem 2). ■

Convolutions of functions on $\mathbb{R}^d$ can be defined similarly. Straightforward extension of Theorem 5 to that case also applies (see also Section 2.2.4).

2.2.3 Inverse Fourier transform

We define next the inverse Fourier transform.

Definition 6 *The inverse Fourier transform of $g : \mathbb{R} \rightarrow \mathbb{C}$ is given by*

$$\mathcal{F}^{-1}(g)(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iux} g(u) du := \lim_{c \rightarrow \infty} \frac{1}{2\pi} \int_{-c}^c e^{-iux} g(u) du,$$

whenever the limit on the right-hand side exists for $x \in \mathbb{R}$.

In analogy to (2.1), the equality

$$\mathcal{F}^{-1}(g)(x) = \frac{1}{2\pi} e^{-iu_0(x-x_0)} \int_{\mathbb{R}} e^{-i(u-u_0)(x-x_0)} e^{-iux_0} g(u) du \quad (2.2)$$

holds for real constants x_0, u_0 .

We wish to provide some simple conditions under which a function can be recovered from its Fourier transform.

Theorem 7 *Suppose f is absolutely integrable. On any compact interval where f is continuous and of finite variation, the inverse Fourier transform of $\mathcal{F}(f)$ is well-defined and*

$$f = \mathcal{F}^{-1}(\mathcal{F}(f)).$$

Proof. See Bochner and Chandrasekharan ((1949), Theorems 4, 8). ■

2.2.4 Multi-dimensional Fourier transforms

Fourier transforms are also defined in multiple dimensions (see Bochner and Chandrasekharan (1949), Theorem 31). For instance, in two dimensions, the Fourier transform of an absolutely integrable function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is given by

$$\mathcal{F}(f)(u) = \int_{\mathbb{R}^2} e^{iu^\top x} f(x) dx. \quad (2.3)$$

The inverse transform of $g : \mathbb{R}^2 \rightarrow \mathbb{C}$ is given by

$$\mathcal{F}^{-1}(g)(x) = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} e^{-iu^\top x} g(u) du. \quad (2.4)$$

2.3 Characteristic function

Based on the principles presented in Section 2.2, we provide next several results which will feature prominently in this thesis. This brings us to the definition of the characteristic function of a random variable as the Fourier transform of its distribution.

Definition 8 *The characteristic function of the $\mathbb{R}^d$ -valued random variable X with probability density f_X is the function $\phi_X : \mathbb{R}^d \rightarrow \mathbb{C}$ given by*

$$\phi_X(u) = E(e^{iu^\top X}) = \int_{\mathbb{R}^d} e^{iu^\top x} f_X(x) dx = \mathcal{F}(f_X)(u).$$

The characteristic function of a random variable uniquely characterizes its law; two random variables with the same characteristic function are identically distributed. A characteristic function satisfies $\phi_X(0) = 1$ and is continuous at $u = 0$, so that $\phi_X(u) \neq 0$ in the neighborhood of $u = 0$. Based on this, one can further define the logarithm of ϕ_X : there exists a unique continuous function $\psi_X(u)$ in the neighborhood of $u = 0$ such that

$$\psi_X(0) = 0 \text{ and } \phi_X(u) = e^{\psi_X(u)}.$$

We call ψ_X the cumulant generating function. Note that if $\phi_X(u) \neq 0$ for all u , $\psi_X(u)$ can be extended to the entire $\mathbb{R}^d$. We define the cumulants of X by

$$c_j(X) := \frac{1}{i^j} \frac{\partial^j \psi_X}{\partial u^j}(0); \tag{2.5}$$

for instance

$$\begin{aligned} c_1(X) &= E(X), \\ c_2(X) &= \text{Var}(X) = E((X - E(X))^2), \\ c_3(X) &= E((X - E(X))^3), \\ c_4(X) &= E((X - E(X))^4) - 3\text{Var}(X)^2. \end{aligned}$$

The normalized versions

$$s(X) := \frac{c_3(X)}{c_2(X)^{3/2}}, \quad \kappa(X) := \frac{c_4(X)}{c_2(X)^2}$$

are called respectively the skewness coefficient and excess kurtosis of X (see Cont and Tankov (2004a), Section 2.2.5).

For a univariate random variable X with a continuous distribution, the density function f_X and the cumulative distribution function F_X can be obtained from ϕ_X using the Gurland (1948) and Gil-Pelaez (1951) inversion formulae

$$\begin{aligned} f_X(x) &= \frac{1}{2\pi} \int_{\mathbb{R}} e^{-iux} \phi_X(u) du, \\ F_X(x) &= \frac{1}{2} - \frac{1}{2\pi} \int_{\mathbb{R}} e^{-iux} \frac{\phi_X(u)}{iu} du. \end{aligned} \quad (2.6)$$

By definition, a distribution function $F_X(x)$ does not decay to zero as $x \rightarrow \infty$; thus, according to Definition 1, the absolute integrability condition for the existence of the Fourier transform of F_X is violated. To fix this, Hughett ((1998), Lemma 9) develops for a continuous function F_X the auxiliary function $\tilde{F}_X$ as

$$\tilde{F}_X(x) := F_X(x) - \frac{1}{2}F_X(x - \xi) - \frac{1}{2}F_X(x + \xi), \quad \xi > 0, \quad (2.7)$$

which is well-behaved in the sense that both $\tilde{F}_X$ and its Fourier transform $\mathcal{F}(\tilde{F}_X)$ decay rapidly to zero. The following theorem then holds.

Theorem 9 *Consider a random variable X which obeys to the law of some continuous distribution F_X with finite variance and has characteristic function ϕ_X . Then, the Fourier transform $\mathcal{F}(\tilde{F}_X)$ of the function $\tilde{F}_X$ defined by (2.7) is given by the continuous function*

$$\mathcal{F}(\tilde{F}_X)(u) = \begin{cases} -\frac{1-\cos(u\xi)}{iu} \phi_X(u), & u \neq 0 \\ 0, & u = 0. \end{cases} \quad (2.8)$$

Proof. See Hughett ((1998), Lemma 9). ■

From Theorem 7, $\tilde{F}_X$ is recovered via

$$\tilde{F}_X = \mathcal{F}^{-1}(\mathcal{F}(\tilde{F}_X)).$$

Then, for sufficiently large $\xi > 0$,

$$F_X(x) \approx \tilde{F}_X(x) + \frac{1}{2}. \quad (2.9)$$

Hughett ((1998), Theorem 10) provides precise bound to the error induced by approximation (2.9) for $|x| \leq \frac{1}{2}\xi$.

2.4 Discrete Fourier transforms

Recall from Section 2.2 the definition of the Fourier transform as a continuous integral. In practice, this is computed by a discrete approximation. To this end, we define the discrete Fourier transform and discuss efficient ways to implement this using the fast Fourier transform algorithm.

Definition 10 *The forward (standard) discrete Fourier transform (DFT) $\hat{\mathbf{a}} = \{\hat{a}_j\}_{j=0}^{n-1}$ of some vector $\mathbf{a} = \{a_k\}_{k=0}^{n-1}$ is given by*

$$\hat{a}_j = \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} e^{i\frac{2\pi}{n}jk} a_k, \quad j = 0, \dots, n-1. \quad (2.10)$$

Furthermore, the inverse (standard) discrete Fourier transform (IDFT) $\hat{\mathbf{b}} = \{\hat{b}_j\}_{j=0}^{n-1}$ of some vector $\mathbf{b} = \{b_k\}_{k=0}^{n-1}$ is given by

$$\hat{b}_j = \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} e^{-i\frac{2\pi}{n}jk} b_k, \quad j = 0, \dots, n-1. \quad (2.11)$$

In short, we write

$$\hat{\mathbf{a}} = \text{dft}(\mathbf{a})$$

and

$$\hat{\mathbf{b}} = \text{idft}(\mathbf{b}).$$

The following result between forward and inverse transforms holds.

Proposition 11 Consider vector $\mathbf{a}$. We have that

$$\text{dft}(\text{idft}(\mathbf{a})) = \mathbf{a},$$

$$\text{idft}(\text{dft}(\mathbf{a})) = \mathbf{a}.$$

Proof. See, for example, Černý ((2004), Appendix). ■

We introduce next the notion of circular convolution for vectors.

Definition 12 Consider two n -dimensional vectors $\mathbf{a} = \{a_l\}_{l=0}^{n-1}$, $\mathbf{b} = \{b_k\}_{k=0}^{n-1}$. We define the circular (cyclic) convolution of $\mathbf{a}$ and $\mathbf{b}$ to be the n -dimensional vector

$$\mathbf{c} = \mathbf{a} \circledast \mathbf{b},$$

such that

$$c_j = \sum_{l=j-k \bmod n} a_l b_k, \quad j = 0, \dots, n-1.$$

A useful application of the discrete Fourier transform is in the computation of the circular convolution.

Proposition 13 Let $\mathbf{c} = \mathbf{a} \circledast \mathbf{b}$, where $\mathbf{a}$ and $\mathbf{b}$ are n -dimensional vectors. Then

$$\text{dft}(\mathbf{c}) = \sqrt{n} \text{dft}(\mathbf{a}) \text{dft}(\mathbf{b}).$$

Proof. See, for example, Černý ((2004), Appendix). ■

The standard DFT and IDFT (see Definition 10) can be implemented fast by means of the fast Fourier transform (FFT) algorithm which is readily available in MATLAB. The forward and inverse FFT in MATLAB are called `fft` and `ifft` respectively. Following these, we calculate

$$\begin{aligned} \text{dft}(\mathbf{a}) &= \sqrt{n} \text{ifft}(\mathbf{a}), \\ \text{idft}(\mathbf{a}) &= \frac{1}{\sqrt{n}} \text{fft}(\mathbf{a}), \end{aligned}$$

for n -dimensional vector $\mathbf{a}$. The length n must be selected with care, otherwise the algorithm may take long time to compute, especially if n is a large prime. It is common practice that

lengths are chosen to be powers of two ($n = 2^p$) or highly composite ($n = 2^p 3^q 5^r$), in order to ensure high-speed computation of the discrete Fourier transform via the FFT algorithm. In fact, Černý (2004) demonstrates that FFT of length $n_1 = 2^{p_1} 3^{q_1} 5^{r_1}$ can be faster than FFT of length $n_2 = 2^{p_2}$ even if $n_1 > n_2$, for suitably chosen p_1, q, r and p_2 values. If the original vector size is not of the desired form, the appropriate number of zeros can be added: a powerful result due to Bluestein (1968) states that a discrete Fourier transform of arbitrary size can be rephrased as a circular convolution and, by zero-padding the input vectors of the circular convolution, this can be then decomposed in terms of three discrete transforms of larger size. This technique speeds up substantially the calculation of an originally prime-sized transform, while the first elements of the re-expressed transform remain unchanged.

Rabiner et al. (1969) apply Bluestein's idea in calculating a more general transform that is termed the chirp z -transform.

Definition 14 Consider vector $\mathbf{a} = \{a_k\}_{k=0}^{n-1}$, and parameters $A, W \in \mathbb{C}$. The chirp z -transform $\hat{\mathbf{a}} = \{\hat{a}_j\}_{j=0}^{\hat{n}-1}$ of vector $\mathbf{a}$ is defined as

$$\hat{a}_j = \sum_{k=0}^{n-1} (AW^{-j})^k a_k, \quad j = 0, \dots, \hat{n} - 1. \quad (2.12)$$

In short, we write

$$\hat{\mathbf{a}} = \text{czt}(\mathbf{a}, A, W, \hat{n}).$$

Next, we explain how the chirp z -transform can be implemented efficiently using Bluestein's decomposition. Let $\hat{n} = n$. Note that $2jk = j^2 + k^2 - (j - k)^2$, expression (2.12) then reads

$$\hat{a}_j = W^{-\frac{j^2}{2}} \sum_{l=|j-k|} b_l c_k, \quad j = 0, \dots, n - 1, \quad (2.13)$$

where

$$\begin{aligned} b_l &= W^{\frac{l^2}{2}}, \quad l = 0, \dots, n - 1, \\ c_k &= W^{-\frac{k^2}{2}} A^k a_k, \quad k = 0, \dots, n - 1. \end{aligned}$$

Bluestein observes that the sum (2.13) can be reformulated as a circular convolution. To this end, we choose $m \geq 2n - 1$. We extend the vectors $\mathbf{b}$ and $\mathbf{c}$ of length n to $\mathbf{b}^*$ and $\mathbf{c}^*$ of length m , such that

$$b_l^* = \begin{cases} W^{\frac{l^2}{2}}, & l = 0, \dots, n-1, \\ 0, & l = n, \dots, m-n, \\ W^{\frac{(l-m)^2}{2}}, & l = m-n+1, \dots, m-1, \end{cases}$$

and

$$c_k^* = \begin{cases} W^{-\frac{k^2}{2}} A^k a_k, & k = 0, \dots, n-1, \\ 0, & k = n, \dots, m-1. \end{cases}$$

Then,

$$\hat{a}_j^* = W^{-\frac{j^2}{2}} d_j^*, \quad j = 0, \dots, m-1 \quad (2.14)$$

with

$$d_j^* = \sum_{l=j-k \bmod m} b_l^* c_k^*, \quad j = 0, \dots, m-1, \quad (2.15)$$

forms a m -point circular convolution. By virtue of Propositions 11 and 13, (2.14-2.15) can be evaluated by utilizing three standard m -point (I)DFTs

$$d^* = \sqrt{m} \text{idft}(\text{dft}(\mathbf{b}^*) \text{dft}(\mathbf{c}^*)). \quad (2.16)$$

By construction,

$$\hat{a}_j = \hat{a}_j^*, \quad j = 0, \dots, n-1.$$

Bailey and Swarztrauber (1991) discuss a special case of the chirp z -transform with $A = 1$ and $|W| = 1$, where $W = e^{-i\frac{2\pi}{n}\alpha}$ for some fractionality coefficient $|\alpha| \in (0, 1]$. They call this the fractional discrete Fourier transform.

Definition 15 Consider vector $\mathbf{a} = \{a_k\}_{k=0}^{n-1}$, and fractionality coefficient $|\alpha| \in (0, 1]$. The fractional discrete Fourier transform $\hat{\mathbf{a}} = \{\hat{a}_j\}_{j=0}^{n-1}$ of vector $\mathbf{a}$ is given by

$$\hat{a}_j = \sum_{k=0}^{n-1} e^{i\frac{2\pi}{n}\alpha jk} a_k, \quad j = 0, \dots, n-1. \quad (2.17)$$

We write

$$\hat{\mathbf{a}} = \text{frft}(\mathbf{a}, \alpha).$$

Let vector $\mathbf{a}$ be of dimension n . By comparing results (2.12) and (2.17), we relate the fractional DFT to the chirp z -transform via

$$\text{frft}(\mathbf{a}, \alpha) = \text{czt}(\mathbf{a}, 1, e^{-i\frac{2\pi}{n}\alpha}, n).$$

By further comparing (2.17) with (2.10) and (2.11), we observe that the fractional DFT reduces to the standard DFT and IDFT when $|\alpha| = 1$, such that

$$\text{dft}(\mathbf{a}) = \frac{1}{\sqrt{n}} \text{frft}(\mathbf{a}, 1), \quad (2.18)$$

$$\text{idft}(\mathbf{a}) = \frac{1}{\sqrt{n}} \text{frft}(\mathbf{a}, -1). \quad (2.19)$$

In this thesis, we will focus on applications of the last three special cases of the chirp z -transform.

2.5 Fourier transforms on a grid

In what follows, we will apply the standard (I)DFT and fractional DFT studied in the previous section to compute continuous Fourier transforms by discrete analogues.

Consider the continuous Fourier transform $\mathcal{F}(f)(u)$ given by (2.1). In practice, this is approximated on a uniform grid $\mathbf{u} = \{u_0 + j\delta u\}_{j=0}^{n-1}$ as

$$e^{i(u-u_0)x_0} \int_{\mathbb{R}} e^{i(u-u_0)(x-x_0)} e^{iu_0x} f(x) dx \approx \mathcal{D}(\mathbf{f}, \mathbf{x}, \mathbf{u}; \alpha) \delta x, \quad (2.20)$$

where

$$\mathcal{D}(\mathbf{f}, \mathbf{x}, \mathbf{u}; \alpha) := e^{i(\mathbf{u}-u_0)x_0} \sum_{k=0}^{n-1} e^{i(\mathbf{u}-u_0)(\mathbf{x}_k-x_0)} e^{iu_0\mathbf{x}_k} f_k, \quad (2.21)$$

and $\mathbf{f} = \{f_k\}_{k=0}^{n-1} := \{f(\mathbf{x}_k)\}_{k=0}^{n-1}$ is the collection of the values of function f on a uniform grid $\mathbf{x} = \{x_0 + k\delta x\}_{k=0}^{n-1}$. Substituting for $\mathbf{x}$ and $\mathbf{u}$ in (2.21) yields

$$\mathcal{D}(\mathbf{f}, \mathbf{x}, \mathbf{u}; \alpha) = e^{i(\mathbf{u}-u_0)x_0} \sum_{k=0}^{n-1} e^{i\frac{2\pi}{n}\alpha jk} e^{iu_0\mathbf{x}_k} f_k; \quad \alpha = \frac{n\delta u\delta x}{2\pi}. \quad (2.22)$$

Similarly, for function values $\mathbf{g} = \{g_j\}_{j=0}^{n-1} := \{g(\mathbf{u}_j)\}_{j=0}^{n-1}$, we approximate the inverse transform $\mathcal{F}^{-1}(g)(x)$ in (2.2) on grid $\mathbf{x} = \{x_0 + k\delta x\}_{k=0}^{n-1}$ as

$$\frac{1}{2\pi} e^{-iu_0(x-x_0)} \int_{\mathbb{R}} e^{-i(u-u_0)(x-x_0)} e^{-iu x_0} g(u) du \approx \frac{1}{2\pi} \mathcal{D}(\mathbf{g}, -\mathbf{u}, \mathbf{x}; -\alpha) \delta u.$$

Proposition 16 *Consider the uniform grids $\mathbf{x} = \{x_0 + k\delta x\}_{k=0}^{n-1}$ and $\mathbf{u} = \{u_0 + j\delta u\}_{j=0}^{n-1}$, the parameter $\alpha = \frac{n\delta u\delta x}{2\pi}$, and the vectors $\mathbf{f} = \{f_k\}_{k=0}^{n-1} = \{f(\mathbf{x}_k)\}_{k=0}^{n-1}$ and $\mathbf{g} = \{g_j\}_{j=0}^{n-1} = \{g(\mathbf{u}_j)\}_{j=0}^{n-1}$. The following identities hold*

$$\mathcal{D}(\mathbf{f}, \mathbf{x}, \mathbf{u}; \alpha) = e^{i(\mathbf{u}-u_0)x_0} \text{frft}(e^{iu_0\mathbf{x}}\mathbf{f}, \alpha), \quad (2.23)$$

$$\mathcal{D}(\mathbf{g}, -\mathbf{u}, \mathbf{x}; -\alpha) = e^{-iu_0(\mathbf{x}-x_0)} \text{frft}(e^{-i\mathbf{u}x_0}\mathbf{g}, -\alpha). \quad (2.24)$$

For $\alpha = 1$, relations (2.23-2.24) reduce to

$$\mathcal{D}(\mathbf{f}, \mathbf{x}, \mathbf{u}; 1) = \sqrt{n} e^{i(\mathbf{u}-u_0)x_0} \text{dft}(e^{iu_0\mathbf{x}}\mathbf{f}), \quad (2.25)$$

$$\mathcal{D}(\mathbf{g}, -\mathbf{u}, \mathbf{x}; -1) = \sqrt{n} e^{-iu_0(\mathbf{x}-x_0)} \text{idft}(e^{-i\mathbf{u}x_0}\mathbf{g}), \quad (2.26)$$

respectively.

Proof. Result (2.23) follows from (2.22) and (2.17). Result (2.24) follows directly from (2.23). Results (2.25) and (2.26) follow from (2.18) and (2.19) respectively. ■

Using the fractional DFT grants us enough flexibility to determine all the three input elements δu , δx , n independently. Instead, the use of the standard (I)DFT is limited by the restriction imposed by $\frac{n\delta u\delta x}{2\pi} = 1$. This has a detrimental effect on the convergence of the right-hand side in (2.20) as n grows. We illustrate this via a practical example in Section 4.6.4. However, for small sizes n , the (I)DFT can still provide results of comparable accuracy at reduced CPU effort; given the decomposition of a n -point fractional DFT into three $2n$ -point standard (I)DFTs, we anticipate an approximately six times faster FFT-implementation of a single n -point DFT. The advantage of the operation-saving (I)DFTs becomes even more pronounced in multi-dimensional settings.

2.5.1 Multi-dimensional discrete approximations

Multi-dimensional Fourier transforms can be approximated by discrete transforms, too. For the purposes of this thesis, we focus on the two-dimensional case.

Definition 17 *The forward discrete Fourier transform $\hat{\mathbf{A}} = \{\hat{A}_{j_1, j_2}\}_{n_1 \times n_2}$ of $\mathbf{A} = \{A_{k_1, k_2}\}_{n_1 \times n_2}$ is given by*

$$\hat{A}_{j_1, j_2} = \frac{1}{\sqrt{n_1 n_2}} \sum_{k_1, k_2} e^{i \frac{2\pi}{n_1} j_1 k_1 + i \frac{2\pi}{n_2} j_2 k_2} A_{k_1, k_2},$$

$$j_1 = 0, \dots, n_1 - 1, j_2 = 0, \dots, n_2 - 1.$$

The inverse discrete Fourier transform $\hat{\mathbf{B}} = \{\hat{B}_{j_1, j_2}\}_{n_1 \times n_2}$ of $\mathbf{B} = \{B_{k_1, k_2}\}_{n_1 \times n_2}$ is given by

$$\hat{B}_{j_1, j_2} = \frac{1}{\sqrt{n_1 n_2}} \sum_{k_1, k_2} e^{-i \frac{2\pi}{n_1} j_1 k_1 - i \frac{2\pi}{n_2} j_2 k_2} B_{k_1, k_2},$$

$$k_1 = 0, \dots, n_1 - 1, k_2 = 0, \dots, n_2 - 1.$$

In short, we write

$$\hat{\mathbf{A}} = \text{dft2}(\mathbf{A})$$

and

$$\hat{\mathbf{B}} = \text{idft2}(\mathbf{B}).$$

Given a $n_1 \times n_2$ matrix $\mathbf{A}$, we employ the MATLAB functions `ifft2` and `fft2` to compute

$$\begin{aligned} \text{dft2}(\mathbf{A}) &= \sqrt{n_1 n_2} \text{ifft2}(\mathbf{A}), \\ \text{idft2}(\mathbf{A}) &= \frac{1}{\sqrt{n_1 n_2}} \text{fft2}(\mathbf{A}). \end{aligned}$$

Consider the uniform grids $\mathbf{x}_i = \{x_{i,0} + k_i \delta x_i\}_{k_i=0}^{n_i-1}$ and $\mathbf{u}_i = \{u_{i,0} + j_i \delta u_i\}_{j_i=0}^{n_i-1}$, $i = 1, 2$. Suppose the values of the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ are given on the two-dimensional grid $\mathbf{x}^* = (\mathbf{x}_1, \mathbf{x}_2)$; we summarize these in the matrix $\mathbf{F} = \{F_{k_1, k_2}\}_{n_1 \times n_2} := \{f(\mathbf{x}_{k_1, k_2}^*)\}_{n_1 \times n_2}$. We approximate the continuous Fourier integral (2.3) by its discrete analogue on grid $\mathbf{u}^* = (\mathbf{u}_1, \mathbf{u}_2)$ as

$$\int_{\mathbb{R}^2} e^{i\mathbf{u}^\top \mathbf{x}} f(\mathbf{x}) d\mathbf{x} \approx \mathcal{D}(\mathbf{F}, \mathbf{x}^*, \mathbf{u}^*) \delta x_1 \delta x_2,$$

where

$$\mathcal{D}(\mathbf{F}, \mathbf{x}^*, \mathbf{u}^*) = \sqrt{n_1 n_2} e^{i(\mathbf{u}_1^\top - u_{1,0})x_{1,0}} e^{i(\mathbf{u}_2 - u_{2,0})x_{2,0}} \text{dfft2}(e^{iu_{1,0}\mathbf{x}_1^\top} e^{iu_{2,0}\mathbf{x}_2} \cdot \mathbf{F}) \quad (2.27)$$

is an extended version of the conversion rule (2.25) to two dimensions. We denote by $\cdot$ the element-wise matrix multiplication. Similarly, for matrix $\mathbf{G} = \{G_{j_1, j_2}\}_{n_1 \times n_2} := \{g(\mathbf{u}_{j_1, j_2}^*)\}_{n_1 \times n_2}$ representing the values of $g : \mathbb{R}^2 \rightarrow \mathbb{C}$ on grid $\mathbf{u}^*$, we approximate the inverse transform (2.4) on grid $\mathbf{x}^*$ as

$$\frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} e^{-i\mathbf{u}^\top x} g(u) du \approx \frac{1}{(2\pi)^2} \mathcal{D}(\mathbf{G}, -\mathbf{u}^*, \mathbf{x}^*) \delta u_1 \delta u_2,$$

where

$$\mathcal{D}(\mathbf{G}, -\mathbf{u}^*, \mathbf{x}^*) = \sqrt{n_1 n_2} e^{-iu_{1,0}(\mathbf{x}_1^\top - x_{1,0})} e^{-iu_{2,0}(\mathbf{x}_2 - x_{2,0})} \text{idft2}(e^{-i\mathbf{u}_1^\top x_{1,0}} e^{-i\mathbf{u}_2 x_{2,0}} \cdot \mathbf{G}). \quad (2.28)$$

Chapter 3

Affine models and option pricing

3.1 Introduction

Affine models have been used traditionally in mathematical finance to model random log-asset price, interest rate and volatility movements. The motivation behind using exponential affine asset price models (excluding Samuelson's geometric Brownian motion) stems from their ability to fit flexibly on empirical observations, and reproduce the volatility skew and smile common amongst other stylized empirical facts in the financial markets (see Cont and Tankov (2004a), Chapter 7 for details). Furthermore, closed-form expressions for the characteristic functions are available for a number of affine models, making them highly tractable in financial applications. Particularly popular is the subclass of Lévy processes which includes the Gaussian diffusion (arithmetic Brownian motion), the jump diffusion models of Merton (1976) and Kou (2002), and pure jump models such as the variance gamma (VG) of Madan and Seneta (1990), Madan and Milne (1991) and Madan et al. (1998), the normal inverse Gaussian (NIG) of Barndorff-Nielsen (1998), and the tempered stable (KoBoL/CGMY) of Koponen (1995), Boyarchenko and Levendorskiĭ (2002), and Carr et al. (2002). Hybrid models with stochastic volatility (e.g., Heston (1993)) and/or stochastic interest rates (e.g., see Chapter 8 of this thesis) also appear in the affine class. Members of the affine class will be used frequently in the next chapters.

Section 3.2 briefly describes the key features and theoretical results relevant to the Lévy processes, and subsequently Section 3.3 focuses on the more general affine processes. It demonstrates how the characteristic function can be derived for this class of models and, by recalling

its interpretation as a Fourier transform, Section 3.4 reviews how European options can be evaluated via inversion of the characteristic function. Popular applications of Fourier transforms to the pricing of options with path-dependence and/or early-exercise features are also discussed.

3.2 Lévy processes

Consider a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \geq 0}, P)$. Then, a càdlàg, $\mathbb{F}$ -adapted, $\mathbb{R}$ -valued process L_t , with $L_0 = 0$, is called a Lévy process if

1. it has independent increments;
2. it has stationary increments;
3. it is stochastically continuous, i.e., for any $t \geq 0$ and $\varepsilon > 0$ we have

$$\lim_{s \rightarrow t} P(|L_t - L_s| > \varepsilon) = 0.$$

Each Lévy process can be characterized by a triplet (μ, σ^2, Π) referred to as the Lévy characteristics of L with drift parameter $\mu \in \mathbb{R}$, diffusion parameter $\sigma \geq 0$, and Lévy density Π satisfying $\Pi(0) = 0$ and $\int_{\mathbb{R} \setminus \{0\}} (1 \wedge |l|^2) \Pi(dl) < \infty$. The local characteristics are given by $(\mu t, \sigma^2 t, \Pi(dl) dt)$. In terms of this triplet, the characteristic function of the Lévy process is given by the celebrated Lévy-Khintchine formula

$$E(e^{iuL_t}) = e^{\psi_L(u)t},$$

where the Lévy exponent ψ_L of the process L is represented as

$$\psi_L(u) = iu\mu - \frac{1}{2}\sigma^2 u^2 + \int_{\mathbb{R} \setminus \{0\}} (e^{iul} - 1 - iul1_{\{|l| \leq 1\}}) \Pi(dl). \quad (3.1)$$

Consider the price of a risky asset S to evolve according to an exponential Lévy model: fix constant $S_0 > 0$ and define

$$S_t = S_0 e^{L_t}$$

under the risk-neutral measure $\mathbb{P}$, i.e., $P = \mathbb{P}$. Assume further the existence of a money market account M which evolves according to $dM_t = rM_t dt$ with $M_0 = 1$ and $r > 0$ being the continuously compounded risk-free interest rate. Then, to satisfy the fundamental theorem of asset pricing (see Delbaen and Schachermayer (1994)), we must guarantee that $S_t/M_t = S_t e^{-rt}$ is a $\mathbb{P}$ -martingale, i.e.,

$$\mathbb{E}(S_t e^{-rt}) = \mathbb{E}(S_0 e^{-rt+L_t}) = S_0,$$

which implies that μ in (3.1) must be chosen to satisfy $\psi_L(-i) = r$. Hence,

$$\mu = r - \frac{1}{2}\sigma^2 - \int_{\mathbb{R} \setminus \{0\}} (e^l - 1 - l1_{\{|l| \leq 1\}}) \Pi(dl).$$

3.3 Affine processes

This class of models has been explored in great detail in Duffie et al. (2000) and Duffie et al. (2003). Following their lead, we consider the d -dimensional diffusion process $\mathbf{X} = (X_1, \dots, X_d)$ in some state space $D \subset \mathbb{R}^d$ satisfying

$$d\mathbf{X}_t = \mu(\mathbf{X}_t) dt + \sigma(\mathbf{X}_t) d\mathbf{W}_t$$

with $\mu : D \rightarrow \mathbb{R}^d$, $\sigma : D \rightarrow \mathbb{R}^{d \times d}$ and $\mathbf{W}$ a standard Brownian motion on $\mathbb{R}^d$. Process $\mathbf{X}$ is called affine if and only if

$$\mu(x_1, \dots, x_d) = m_0 + \sum_{j=1}^d x_j m_j, \text{ for } m_j \in \mathbb{R}^d, \quad (3.2)$$

$$\sigma(x_1, \dots, x_d) \sigma(x_1, \dots, x_d)^\top = s_0 + \sum_{j=1}^d x_j s_j, \text{ for } s_j \in \mathbb{R}^{d \times d}. \quad (3.3)$$

Then Duffie et al. (2003) uniquely characterize the distribution law of a regular affine process $\mathbf{X}$ by the characteristic function

$$E(e^{i\mathbf{u}^\top \mathbf{X}_t}) = e^{\Psi_0(\mathbf{u};t) + \Psi_{(1,\dots,d)}(\mathbf{u};t)^\top \mathbf{X}_0}, \quad \mathbf{u} \in \mathbb{C}^d,$$

where Ψ_0 and $\Psi_{(1,\dots,d)} = (\Psi_1, \dots, \Psi_d)$ are respectively $\mathbb{C}$ - and $\mathbb{C}^d$ -valued functions solving the following system of generalized Riccati equations

$$\frac{\partial \Psi_j(\mathbf{u}; t)}{\partial t} = \psi_j(-i\Psi_{(1,\dots,d)}(\mathbf{u}; t)), \quad j = 0, \dots, d, \quad (3.4)$$

$$\Psi_0(\mathbf{u}; 0) = 0, \quad \Psi_{(1,\dots,d)}(\mathbf{u}; 0) = i\mathbf{u}, \quad (3.5)$$

with

$$\psi_j(\mathbf{u}) = i\mathbf{u}^\top m_j - \frac{1}{2}\mathbf{u}^\top s_j \mathbf{u}.$$

In fact, Ψ_0 is determined by $\Psi_{(1,\dots,d)}$ via the simple integration

$$\Psi_0(\mathbf{u}; t) = \int_0^t \psi_0(-i\Psi_{(1,\dots,d)}(\mathbf{u}; s)) ds.$$

In some applications, explicit solutions for the complex-valued ODEs (3.4-3.5) can be found, whereas in others, solutions may only be found numerically by using, for example, the Runge-Kutta method.

We consider next two examples from the affine diffusion class which will feature in this thesis: the Heston stochastic volatility model (Chapter 6) and the asset diffusion with Vaříček stochastic interest rate hybrid model (Chapter 8). In the Heston framework the stochastic volatility is modelled by the same mean-reverting square-root process that is used for the interest rate in Cox et al. (1985). In particular, the model is of the form

$$\begin{aligned} dX_t &= (\mu - v_t/2) dt + \sqrt{v_t} dW_t, \\ dv_t &= \alpha(\beta - v_t) dt + \eta\sqrt{v_t} dW_{v,t}, \end{aligned}$$

where W and W_v are correlated standard Brownian motions with constant correlation ρ , and μ, α, β, η are constant parameters. In these SDEs, $X = \ln S$ represents the logarithm of the asset S with stochastic volatility $\sqrt{v}$. Viewing jointly (v, X) as the state variables yields an

affine process in the spirit of Duffie et al. (2000) and Duffie et al. (2003). In particular, we get

$$\begin{aligned} (m_0, s_0) &= \left(\begin{pmatrix} \alpha\beta \\ \mu \end{pmatrix}, 0 \right), \\ (m_1, s_1) &= \left(\begin{pmatrix} -\alpha \\ -\frac{1}{2} \end{pmatrix}, \begin{pmatrix} \eta^2 & \eta\rho \\ \rho\eta & 1 \end{pmatrix} \right), \\ (m_2, s_2) &= (0, 0). \end{aligned}$$

In the second example, we assume Gaussian interest rate movements according to the model by Vašíček (1977) and constant asset volatility. Denoting by X and r the log-asset price and interest rate processes respectively, this hybrid model takes form

$$dX_t = (r_t - \sigma^2/2) dt + \sigma dW_t, \quad (3.6)$$

$$dr_t = \kappa(\mu_r - r_t) dt + \sigma_r dW_{r,t}, \quad (3.7)$$

where W and W_r are correlated standard Brownian motions, ρ is the constant correlation coefficient, and parameters σ , κ , μ_r , σ_r are constant. The form of the bivariate model (r, X) agrees with the affine structure (3.2-3.3) such that

$$\begin{aligned} (m_0, s_0) &= \left(\begin{pmatrix} \kappa\mu_r \\ -\frac{1}{2}\sigma^2 \end{pmatrix}, \begin{pmatrix} \sigma_r^2 & \sigma_r\sigma\rho \\ \sigma\sigma_r\rho & \sigma^2 \end{pmatrix} \right), \\ (m_1, s_1) &= \left(\begin{pmatrix} -\kappa \\ 1 \end{pmatrix}, 0 \right), \\ (m_2, s_2) &= (0, 0). \end{aligned}$$

In terms of mathematical tractability, both settings are advantageous since they permit closed-form characteristic functions as we illustrate in Chapters 6 and 8.

3.4 Option pricing with Fourier transforms

Heston (1993) provides the first well-established pricing model in the literature for European plain vanilla options based on Fourier-inversion of the characteristic function of the

log-increment of an underlying asset with stochastic volatility. Heston's approach allows for correlated variance and asset price processes, in contrast to the previous attempt by Stein and Stein (1991) which heavily relies on independent processes, and is further applicable to any model for the log-asset dynamics providing that the characteristic function is known in closed form. Since Heston's seminal paper, the pricing of derivative contracts using Fourier-inversion techniques has raised the interest of several authors due to their reported accuracy and speed.

Consider a generic European contingent claim C written on a single asset S , with maturity T , and terminal payoff $p(X_T)$ where $X_T = \ln S_T$. From the fundamental theorem of asset pricing, the forward price of the contingent claim is given under the risk-neutral measure $\mathbb{P}$ by

$$C_0 = \mathbb{E}(p(X_T)). \quad (3.8)$$

Heston (1993) deals with the pricing problem (3.8) by decomposing the expectation in terms of cumulative probabilities of the underlying asset. We illustrate his approach by applying on a plain vanilla call option with payoff $p(X_T) = (e^{X_T} - K)^+$, where $K > 0$ is the strike price. The forward price of the option reads

$$C_0 = e^{X_0 + rT} \bar{\mathbb{P}}(X_T > k) - K \mathbb{P}(X_T > k), \quad (3.9)$$

where $k = \ln K$ and $r > 0$ is the risk-free interest rate. $\bar{\mathbb{P}}$ indicates the equivalent spot measure induced by taking S_T as the numéraire (see Appendix 6.B). Using the Gil-Pelaez formula (2.6), both probabilities in (3.9) are given by

$$\begin{aligned} \bar{\mathbb{P}}(X_T > k) &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} e^{-iuk} \frac{\phi_{X_T}(u-i)}{iu\phi_{X_T}(-i)} du, \\ \mathbb{P}(X_T > k) &= \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} e^{-iuk} \frac{\phi_{X_T}(u)}{iu} du, \end{aligned}$$

where

$$\phi_{X_T}(u) = \mathbb{E}(e^{iuX_T})$$

is the characteristic function of the log-asset price X_T . Note that method (3.9) requires that we calculate two inverse transforms.

Carr and Madan (1999) suggest an alternative representation for the forward price of a European call option in terms of its Fourier transform with respect to the log-strike k . In particular, they express the option payoff in the integral form

$$(e^x - e^k)^+ = \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} e^{-iuk} \frac{e^{x(iu+1)}}{iu(iu+1)} du,$$

for arbitrary constant $R < 0$. The risk-neutral pricing formula (3.8) then reads

$$C_0(k) = \mathbb{E} \left((e^{X_T} - e^k)^+ \right) = \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} e^{-iuk} \frac{\phi_{X_T}(u-i)}{iu(iu+1)} du. \quad (3.10)$$

Constant R is chosen to ensure that $C_0(k)$ is absolutely integrable as $k \rightarrow -\infty$ (see Carr and Madan (1999), Section 3).

Raible (2000) and Lewis (2001) consider a similar approach to Carr and Madan (1999), except that they express the (forward) option price in terms of Fourier transforms taken with respect to the log-forward and log-spot prices respectively. Their approach is general in that it can be adapted to a wide range of European payoff functions, providing their Fourier transforms exist. Following them, we represent continuous payoff functions $p(X_T)$ in the integral form

$$p(X_T) = \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} e^{-iuX_T} \mathcal{F}(p)(u) du, \quad (3.11)$$

where $\mathcal{F}(p)(u)$ is the Fourier transform of the payoff function with respect to X_T , in consistency with Definition 1. For example, $\mathcal{F}(p)(u) = \frac{K^{iu+1}}{iu(iu+1)}$ with $R \in (1, \infty)$ corresponds to a plain vanilla call option with payoff $p(X_T) = (e^{X_T} - K)^+$. Substituting the integral (3.11) for $p(X_T)$ in the risk-neutral valuation formula (3.8) yields

$$C_0(X_0) = \mathbb{E}(p(X_T)) = \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} \phi_{X_T}(-u) \mathcal{F}(p)(u) du. \quad (3.12)$$

In practice, the three price representations (3.9), (3.10) and (3.12) are calculated using discrete approximations for the continuous Fourier transforms as discussed in the previous chapter. In particular, Carr and Madan (1999) compute (3.10) on a grid of strikes by utilizing a standard DFT approximation, whereas Chourdakis (2004) implements a fractional DFT approximation. The same numerical techniques can also be used to compute (3.12). A more recent technique

by Fang and Oosterlee (2008a) replaces DFTs by Fourier-cosine series expansions truncated at $N < \infty$ points. For a European call/put option with payoff function $(\pm e^x \mp K)^+$, the approximate pricing formula reads

$$C_0 \approx \sum_{j=0}^{N-1} V_j(\zeta, K) \operatorname{Re} e^{-i \frac{jL\pi}{U-L}} \phi_{X_T} \left(\frac{j\pi}{U-L}; X_0 \right), \quad (3.13)$$

where $\sum'$ indicates that the first term in the summation is weighted by a half. ζ takes value 1 (-1) for a call (put) option. Coefficients V_j are given by

$$V_j(\zeta, K) = \frac{2}{U-L} \int_L^U (\zeta (e^x - K))^+ \cos \left(j\pi \frac{x-L}{U-L} \right) dx$$

for some closed interval $[L, U] \subset \mathbb{R}$. Coefficients $\{V_j\}_{j=0}^{N-1}$ exist in closed form for payoffs of the form $(\zeta (e^x - K))^+$ (see Fang and Oosterlee (2008a), Section 3). Fang and Oosterlee (2008a) prove that the error of the Fourier-cosine series approximation (3.13) decays exponentially to zero as N grows to infinity for smooth probability densities, whereas Carr and Madan's implementation of (3.10) by DFT achieves only fourth-order convergence.

In addition to plain vanilla options, exotic contracts with possible early exercise and/or path-dependence form another important class. Carverhill and Clewlow (1990) are the first to make use of FFT applications in pricing arithmetic average (Asian) options written on some asset S , whose price is recorded at points $\{t_j\}_{j=1}^n$ on the time line $[0, T]$, with $t_0 = 0$ and $t_n = T$. In general terms, their scheme retrieves the density of the terminal average asset price by employing, on a reduced state space, forward recursive convolutions of the density of the running average at a monitoring date t_{j-1} with the density of the asset log-return $Z_j = X_j - X_{j-1} = \ln S_j - \ln S_{j-1}$ over the next sub-period, until maturity. We refer to Section 4.4 for more details. Instead, Eydeland (1994) computes backward sequential convolutions of the payoff function $\bar{C}_j$ at a monitoring date t_j with the density of the asset log-return Z_j , to provide the option value function C_{j-1} at t_{j-1} . Ultimately, the convolution algorithm provides the option value at t_0 . For example, a Bermudan vanilla call option with strike K has payoff function $\bar{C}_j(X_j) = \max(C_j(X_j), e^{X_j} - K)$ at t_j , $1 \leq j < n$, and $\bar{C}_n(X_n) = \max(e^{X_n} - K, 0)$ at t_n . For n equidistant monitoring dates with step size δt , the problem is formulated under

Lévy log-returns as

$$C_{j-1}(X_{j-1}) = e^{-r\delta t} \int_{\mathbb{R}} \bar{C}_j(X_{j-1} + z) f_j(z) dz, \quad j = 1, \dots, n, \quad (3.14)$$

where $r > 0$ denotes the risk-free interest rate, and f_j the risk-neutral density of the asset log-return Z_j . Convolution (3.14) is then discretized and extended to a circular convolution for efficient computation using Bluestein's decomposition (2.16). To skip explicit use of the densities of the log-returns, Lord et al. (2008) directly compute the Fourier transform of the convolution (3.14), and subsequently invert this to obtain the option value at t_{j-1} . More specifically, following Theorems 5 and 7, they express

$$C_{j-1} = \mathcal{F}^{-1}(\mathcal{F}(\bar{C}_j)\varphi_j),$$

where $\varphi_j = \mathbb{E}(e^{-iuZ_j}) = \int_{\mathbb{R}} e^{-iuz} f_j(z) dz$. Whereas Lord et al. (2008) employ standard (I)DFTs to approximate the continuous Fourier transforms subject to polynomially decaying errors, Fang and Oosterlee (2008b) achieve exponentially decaying errors by utilizing Fourier-cosine series expansions. Alternatively, Feng and Linetsky (2008) and Feng and Lin (2009) replace Fourier transforms by Hilbert transforms, and subsequently develop discrete approximations based on fast-convergent Sinc expansions for faster convergence. It is also worth to mention the method by Broadie and Yamamoto (2005), which combines the double exponential quadrature rule and the fast Gauss transform introduced earlier in Broadie and Yamamoto (2003). Their technique is remarkably fast, as it is linear in both the monitoring dates n and the number of quadrature points N used in the log-asset price dimension, i.e., it is of order $O(nN)$, in contrast to all the other techniques with computational complexity $O(nN \log_2 N)$. Still, the applicability of Broadie and Yamamoto's method is restricted to return distributions which are mixtures of independent Gaussians, as in the Merton jump diffusion in addition to the Black-Scholes-Merton model.

The backward techniques for exotic products discussed above have been employed to evaluate mainly discrete barrier and Bermudan vanilla options. The case of discrete arithmetic Asian options is more complicated, since the value function at each monitoring date not only depends on the state of the underlying asset, but also the running average level; this raises

substantially the dimensionality of the pricing problem. In the next chapter, we exploit the state-space reduction of Carverhill and Clewlow (1990) to construct an efficient backward price convolution for discrete arithmetic Asian options.

Chapter 4

A backward convolution algorithm for discretely sampled arithmetic Asian options

4.1 Introduction

First introduced in the Tokyo Stock Exchange, Asian options are path-dependent derivatives whose payoff depends on the average price of the underlying asset monitored over a predetermined period of time. The fact that the averaging reduces the impact of the volatility of the underlying asset leads to lower option prices, rendering these favourable to the traders. Also, the dependence of the option payout on the average value of the underlying, rather than a single snapshot, makes them more robust against price manipulation. For more on the history and evolution of Asian options, we refer to Boyle and Boyle (2001).

The way the average is defined (geometric versus arithmetic and discrete versus continuous monitoring) plays a critical role in the analytical tractability of the option. Although a true pricing formula exists in closed form for the geometric average under the standard Black-Scholes-Merton market assumptions (see Kemna and Vorst (1990), Conze and Viswanathan (1991), Turnbull and Wakeman (1991)), this is not the case for the more popular discrete and

Chapter 4 draws heavily on the forthcoming paper Černý and Kyriakou (2010).

arithmetic average. For this, several approaches have been proposed in the literature, including analytical approximations, numerical integration methods, partial differential equations (PDEs) and Monte Carlo simulation. The first category encompasses analytical expressions resulting from approximations of the distribution of the average by fitting different distributions, analytical representations in terms of (e.g., Laplace) transformed functionals that require numerical evaluation, and lower and upper bounds for the option price. So far, numerical integration methods have been used to produce the exact density of the average price by forward-in-time recursive integration, either direct or via transforms to the Fourier or Laplace spaces, for use in the computation of the price. Pricing PDEs implemented numerically by finite difference schemes occupy a significant part in the literature, while control variate Monte Carlo strategies are common, especially following recent generalization to any Lévy model for the log-returns, and simple to implement. All four pricing approaches are revisited in more detail in Section 4.2.

We recognize that most Asian options are not monitored continuously, indeed it is typical for the underlying asset value to be recorded at discrete points in time, e.g., on a daily, weekly, monthly basis, etc. With the focus on the numerical (recursive) integration methods, we mention three existing contributions in the literature for discretely sampled arithmetic Asians, Carverhill and Clewlow (1990), Benhamou (2002) and Fusai and Meucci (2008), which have been developed on a reduced state space using the so-called Carverhill-Clewlow-Hodges factorization. The idea of these works is to evaluate the density of the arithmetic average by employing forward recursive density convolutions. All three papers have a significant advantage over the abovementioned approaches in that they can easily be adapted to non-Gaussian Lévy log-returns. In what follows, we illustrate how to replace the forward density convolution by a backward price convolution. We show that this has substantial numerical and theoretical merits.

The remainder of the chapter is structured as follows. Section 4.2 reviews previous contributions in the field of arithmetic Asian options pricing. Section 4.3 presents the Carverhill-Clewlow-Hodges factorization and, given this, Section 4.4 focuses on the existing forward density convolution schemes and their limitations. Section 4.5 develops the main theoretical results for the backward price convolution scheme, and Section 4.5.1 discusses its implementation via dis-

crete Fourier transform. Section 4.6 describes parameterizations of the log-return distribution and illustrates speed-accuracy comparisons of our scheme with previous studies, and Section 4.7 concludes the chapter.

4.2 Pricing approaches to arithmetic Asian options

Even in the simple Black-Scholes-Merton model, arithmetic Asian options do not admit an exact formula in closed form. This is because the sum of correlated lognormal asset prices is not lognormal anymore. With the focus on the continuous arithmetic average asset price, Geman and Yor (1993) are the first to write the price of the Asian option as the inverse of its Laplace transform which they derive in analytical form. Thereafter several authors have attempted to compute the inverse transform using standard numerical approaches, including Fourier series expansion, Laguerre series expansion, sequence of Gaver functionals, and deformation of Bromwich contour (for a thorough review of these techniques, see Davies (2002), Chapter 19), and all have encountered significant numerical instabilities for short maturities and low volatilities (see Dufresne (2000), Linetsky (2004)). These limitations have been attributed to the slow convergence of the inversion algorithms and computational difficulties related to the Kummer confluent hypergeometric function appearing in the Laplace transform. Instead, Fusai (2004) and Cai and Kou (2010) obtain analytical expressions for the double Laplace transform of the option price, which they invert numerically using a two-sided Euler inversion algorithm. Although the two methods share similarities, Cai and Kou's inversion technique is faster, for given accuracy, for low asset volatility, e.g., smaller than 0.1, and performs better under jump diffusion model assumptions.

Other authors choose to approximate the unknown distribution law of the arithmetic (either continuous or discrete) average by fitting different distributions, and subsequently deduce approximate analytical formulae for the option price in the Black-Scholes-Merton economy: Turnbull and Wakeman (1991) and Levy (1992) employ Edgeworth series expansions to approximate the true density of the average with a lognormal density. While this method works well for short maturities, longer maturities have a detrimental effect on the quality of the approximation. Turnbull and Wakeman (1991) also provide an algorithm to compute the moments of the

true distribution of the average. It is observed that the performance of the method is affected when the third and fourth moments differ significantly from the ones implied by the lognormal distribution (the first two are matched by construction).¹ Milevsky and Posner (1998) instead use moment-matching to approximate the density of the average with a reciprocal gamma density. This method yields poor results when a small number of asset variables is used in the average (low monitoring frequency), since the density of the average is then far from its asymptotic limit (the reciprocal gamma density). Ju (2002) fits a lognormal distribution to the average. A Taylor expansion is then employed around zero volatility to approximate the ratio of the characteristic function of the average to that of the approximating lognormal variable. Based on this, an approximation to the density of the average is provided, which further allows for a closed-form pricing formula; this is observed to work particularly well for low volatilities. More recently, Lord (2006a) determines the Black-Scholes-Merton price of a discretely sampled Asian option as the composition of an exact part and a part that is approximated using conditional moment-matching arguments from Curran (1994a), and further shows that the total price lies between the sharp lower bound of Rogers and Shi (1995) and a sharpening of their upper bound by Nielsen and Sandmann (2003) and Vanmaele et al. (2006). This is known as a partially exact and bounded approximation. In approximating the conditional distribution law of the arithmetic average, the geometric average serves as an optimal conditioning variable. In fact, with his work, Lord (2006a) fixes the divergence of the original approximation of Curran (1994a) for large strike prices.

Significant contributions to the pricing of Asians also rely on PDE approaches. The PDE setup for the Asians is complicated by the fact that one wishes to achieve a reduction in the number of state variables. This reduction, foreshadowed in Ingersoll (1987) and employed in Rogers and Shi (1995), Andreasen (1998) and Večer (2001), (2002), follows from homogeneity of degree 1 of the option payoff and a change to the spot measure. In particular, all the previous PDEs, excluding Večer (2002), suffer from instability under standard (explicit, implicit, Crank-Nicolson) finite difference schemes; this is because the drift dominates the diffusion term in some regions of the grid. To deal with this, Zhang (2001) adjusts the diffusion term and

¹For applications in Lévy economies similar in spirit to Turnbull and Wakeman (1991) and Levy (1992), see, for example, Albrecher and Predota (2002) (use a variance gamma distribution for the average), Albrecher and Predota (2004) (normal inverse Gaussian), Ballotta (2010) (exponential variance gamma).

obtains an analytical solution to the modified PDE as a first-order approximation to the true price, which is however not as accurate as the one provided by Ju's method. Accuracy can be improved by adding a correction term satisfying another PDE which can be solved numerically to high precision. Alternatively, Večer (2001) sets up a new PDE for Asian options based on techniques developed in Shreve and Večer (2000) for pricing options on a traded account, while Večer (2002) provides an even simpler two-term PDE which can be solved to give fast and accurate results, rendering this the most numerically competitive one. Večer and Xu (2004) further extend in pricing Asian options in a semimartingale model and derive a PIDE, which is later implemented numerically for continuously monitored options under jump diffusions in Bayraktar and Xing (2011). Although the previous PDE techniques can be modified to accommodate discrete sampling (see Andreasen (1998), Večer (2002)), this has a side effect on the finite difference algorithms by making the PDE coefficients discontinuous and therefore impacting the quadratic convergence in time of the Crank-Nicolson scheme.

Monte Carlo is typically too slow to compete with the other methods at low dimensions. However, in the Asian case this is not a foregone conclusion since geometric Asians provide a control variate technique that works effectively in the simulation of the arithmetic Asians (see Kemna and Vorst (1990)). Fusai and Meucci (2008) extend the work of Kemna and Vorst (1990) on Gaussian log-returns by deriving the characteristic function of the log-geometric average distribution law under non-Gaussian Lévy log-returns. This in turn yields the price of the geometric Asian option as a Fourier integral which can be computed very efficiently by numerical means (see Section 3.4).

Methods based on recursive integration are specifically adapted to discrete monitoring and can be adjusted to any Lévy assumption for the log-asset returns by simply switching to the relevant characteristic function. Following necessary definitions in the next section, we present in more detail in Section 4.4 the Carverhill-Clewlow forward density convolutions. Subsequently, we provide a new backward price convolution scheme, and demonstrate its numerical and theoretical advantages over the density convolutions, the partially exact and bounded approximation and the control variate Monte Carlo.

Option type	λ_0	$\lambda_1, \dots, \lambda_{n-1}$	λ_n
Fixed-strike call/put	$\zeta \left(\frac{\gamma}{n+\gamma} - \frac{K}{S_0} \right)$	$\frac{\zeta}{n+\gamma}$	$\frac{\zeta}{n+\gamma}$
Floating-strike call/put	$-\frac{\zeta \gamma \alpha^*}{n+\gamma}$	$-\frac{\zeta \alpha^*}{n+\gamma}$	$\zeta \left(1 - \frac{\alpha^*}{n+\gamma} \right)$

Table 4.1: Choice of λ corresponding to different types of Asian options. $\alpha^* > 0$ is the coefficient of partiality for floating-strike options. Coefficient γ takes value 1 (0) when S_0 is (is not) included in the average. Coefficient ζ takes value 1 (-1) for the call (put) option.

4.3 Modelling on reduced state space

Consider the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Define the collection of independent random variables $\{Z_k\}_{k=1}^n$, $n \in \mathbb{N}^*$, as the log-returns, $\ln \frac{S_k}{S_{k-1}}$, on some asset S , with $S_0 > 0$, over sub-periods $\{[t_{k-1}, t_k]\}_{k=1}^n$ of the time line $[0, T]$. Assume $t_k - t_{k-1} = \delta t$ for all k , $t_0 = 0$ and $t_n = T$ (the maturity). Let also $\mathbb{F} = \{\mathcal{F}_k\}_{k=1}^n$ be the information filtration generated by $\{Z_k\}$, with $\mathcal{F}_0$ trivial. If we interpret $\mathbb{P}$ as a risk-neutral measure, we have that under this measure

$$\mathbb{E}(e^{Z_k}) = e^{r(t_k - t_{k-1})},$$

where $r > 0$ is the continuously compounded risk-free interest rate.

The forward price of an Asian option is provided by the unifying form

$$\mathbb{E}(\Lambda_n^+), \tag{4.1}$$

where $\Lambda_k = \sum_{j=0}^k \lambda_j S_j$, for some deterministic process λ originally described in Večer (2002). Different choices of the process λ (see Table 4.1) reflect different type of contracts (call or put; fixed or floating strike price). The difficulty in the computation of the expectation (4.1) arises from the fact that Λ is not a Markov process under $\mathbb{P}$. Instead, (S, Λ) jointly form a Markov system which means, when evaluating (4.1) recursively, that the conditional expectation $\mathbb{E}(\Lambda_n^+ | \mathcal{F}_k)$, $k < n$, depends on both S_k and Λ_k . This implies pricing must be performed on a two-dimensional (excluding time) grid (S, Λ) , raising significantly the computational workload, since, for N grid points in the S dimension, there is a vast of N^n possible averages to the contract expiration (see Andricopoulos et al. (2007) for pricing on a two-dimensional grid using quadrature).

To reduce dimensionality, we define the one-dimensional process

$$\bar{X}_k = \frac{\Lambda_k}{S_k} = \frac{\sum_{j=0}^k \lambda_j S_j}{S_k} = \lambda_k + \frac{\Lambda_{k-1}}{S_k} = \lambda_k + \frac{\Lambda_{k-1}}{S_{k-1}} e^{-Z_k} = \lambda_k + \bar{X}_{k-1} e^{-Z_k} \quad (4.2)$$

for $k = 1, \dots, n$, and

$$\bar{X}_0 = \lambda_0.$$

From (4.2), process $\bar{X}$ is adapted to filtration $\mathbb{F}$ and is Markov under measure $\mathbb{P}$. Additionally, we define a new filtration $\mathbb{G} = \{\mathcal{G}_k\}_{k=1}^n$ with

$$\mathcal{G}_k = \sigma\{Z_n, Z_{n-1}, \dots, Z_{n+1-k}\}.$$

Intuitively, filtration $\mathbb{G}$ is a filtration in which we first observe the log-return in the last time period, then the log-return in the last but one period etc. Then, we define the process X

$$\begin{aligned} X_k &= \lambda_{n-k} + X_{k-1} e^{Z_{n+1-k}}, \quad 0 < k \leq n, \\ X_0 &= \lambda_n, \end{aligned}$$

which is adapted to filtration $\mathbb{G}$ and is Markov under measure $\mathbb{P}$. Having $\lambda_k > 0$ for $k = 1, \dots, n$ implies $X_k > 0$ for $0 \leq k < n$ and $Y_k = X_k - \lambda_{n-k} > 0$ for $0 < k \leq n$, such that

$$\ln Y_k = \ln(Y_{k-1} + \lambda_{n+1-k}) + Z_{n+1-k}, \quad 1 < k \leq n, \quad (4.3)$$

$$\ln Y_1 = \ln \lambda_n + Z_n, \quad (4.4)$$

which corresponds to the so-called Carverhill-Clewlow-Hodges factorization. From (4.3-4.4), it follows by recursive substitution that

$$\Lambda_n = \sum_{k=0}^n \lambda_k S_k = S_0 (Y_n + \lambda_0) = S_0 \left(e^{\ln Y_n} + \lambda_0 \right),$$

hence the valuation equation (4.1) can be restated as

$$\mathbb{E}(\Lambda_n^+) = S_0 \mathbb{E} \left(\left(e^{\ln Y_n} + \lambda_0 \right)^+ \right). \quad (4.5)$$

Expectation (4.5) can now be evaluated iteratively on the one-dimensional grid $\ln Y$.

4.4 Pricing of Asian options by convolution

Carverhill and Clewlow (1990), Benhamou (2002) and Fusai and Meucci (2008) use the transition equation (4.3) to compute the unconditional risk-neutral density of $\ln Y_n$, which they subsequently apply to compute the pricing expectation for the Asian option. The target density is retrieved via recursive evaluation of the density of $\ln Y_k$ as the convolution of the densities of the independent variables $\ln(Y_{k-1} + \lambda_{n+1-k})$ and Z_{n+1-k} in line with equation (4.3). In the first two papers the convolution is computed by Fourier transform, while in the third this is computed directly with numerical integration.

In all three papers the difficulty stems from the fact that the density of $\ln Y_k$ accumulates as k increases and is no longer related to the density of Z_n beyond $k = 1$. Ideally, one should use a dense and narrow grid for $\ln Y_1$, and wide and relatively sparse grid for $\ln Y_n$. Clewlow and Carverhill (1990) use the same equidistantly spaced grid for all variables $\ln Y_k$. On this ground, Benhamou (2002) argues on the slow convergence of their algorithm. To speed up convergence, he models re-centred variables $\ln Y_k - \mathbb{E}(\ln Y_k)$ on a common grid based on the approximation

$$\mathbb{E}(\ln Y_k) \approx \ln \mathbb{E}(e^{\ln Y_{k-1} + \lambda_{n+1-k}}) + \mathbb{E}(Z_{n+1-k}) \quad (4.6)$$

from (4.3). The approximate re-centring technique (4.6), which misses a convexity-adjusting term by Jensen's inequality², becomes weaker especially at high volatilities. More recently, Fusai and Meucci (2008) adopt the original transition mechanism (4.3) (without re-centring) and construct forward-recursive convolution integrals which they evaluate by recursive Gaussian quadratures on a non-equidistant grid, skipping the Fourier transform route. In fact, they only utilize Fourier-inversion to retrieve the densities of $\{Z_{n+1-k}\}_{k=1}^n$ known through their characteristic functions. Although in theory Gaussian integration promises faster convergence than the trapezoidal or Simpson integration, the price pattern of Fusai and Meucci (2008) exhibits non-monotone convergence in the number of integration points when tested on different strikes, sampling frequencies and model assumptions for the log-returns.

²In fact, Jensen's inequality implies that $\mathbb{E} \ln(e^{\ln Y_{k-1} + \lambda_{n+1-k}}) \geq \ln \mathbb{E}(e^{\ln Y_{k-1} + \lambda_{n+1-k}})$.

The main difference between our approach and the foregoing papers is that we model the price of the Asian option directly, rather than the density of the underlying average. Our method resembles backward pricing on a lattice, as opposed to the forward density convolution. Note that smoothness of the densities of $\{Z_{n+1-k}\}_{k=1}^n$ is not required here, as opposed to the aforementioned density convolutions, therefore our result is applicable also to models where the density has a singularity. This occurs, for example, in the variance gamma and normal inverse Gaussian distributions for very short time intervals.

In the next section we write down our recursive pricing algorithm.

4.5 The backward price convolution algorithm

From Section 4.3, expectation $\mathbb{E}(\Lambda_n^+) = S_0 \mathbb{E}\left((e^{\ln Y_n} + \lambda_0)^+\right)$ can be expressed iteratively in filtration $\mathbb{G}$ as

$$S_0 \mathbb{E}\left(\mathbb{E}\left(\mathbb{E}\left(\mathbb{E}\left((e^{\ln Y_n} + \lambda_0)^+\right)\middle|\mathcal{G}_{n-1}\right)\cdots\middle|\mathcal{G}_1\right)\middle|\mathcal{G}_0\right) \quad (4.7)$$

by virtue of the law of iterated expectations. We take $\ln Y$ as the state variable, and express (4.7) in terms of the following recursion for $1 < k \leq n$,

$$p_n(\ln Y_n) = (e^{\ln Y_n} + \lambda_0)^+, \quad (4.8)$$

$$h_{k-1}(\ln Y_{k-1}) = \ln(e^{\ln Y_{k-1}} + \lambda_{n+1-k}),$$

$$q_{k-1}(h_{k-1}(\ln Y_{k-1})) = \mathbb{E}(p_k(\ln Y_k) | \mathcal{G}_{k-1}) \quad (4.9)$$

$$= \mathbb{E}(p_k(h_{k-1}(\ln Y_{k-1}) + Z_{n+1-k}) | \mathcal{G}_{k-1}) \quad (4.10)$$

$$= \int_{\mathbb{R}} p_k(h_{k-1}(\ln Y_{k-1}) + z) f_k(z) dz = (p_k * f_k(-z))(h_{k-1}(\ln Y_{k-1})), \quad (4.11)$$

$$p_{k-1}(\ln Y_{k-1}) = q_{k-1}(h_{k-1}(\ln Y_{k-1})).$$

Equation (4.8) follows from the inner part of the expectation (4.7) and initializes the recursion. Equation (4.9) follows from the law of iterated expectations, equality (4.10) follows from (4.3), while equality (4.11) holds by virtue of the Markov property of the process $\ln Y$ where f_k is the density of the log-return Z_{n+1-k} for any $k = 1, \dots, n$ and $*$ denotes a convolution (see Definition 4). The integral (4.11) is computed recursively for $k = n, \dots, 1$ to provide eventually the forward price of the option

$$S_0 q_0(\ln \lambda_n) \quad (4.12)$$

by virtue of (4.4). We apply to the case of a fixed-strike Asian call with $\lambda_1 = \dots = \lambda_n = \frac{1}{n+\gamma} > 0$ (see Table 4.1). A rigorous proof of result (4.12) is given in Černý and Kyriakou ((2010), Theorem 3.1).

Given the equation (4.11), we obtain from Theorem 5 that

$$\mathcal{F}(p_k * f_k(-z)) = \mathcal{F}(p_k)\mathcal{F}(f_k(-z)) = \mathcal{F}(p_k)\varphi_k,$$

where $\mathcal{F}$ denotes the Fourier transform, and φ_k the complex conjugate of the characteristic function ϕ_k of Z_{n+1-k}

$$\varphi_k(u) = \phi_k(-u) = \mathbb{E}(e^{-iuZ_{n+1-k}}) = \int_{\mathbb{R}} e^{-iuz} f_k(z) dz.$$

q_{k-1} is recovered via

$$q_{k-1} = \mathcal{F}^{-1}(\mathcal{F}(p_k)\varphi_k). \quad (4.13)$$

Note that functions $p_k(y)$ and $q_k(x)$ are not absolutely integrable on the entire real line, therefore their Fourier transforms do not exist (see Definition 1). Černý and Kyriakou ((2010), Theorem 3.2) show that the integrability condition over the negative axis is assisted when y and x are restricted to some compact intervals. In practice, this is taken into account in the approximation of the continuous Fourier transforms by truncated, discrete transforms (see Section 4.5.1).

4.5.1 Numerical implementation

Preliminaries. To evaluate numerically $\mathcal{F}(p_k)$ and subsequently $q_{k-1} = \mathcal{F}^{-1}(\mathcal{F}(p_k)\varphi_k)$, we select evenly spaced grids $\mathbf{u} = \{u_0 + j\delta u\}_{j=0}^{N-1}$, $\mathbf{y} = \{y_0 + l\delta y\}_{l=0}^{N-1}$ and $\mathbf{x} = \{x_0 + m\delta x\}_{m=0}^{N-1}$ with N grid points and spacings δu and $\delta y = \delta x$. More precisely, grid $\mathbf{u}$ is chosen to be symmetric around zero such that $\mathbf{u}_j = (j - N/2)\delta u$ for $j = 0, \dots, N-1$. The range of values of $\mathbf{u}$ is determined to ensure that $|\phi_k| < 10^{-\rho}$ outside $\mathbf{u}$, where the ρ value is guided by the targeted precision, e.g., $\rho = 7$ corresponds to 7 decimal places of accuracy. Note that for higher ρ the range of $\mathbf{u}$ becomes wider, thus the spacing δu becomes larger for fixed N , which is undesirable in numerical integration. For this, ρ has to be chosen with care. Černý and Kyriakou ((2010),

Theorem 3.2) additionally derive a rule for the ranges of $\mathbf{y}$ and $\mathbf{x}$ to achieve a predetermined pricing error caused by curtailment of the integration range in (4.11). Alternatively, one can try several ranges arbitrarily and pick the narrowest one that guarantees a smoothly convergent price pattern with increasing N . Furthermore, to achieve monotone convergence, it is necessary to construct the grid $\mathbf{y}$ with range $[L, U] \subset \mathbb{R}$ such that

$$\mathbf{y}_l = \ln(-\lambda_0) + \left(\left\lceil \frac{L - \ln(-\lambda_0)}{\delta y} \right\rceil + l \right) \delta y,$$

which allows us to place the point of nonlinearity $y = \ln(-\lambda_0)$ (see equation (4.8)) upon a node of the grid.

In the next two steps, we summarize the recursive part of the numerical scheme which is applied n times (as many as the number of monitoring dates), starting from maturity and moving backwards to provide the option price at inception.

1. **Swapping between the state and Fourier spaces.** Suppose the values approximating p_k at the k^{th} monitoring date are given on grid $\mathbf{y}$ and the values of φ_k are given on grid $\mathbf{u}$. We denote these by $\mathbf{p}_k$ and $\boldsymbol{\varphi}_k$ respectively. Using the conversion (2.23), we evaluate the fractional DFT $\mathbf{P}_k = \mathcal{D}(\mathbf{p}_k \mathbf{w}, \mathbf{y}, \mathbf{u}; N \delta u \delta y / 2\pi) \delta y$ as a discrete approximation of the Fourier transform $\mathcal{F}(p_k)$ on grid $\mathbf{u}$, where $\mathbf{w}$ denotes some low-order Newton-Côtes integration weights, e.g., for the trapezoidal rule $\mathbf{w}_l = 1 - \frac{1}{2}(\delta_l + \delta_{N-1-l})$ where the Kronecker delta δ_ϵ takes value 1 (0) for $\epsilon = 0$ ($\epsilon \neq 0$). We then approximate the inverse Fourier transform (4.13) on grid $\mathbf{x}$ by computing $\mathbf{q}_{k-1} = \frac{1}{2\pi} \mathcal{D}(\mathbf{P}_k \boldsymbol{\varphi}_k, -\mathbf{u}, \mathbf{x}; -N \delta u \delta x / 2\pi) \delta u$, following the conversion (2.24).
2. **From q to p .** We calculate $\mathbf{p}_{k-1} = q_{k-1}(h_{k-1}(\mathbf{y}))$: we approximate q_{k-1} inside grid $\mathbf{x}$ by fitting a cubic interpolating spline to the nodes $(\mathbf{x}, \mathbf{q}_{k-1})$ by utilizing the MATLAB built-in function INTERP1. Outside $\mathbf{x}$, we approximate q_{k-1} by extrapolating linearly in e^x .

4.6 Numerical study

For the purposes of this study, we opt for an Asian call option (in consistency with the notation in Table 4.1: $\zeta = 1, \gamma = 1$) with fixed strike price K and maturity after a year's time, i.e., $T = 1$. Sampling frequency is n . All numerical experiments are coded in MATLAB R15 on a Dell Latitude 620 Intel Core 2 Duo T7200 PC 2.00 GHz with 2.0 GB RAM.

4.6.1 Models

We price options numerically based on three distributions of log-returns from the Lévy class: Gaussian, normal inverse Gaussian and tempered stable. From Section 3.2, the characteristic function of some Lévy process L_t under the risk-neutral measure is given by

$$\begin{aligned}\mathbb{E}(e^{iuL_t}) &= e^{\psi_L(u)t}, \\ \psi_L(u) &= i(r - \chi_L(-i))u + \chi_L(u),\end{aligned}$$

where t is the time horizon (in years). The functions χ_L for the different models are

$$\begin{aligned}\chi_G(u) &= -\sigma^2 u^2 / 2, \\ \chi_{\text{NIG}}(u) &= (1 - \sqrt{1 - 2i\theta\nu u + \nu\sigma^2 u^2}) / \nu, \\ \chi_{\text{CGMY}}(u) &= C\Gamma(-Y)((M - iu)^Y - M^Y + (G + iu)^Y - G^Y).\end{aligned}\tag{4.14}$$

The exact cumulants $\mathbb{E}(L_t)$, $\text{Var}(L_t)$, $c_3(L_t)$, $c_4(L_t)$ are derived for all models by differentiating the corresponding cumulant generating functions, $\psi_L(u)t$, and evaluating at zero, according to equation (2.5).

We calibrate the three models to achieve $\text{vol} := \text{Var}(L_1)^{1/2} \in \{0.1, 0.3, 0.5\}$ and, for the non-Gaussian distributions, further $s(L_1) = -0.5$ and $\kappa(L_1) = 0.7$. These values are broadly consistent with risk-neutral densities fitted to option price data in Madan et al. (1998). The fitted parameters, rounded to four leading digits, are presented in Table 4.2.

Gaussian σ	normal	inverse Gaussian σ	θ	tempered stable			
	ν			C	G	M	Y
0.1	0.1222	0.0879	-0.1364	0.2703	17.56	54.82	0.8
0.3	0.1222	0.2637	-0.4091	0.6509	5.853	18.27	0.8
0.5	0.1222	0.4395	-0.6819	0.9795	3.512	10.96	0.8

Table 4.2: Calibrated model parameters

4.6.2 Pricing in the Black-Scholes-Merton economy

We investigate how the results from the backward price convolution with precision $\pm 10^{-7}$ compare with the original ones from the forward density convolutions of Carverhill and Clewlow (1990) and Benhamou (2002) for different strikes and volatilities. Table 4.3 illustrates that Benhamou (2002) provides improvement over Carverhill and Clewlow (1990), especially for high volatilities. Note that for $\sigma \in \{0.1, 0.3\}$, Carverhill and Clewlow (1990) achieve in three cases smaller absolute % error than Benhamou (2002). In general, both density convolutions tend to suffer less for lower volatilities and in/at-the-money options. Comparing with the price convolution, we observe a variable precision of 1–3 decimal places for both density convolutions.

σ	K	Backward convolution prec. $\pm 10^{-7}$	Benhamou	% error	Carverhill & Clewlow	% error
0.1	80	22.7771749	22.7838	-0.029	22.78	-0.013
	90	13.7337773	13.7347	-0.007	13.73	0.028
	100	5.2489927	5.2438	0.099	5.25	-0.019
	110	0.7238324	0.7211	0.379	0.72	0.532
	120	0.0264092	0.0336	-21.401	0.02	32.046
0.3	80	23.0914378	23.0733	0.079	23.09	0.006
	90	15.2207610	15.2231	-0.015	15.29	-0.453
	100	9.0271888	9.0110	0.180	9.08	-0.582
	110	4.8349071	4.8338	0.023	4.86	-0.516
	120	2.3682854	2.3545	0.586	2.4	-1.321
0.5	80	24.8242581	24.8324	-0.033	25.01	-0.743
	90	18.3316740	18.3207	0.060	18.5	-0.910
	100	13.1580456	13.1811	-0.175	13.47	-2.316
	110	9.2345134	9.2300	0.049	9.45	-2.280
	120	6.3719536	6.3615	0.164	6.68	-4.612

Table 4.3: Fixed-strike Asian call option ($\zeta = 1$, $\gamma = 1$, $T = 1$, $n = 50$): comparison with Benhamou (2002) and Carverhill and Clewlow (1990) for Gaussian log-returns. Error expressed as a percentage of the backward convolution price (precision $\pm 10^{-7}$). Other parameters: $r = 0.1$, $S_0 = 100$.

Following successive grid refinements, the price convolution scheme reaches monotone con-

vergence, permitting, as a consequence, high accuracies across strikes and volatilities. Inability of the density convolution schemes to maintain regular convergence in the number of grid points makes it hard to judge on the precision of their outcome. The same applies to the most recent approach by Fusai and Meucci (2008), which is, though, a substantial improvement over Carverhill and Clewlow (1990) and Benhamou (2002). Table 4.4 presents the original results from Fusai and Meucci (2008) for 1,000, 5,000 and 10,000 quadrature points: we observe that the absolute % error for 1,000 points is in two cases smaller than the error for 5,000 and 10,000 points, while the error for 5,000 points is smaller than the error for 10,000 points in further four cases. It is also indicated that, for 5,000 and 10,000 points, the absolute % error reduces as the option moves into-the-money; we then observe precision to 4 decimal places. Instead, an implementation of Fusai and Meucci's algorithm with 1,000 points (the least CPU-demanding for given n) typically yields precision to 3 decimal places in 5 seconds, as opposed to 1 second for guaranteed 5 decimal-place accuracy with our pricing procedure (see Table 4.7).³

n	K	Backward convolution	Fusai & Meucci			% error		
		prec. $\pm 10^{-7}$	10,000	5,000	1,000	10,000	5,000	1,000
12	90	11.9049157	11.90497	11.90498	11.90428	-0.0005	-0.0006	0.0054
	100	4.8819616	4.88210	4.88212	4.88199	-0.0028	-0.0033	-0.0006
	110	1.3630380	1.36314	1.36314	1.36371	-0.0075	-0.0075	-0.0493
50	90	11.9329382	11.93301	11.93299	11.93339	-0.0006	-0.0004	-0.0038
	100	4.9372028	4.93736	4.93738	4.93711	-0.0032	-0.0036	0.0019
	110	1.4025155	1.40264	1.40262	1.40199	-0.0089	-0.0075	0.0375
250	90	11.9405632	11.94068	11.94069	11.94137	-0.0010	-0.0011	-0.0068
	100	4.9521569	4.95233	4.95239	4.94942	-0.0035	-0.0047	0.0553
	110	1.4133670	1.41351	1.41350	1.41290	-0.0101	-0.0094	0.0331

Table 4.4: Fixed-strike Asian call option ($\zeta = 1$, $\gamma = 1$, $T = 1$): comparison with Fusai and Meucci (2008) for Gaussian log-returns. Error expressed as a percentage of the backward convolution price (precision $\pm 10^{-7}$). Numbers 1,000, 5,000, 10,000 in the last six columns signify the number of grid points used by Fusai and Meucci (2008). Other parameters: $\sigma = 0.17801$, $r = 0.0367$, $S_0 = 100$.

With the focus on the Black-Scholes-Merton framework, PDE methods (see Večer (2002)) and partially exact and bounded (PEB) analytical approximations (see Lord (2006a)) provide additionally efficient means for valuing discretely sampled Asian options. We deal here only

³The CPU time is in our favour, since our control variate Monte Carlo executes 1,000,000 trials in 190 seconds, whereas Fusai and Meucci's takes 130 seconds, for $n = 50$.

with the PEB method, as Večer's PDE has been studied in detail in Černý and Kyriakou (2010). The PEB approximation originates from the work of Curran (1994a) who decomposes the Asian option price q_0 into the parts $q_{a,0}$ and $q_{b,0}$, such that

$$\begin{aligned} q_0 &= e^{-rt_n} (q_{a,0} + q_{b,0}); \\ q_{a,0} &= \mathbb{E} \left((A_n - K)^+ 1_{\{G_n < K\}} \right), \\ q_{b,0} &= \mathbb{E} \left((A_n - K)^+ 1_{\{G_n \geq K\}} \right), \end{aligned}$$

where $A_n = \frac{1}{n+\gamma}(\gamma S_0 + \sum_{k=1}^n S_k)$ and $G_n = (S_0^\gamma \prod_{k=1}^n S_k)^{1/(n+\gamma)}$. Using that $A_n \geq G_n$, $q_{b,0}$ simplifies to

$$q_{b,0} = \mathbb{E} \left((A_n - K) 1_{\{G_n \geq K\}} \right),$$

which is computed in closed form (see Curran (1994a), Section 2.2). This is not the case with $q_{a,0}$, since, simply knowing that $G_n \leq K$, is not enough to say whether the arithmetic option finishes in-the-money or not. To compute $q_{a,0}$, Curran ((1994a), Section 2.3) suggests approximating A_n conditional on G_n by some nonnegative random variable $\Psi(G_n)$, such that

$$q_{a,0} = \mathbb{E} \left(\mathbb{E} \left((A_n - K)^+ 1_{\{G_n < K\}} | G_n \right) \right) = \int_0^K \mathbb{E} \left((\Psi(K') - K)^+ \right) dF_{G_n}(K'), \quad (4.15)$$

where F_{G_n} denotes the distribution function of G_n and the expectation inside the integral reflects the expression for the price of a European plain vanilla call option. In computing (4.15), Curran (1994a) determines the distribution of $\Psi(G_n)$ which satisfies

$$\mathbb{E}(A_n | G_n = K') = \mathbb{E}(\Psi(K')), \quad (4.16)$$

$$\mathbb{V}\text{ar}(A_n | G_n = K') = \mathbb{V}\text{ar}(\Psi(K')) \quad (4.17)$$

for $K' = K$, where the left-hand sides of (4.16-4.17) are known. Then, (4.15) is evaluated numerically by quadrature on the grid K' . Lord ((2006a), Theorem 5) proves that Curran's approximation diverges as $K \rightarrow \infty$. To fix divergence for large strikes, Lord ((2006a), Theorem 4) suggests matching the conditional moments (4.16) and (4.17) at all the grid points $K' \leq K$. This also guarantees that the resulting approximation q_0 lies between sharp lower and upper

bounds (see Lord (2006a), Section 6).

From Table 4.5, Lord’s PEB approximation performs extremely well for $n = 50$ and $\sigma \in \{0.1, 0.3\}$, by exhibiting 5 decimal-place precision across strikes in 0.3 seconds, dominating our method which requires, instead, 1 second for $\sigma = 0.1$ (see Table 4.7). However, the PEB approximation becomes less competitive for high volatility, since it only converges to the fourth decimal place for the in-the-money option. Precision restores to 5 decimal places as K increases. For $n = 12$, and $\sigma = 0.5$, $K = 90$ (worst-case scenario for PEB), we achieve precise results at $\pm 10^{-5}$ in 0.07 and 0.3 seconds with the price convolution scheme and the PEB approximation respectively. Both methods agree to the same accuracy in 0.15 seconds when σ reduces to 0.1. Raising, instead, $n = 250$ sees the PEB approximation as the winner (in terms of speed) for precision $\pm 10^{-5}$ for all strikes and volatility levels, still our method has an extra edge for higher precision levels.

For a larger number of sampling dates (tending to infinity), we expect the PDE by Večer (2002) implemented with a Crank-Nicolson scheme to be the best-performing alternative to both methods. This is investigated in greater detail in Černý and Kyriakou (2010).

σ	K			CPU (s)
	90	100	110	
0.1	11.58113	3.33861	0.27375	0.3
0.3	13.66981	7.69859	3.89638	0.3
0.5	17.19241	12.09154	8.31440	0.3

Table 4.5: Fixed-strike Asian call option ($\zeta = 1$, $\gamma = 1$, $T = 1$, $n = 50$) for Gaussian log-returns: results of the PEB approximation implemented with G_n as the conditioning variable. $\Psi(G_n) = G_n + \exp(\mu_{A|G}(G_n) + \sigma_{A|G}(G_n)Z)$; $Z \sim \mathcal{N}(0, 1)$. Parameters $\mu_{A|G}$, $\sigma_{A|G}$ determined via moment-matching (see equations (4.16), (4.17)) for each grid point $G_n \leq K$. Other parameters: $r = 0.04$, $S_0 = 100$. CPU times in seconds (s).

4.6.3 Pricing in Lévy economies

We do not detail numerical comparisons with the existing density convolutions for Lévy log-returns, as these appear to have a detrimental effect on the error convergence of the scheme. We compare, instead, against the outcome from Monte Carlo simulation accelerated by the geometric Asian control variate (CVMC), as proposed in Fusai and Meucci (2008). We simulate NIG trajectories using standard time-change Brownian representation for the NIG process (see

Glasserman (2004), Section 3.5.2), while for the CGMY process we employ the joint Monte Carlo-Fourier transform scheme developed in the follow-up Chapter 7. Exact prices for geometric Asian options are obtained from the analytical formula (3.13). In implementing the control variate technique, the control variate estimate $\bar{q}_{CV}^{aa}$ for the arithmetic Asian is obtained as

$$\bar{q}_{CV}^{aa} = \bar{q}_{MC}^{aa} - \beta (\bar{q}_{MC}^{ga} - q^{ga}),$$

where q^{ga} is the exact geometric Asian option price, and $\bar{q}_{MC}^{aa}, \bar{q}_{MC}^{ga}$ the (crude) Monte Carlo estimates of the arithmetic and geometric Asian option prices respectively. The optimal coefficient $\beta = \beta^* = \text{Cov}(\bar{q}_{MC}^{aa}, \bar{q}_{MC}^{ga}) / \text{Var}(\bar{q}_{MC}^{ga})$ is chosen to minimize the variance of the estimator $\bar{q}_{CV}^{aa}$. Since β^* itself is unknown, we must pre-estimate and fix this by regressing m^* simulations of the arithmetic price against the respective geometric price (before applying to price estimation using a new set of m simulations), to avoid generating an undesirable amount of bias in the final $\bar{q}_{CV}^{aa}$ estimate (see Glasserman (2004), p. 195-196, and Section 4.1.3).

In Table 4.6 we report standard errors and CPU timings for the control variate method across strikes and volatilities, for each Lévy model. CVMC appears to perform best (lowest standard error) for $K = 110$, $\text{vol} = 0.1$ for all three models. Comparing with our convolution method, we conclude that even with a very effective control variate the Monte Carlo method is not competitive as it requires 19, 26, 95 seconds to achieve at best 4 decimal places of accuracy (at 99% confidence level) for the Gaussian, NIG, CGMY models respectively, whereas our method needs 1, 3.7, 8.5 seconds to obtain 5 decimal places for the same cases (see Table 4.7).

In Table 4.7 we report prices from the backward price convolution algorithm with precision $\pm 10^{-5}$. We can achieve higher precision (up to 10^{-8}) by exploiting the regular second-order convergence of our scheme in the number of grid points. Beyond 10^{-8} we require too many grid points to keep the grid spacing δu sufficiently small and maintain smooth convergence, raising significantly the computational time (see Section 4.5.1).

Two comments are in order. Firstly, comparing Gaussian with leptokurtic prices, we find out-of-the-money options to be more expensive, while in-the-money options slightly cheaper. This pattern is attributed to a combination of negative skewness and excess kurtosis effects in

model	vol	K						CPU (s)
		90		100		110		
		CVMC	std $\times 10^{-5}$	CVMC	std $\times 10^{-5}$	CVMC	std $\times 10^{-5}$	
Gaussian	0.1	11.58129	25	3.33880	20	0.27373	15	19
	0.3	13.66806	168	7.69944	157	3.89814	146	
	0.5	17.19361	495	12.09046	474	8.31389	439	
NIG	0.1	11.63998	24	3.32381	15	0.15821	11	26
	0.3	13.70351	149	7.34122	123	3.28026	106	
	0.5	16.75916	391	11.23206	345	7.17027	323	
CGMY	0.1	11.63994	24	3.32448	15	0.15774	11	95
	0.3	13.70075	150	7.34662	119	3.28289	105	
	0.5	16.76107	384	11.23964	343	7.17578	310	

Table 4.6: Fixed-strike Asian call option ($\zeta = 1$, $\gamma = 1$, $T = 1$, $n = 50$): results of control variate Monte Carlo (CVMC) strategy with 100,000 trials and Lévy log-returns. “vol”: standard deviation of log-returns for a one-year time horizon, “std”: standard error of CVMC estimator. Model parameters: Table 4.2. Other parameters: $r = 0.04$, $S_0 = 100$. CPU times in seconds (s) (excl. computational time for the true geometric Asian price).

the risk-neutral distribution. Secondly, the prices generated by the two Lévy models coincide to penny accuracy. This suggests that the skewness and excess kurtosis of the risk-neutral distribution are the primary factors driving option prices, rather than the actual model choice itself.

4.6.4 Standard DFT versus fractional DFT

In the numerical implementation of the algorithm in Section 4.5.1, in step 1, one may instead consider evaluating the discrete approximations of the Fourier transforms by utilizing standard (I)DFTs, i.e., by computing $\mathbf{P}_k = \mathcal{D}(\mathbf{p}_k \mathbf{w}, \mathbf{y}, \mathbf{u}; 1) \delta y$ and $\mathbf{q}_{k-1} = \frac{1}{2\pi} \mathcal{D}(\mathbf{P}_k \boldsymbol{\varphi}_k, -\mathbf{u}, \mathbf{x}; -1) \delta u$ using the conversions (2.25) and (2.26) respectively.

As discussed in Section 2.5, to apply the (I)DFTs, it is necessary that the restriction

$$\delta u \delta y = \delta u \delta x = \frac{2\pi}{N} \quad (4.18)$$

model	vol	K			CPU (s)
		90	100	110	
Gaussian	0.1	11.58113	3.33861	0.27375	1.0
	0.3	13.66981	7.69859	3.89639	0.3
	0.5	17.19239	12.09153	8.31441	0.3
NIG	0.1	11.64024	3.32385	0.15835	3.7
	0.3	13.70084	7.34265	3.27860	1.8
	0.5	16.76306	11.23586	7.16836	1.8
CGMY	0.1	11.63988	3.32458	0.15787	8.5
	0.3	13.70160	7.34742	3.28308	4.1
	0.5	16.76835	11.24424	7.17624	2.1

Table 4.7: Fixed-strike Asian call option ($\zeta = 1$, $\gamma = 1$, $T = 1$, $n = 50$) for Lévy log-returns: results of the backward price convolution, precision $\pm 10^{-5}$. “vol”: standard deviation of log-returns for a one-year time horizon. Model parameters: Table 4.2. Other parameters: $r = 0.04$, $S_0 = 100$. CPU times in seconds (s).

is in force. Equation (4.18) allows us to state

$$\delta u = \frac{a_1}{\sqrt{N}}, \quad \delta y = \delta x = \frac{a_2}{\sqrt{N}} \quad (4.19)$$

with

$$a_1 a_2 = 2\pi.$$

Equation (4.19) suggests that as N grows to infinity, the grid spacings δu , δy (and δx) reduce to zero, whereas the grid widths $N\delta u = \sqrt{N}a_1$ and $N\delta y = N\delta x = \sqrt{N}a_2$ expand notably. Extremely wide grids result to significantly small and large values of φ_k and p_k respectively, causing numerical instability of the (I)FFT routines, consequently affecting the smooth convergence $\mathbf{P}_k \rightarrow \mathcal{F}(p_k)$ and $\mathbf{q}_k \rightarrow q_k$ as $N \rightarrow \infty$. Similar phenomena are observed when we fix one of δu , δy for all N , instead; for instance, fixing $\delta u = b$ ensures that $\delta y = \delta x = \frac{2\pi}{Nb} \rightarrow 0$ as $N \rightarrow \infty$, such that $N\delta y = N\delta x = \frac{2\pi}{b}$ remains fixed, whereas $N\delta u = Nb \rightarrow \infty$. Alternatively, the fractional DFT permits full control by the user over all δu , δy , δx , N , therefore allows to refine grids $\mathbf{u}$, $\mathbf{y}$, $\mathbf{x}$ by increasing N , while simultaneously maintaining their widths fixed.

Figure 4-1 illustrates the error convergence for both standard and fractional DFT implementations in pricing an at-the-money Asian call ($\zeta = 1$, $\gamma = 1$, $T = 1$, $n = 50$), with $r = 0.04$, $S_0 = 100$, and Gaussian log-returns with $\sigma = 0.1$. We observe that the fractional DFT converges smoothly, guaranteeing high-level precision for sufficiently large N . For low precision




Figure 4-1: Standard DFT versus fractional DFT implementation of the backward convolution algorithm: error convergence in the number of grid points N . “Error”: consecutive price differences computed for $N = N^*$ and $N = 2N^*$. (I)DFTs & fractional DFT: fixed $\mathbf{y}$, $\mathbf{x}$ ranges for all N . Fractional DFT: fixed $\mathbf{u}$ range for all N . (I)DFTs: fixed spacing $\delta u = b$, $\delta y = \delta x = \frac{2\pi}{Nb}$ via (4.18).

targets, the standard DFT is preferable for faster results (see Section 2.5).

4.7 Concluding remarks

We have presented a backward price convolution for arithmetic Asian options with discrete sampling. Our algorithm is a major improvement over existing forward density convolutions which exhibit non-monotone convergence when implemented numerically. With a fractional DFT implementation of our scheme, we observe regular second-order convergence in the number of grid points for different strikes and volatilities, which therefore can be accelerated by utilizing Richardson extrapolation. Furthermore, our numerical study has shown that even a very effective control variate Monte Carlo method is not competitive with the convolution method. On the other hand, Lord’s PEB approximation becomes more competitive for low asset volatility and high sampling frequencies. Still, our pricing procedure has an extra edge for high volatility and high precision targets. Furthermore, the PEB approximation is not applicable to non-Gaussian log-returns, as opposed to our method.

In what follows, we investigate a first extension of the backward convolution method to obtain the option price sensitivities, and a second extension to compute the option prices under stochastic volatility.

Chapter 5

Computation of the Asian option price sensitivities

5.1 Introduction

Practitioners are interested in efficient ways to calculate the prices of derivative contracts, but also their sensitivities (popularly known as the “Greeks”) obtained by differentiating the contract price with respect to parameters of interest. Estimation of the price sensitivities is important, since, apart from their use for risk management and hedging, they can also serve as a measure of the pricing error resulting from parameter values that may be inappropriate, or may vary during the life of the contract. Moreover, price sensitivities contribute directly to the price quotes, since the bid-ask spread is often proportional to some Greeks.

The literature distinguishes between finite difference approximations for the sensitivities and methods based on direct differentiation of a pricing formula with respect to some parameter of interest. The first approach suggests calculating a price sensitivity indirectly, by recalculating the price at a perturbed parameter value and approximating the corresponding sensitivity using a finite difference. Although conceptually simple and intuitive, this formulation suffers from practical issues, including the lack of standard rules for choosing the optimal perturbation size for different payoffs and sensitivities sought and, therefore, the inability to gauge the precision of the outcome. It also doubles the computational time by the need to calculate the price twice (for a first-order sensitivity).

To deal with these issues, two approaches based on direct differentiation of the risk-neutral valuation formula have been considered. The first one, employed originally by Curran (1994b), (1998) and Broadie and Glasserman (1996) in option pricing applications, relies on the relationship between the contingent claim payoff and the parameter of interest. In the Monte Carlo literature, this is known as the pathwise (PW) approach. In particular, for a continuous payoff as a function of the parameter of differentiation, the relevant sensitivity is obtained by differentiating the payoff function inside the pricing expectation. However, payoff functions are typically not smooth, as opposed to the probability density of the underlying asset vector as a function of its parameters. By exploiting this very useful property of the probability density, Broadie and Glasserman (1996) advocate to differentiate the probability density, instead. In the Monte Carlo context, this is frequently termed the likelihood ratio (LR) method. Favourably, the latter technique provides an efficient tool for computing required sensitivities for various contracts, while suppresses the need for recalculating their prices.

So far, Monte Carlo has been the method of choice for the estimation of the Greeks using the techniques mentioned above, mainly due to its ease of implementation and ability to adapt flexibly to various payoff structures. Alternatively, sensitivities have been computed by means of PDEs (see Norberg (2006)). More recently, recursive integration-based methods have been replacing the PDE methods, aiming to establish procedures of comparable speed which can fit flexibly to Lévy economies, but also the Monte Carlo methods which experience slow convergence and limitation in handling early-exercise features (see Lord et al. (2008), Fang and Oosterlee (2008b) for discretely monitored Bermudan vanilla and barrier options, and Feng and Linetsky (2008) for barrier options).

In the Asians case, several analytical results for the Greeks have been obtained via direct differentiation of original price representations: with the focus on an arithmetic Asian option on a lognormal underlying, we mention, from Section 4.2, Geman and Yor (1993), Ju (2002), Lord (2006a) and Zhang (2001). For an underlying driven by an exponential Lévy model, forward density convolution algorithms (see Carverhill and Clewlow (1990), Benhamou (2002), Fusai and Meucci (2008)) have been employed to compute the density of the discrete arithmetic average price. This density is then utilized to compute the pricing expectation of the Asian option. By directly differentiating inside the expectation, Fusai and Meucci (2008) obtain also

modified expectations for the Greeks. In practice, however, forward density convolutions tend to suffer by non-smooth convergence to the true option price (or sensitivity), as discussed in Chapter 4. In the light of this limitation, we focus attention on the efficient backward price convolution of Černý and Kyriakou (2010) (see also Chapter 4), which bypasses the calculation of the density of the average price, and extend this to the computation of the Greeks.

The chapter is organized as follows: with the focus on a generic European contract, we present in Section 5.2 the two methods for deriving price sensitivities via direct differentiation of the risk-neutral valuation formula. We further derive integral representations for the price sensitivities which are applicable to a wide range of European payoffs and models for the asset log-returns from the affine class. Based on the same principles, we develop in Section 5.3 the backward convolution algorithm for the price sensitivities of discretely sampled arithmetic Asian options. In Section 5.4 we revisit standard Monte Carlo methods for the estimation of the Greeks and show how one of these, the likelihood ratio method, can be adapted to any Lévy assumption for the asset log-returns. Section 5.5 presents numerical comparisons between our convolution scheme and the Monte Carlo equipped with a geometric Asian option as control variate. Section 5.6 concludes the chapter.

5.2 Price sensitivities via direct differentiation

Consider the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We interpret $\mathbb{P}$ as a risk-neutral probability measure. Let function $C = q(S)$ be the payoff of a generic contingent claim written on the asset vector $S = (S_1, S_2, \dots, S_n)$. Consider parameter $\beta \in \mathcal{B}$, for some bounded interval $\mathcal{B}$ on the real axis. Fix $\omega \in \Omega$ and think of $S = h(\beta, \omega)$, such that $C = q(h(\beta, \omega))$. For a smooth function $q(h(\cdot, \omega))$ on the interval $\mathcal{B}$, we obtain, under appropriate regularity conditions, the (forward) sensitivity as

$$\frac{\partial \mathbb{E}(C)}{\partial \beta} := \frac{\partial \mathbb{E}(q(h(\beta, \omega)))}{\partial \beta} = \mathbb{E}\left(\frac{\partial q(h(\beta, \omega))}{\partial \beta}\right) =: \mathbb{E}\left(\frac{\partial C}{\partial \beta}\right). \quad (5.1)$$

Approach (5.1) is commonly used in the Monte Carlo context to generate the so-called pathwise (PW) estimators for the Greeks, which are unbiased.

We illustrate method (5.1) by applying on a European plain vanilla call option which expires

at T with payoff

$$C = q(S_T) = (S_T - K)^+ = (S_T - K) 1_{\{S_T > K\}},$$

where $K > 0$. For concreteness, let S be modelled as

$$S_T = h(S_0, \omega) = S_0 e^{(r-\varpi)T + L_T(\omega)}$$

with S_0 , r , ϖ positive constants and L a Lévy process. The delta of the option is given by the partial derivative of its price with respect to S_0 . Viewed as a function of S_T , C is continuous and piecewise differentiable, and has derivative $\frac{\partial C}{\partial S_T} = 1_{\{S_T > K\}}$, while S_T is linear in S_0 with derivative $\frac{\partial S_T}{\partial S_0} = \frac{S_T}{S_0}$ for fixed ω . Therefore, applying the chain rule for differentiation yields

$$\frac{\partial C}{\partial S_0} = \frac{\partial C}{\partial S_T} \frac{\partial S_T}{\partial S_0} = 1_{\{S_T > K\}} \frac{S_T}{S_0}. \quad (5.2)$$

To obtain the (forward) delta of the option, one needs to compute

$$\mathbb{E} \left(\frac{\partial C}{\partial S_0} \right) = \mathbb{E} \left(1_{\{S_T > K\}} \frac{S_T}{S_0} \right).$$

The deltas of single-asset options with continuous payoffs, e.g., Asian, lookback options, but also of multi-asset options, e.g., spreads and options on the minimum/maximum of several assets, can be obtained in a similar way. Approach (5.1) fails when used to compute sensitivities for products with discontinuous payoffs, e.g., digital and barrier options. The same issue arises when applying this method to payoffs with non-smooth derivatives in order to compute higher-order price sensitivities. This is the case, for instance, with $\frac{\partial C}{\partial S_0}$ in (5.2) which presents a discontinuity at $S_T = K$. This prohibits the use of method (5.1) to compute the so-called (forward) gamma; that is, the second-order sensitivity with respect to S_0 .

In contrast to a payoff function, the probability distribution of the underlying asset is typically a smooth function of its parameters. This allows us to rephrase the problem of calculating sensitivities by passing the dependence on parameter β from the payoff to the underlying probability law. In this setting, the forward price of the contingent claim is given by

$$\mathbb{E}_\beta(C) = \int_{\mathbb{R}^n} q(s) g(s; \beta) ds,$$

where g is the density of S . We call this the *distribution-based* approach. Furthermore, for asset log-return Z obeying to the probability law of an affine model, we write $Z = \ln \frac{S}{S_0}$ where $S_0 > 0$. We define generic parameter $\theta > 0$ related to the probability distribution of Z , e.g., asset volatility σ in the Gaussian model, hence $\theta \neq S_0$. The density of Z reads

$$f(z; \theta) := S_0 e^z g(S_0 e^z; \beta),$$

where $\beta = \{S_0, \theta\}$.

Consider a European, path-independent contingent claim on a single asset S which expires at T . Based on explicit integral representation for the price of this claim from Section 3.4, we will derive integral representations for the price sensitivities in the distribution-based setting.

5.2.1 Distribution-based approach for price sensitivities in affine models

Assume $Z_T = \ln \frac{S_T}{S_0}$ has density function $f(z; \theta)$ and characteristic function

$$\phi(u; \theta) = \int_{\mathbb{R}} e^{iuz} f(z; \theta) dz. \quad (5.3)$$

Recall from equation (3.11) the integral representation for the forward price of a contingent claim with general payoff $p(\ln S_T)$

$$V_0(S_0, \theta) = \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} S_0^{-iu} \phi(-u; \theta) \varrho(u) du, \quad (5.4)$$

where $\varrho(u) = \int_{\mathbb{R}} e^{iux} p(x) dx$.

Under appropriate regularity conditions which justify differentiation under the integral sign, we derive sensitivities by differentiating with respect to parameters of interest. For example, the (forward) delta is given by

$$\frac{\partial V_0(S_0)}{\partial S_0} = \frac{-S_0^{-1}}{2\pi} \int_{iR-\infty}^{iR+\infty} S_0^{-iu} iu \phi(-u) \varrho(u) du. \quad (5.5)$$

Given (5.3) and for $f'(z) = \frac{\partial f(z)}{\partial z}$, we have from Theorem 3 that

$$\frac{\partial V_0(S_0)}{\partial S_0} = \frac{-S_0^{-1}}{2\pi} \int_{iR-\infty}^{iR+\infty} S_0^{-iu} \left(\int_{\mathbb{R}} e^{-iuz} f'(z) dz \right) \varrho(u) du = -S_0^{-1} \int_{\mathbb{R}} p(\ln S_0 + z) f'(z) dz, \quad (5.6)$$

where the second equality follows by changing the order of integration and applying result (3.11). The derivation of the (forward) gamma is based on the same arguments:

$$\begin{aligned} \frac{\partial^2 V_0(S_0)}{\partial S_0^2} &= \frac{1}{2\pi} \int_{iR-\infty}^{iR+\infty} iu(iu+1) S_0^{-iu-2} \phi(-u) \varrho(u) du \\ &= \frac{S_0^{-2}}{2\pi} \int_{iR-\infty}^{iR+\infty} S_0^{-iu} (-u^2) \phi(-u) \varrho(u) du - S_0^{-1} \frac{\partial V_0(S_0)}{\partial S_0}, \end{aligned} \quad (5.7)$$

where the second term in (5.7) follows from (5.5). From Theorem 3 and a change in the order of integration, we get

$$\frac{\partial^2 V_0(S_0)}{\partial S_0^2} = S_0^{-2} \int_{\mathbb{R}} p(\ln S_0 + z) f''(z) dz - S_0^{-1} \frac{\partial V_0(S_0)}{\partial S_0}, \quad (5.8)$$

where $f''(z) = \frac{\partial^2 f(z)}{\partial z^2}$.

In practice, the integrals (5.6) and (5.8) can be computed by utilizing numerical techniques discussed in Section 3.4. Apart from the delta and gamma, expressions for other sensitivities with respect to $\beta = \theta$ can be obtained via dependence of the characteristic function in (5.4) on θ .

Two comments are in order. Firstly, formulations (5.6) and (5.8) can be used to compute sensitivities of certain path-dependent options, too. These include discretely sampled geometric Asian options, whose payoffs depend on $G_n = (\prod_{k=1}^n S_k)^{1/n}$ where $\{S_k\}_{k=1}^n$ are the values of the underlying asset recorded at times $\{t_k\}_{k=1}^n$ on the time line $[0, T]$. The distribution law of G_n is known explicitly for Gaussian asset log-returns and implicitly via its characteristic function for log-return models from the affine class (see Section 6.6).

Secondly, we get from (5.6)

$$-S_0^{-1} \int_{\mathbb{R}} p(\ln S_0 + z) f'(z) dz = \int_{\mathbb{R}} p(\ln S_0 + z) \ell(z; S_0) f(z) dz = \mathbb{E}(p(\ln S_0 + Z_T) \ell(Z_T; S_0)), \quad (5.9)$$

where

$$\ell(z; S_0) = \frac{-1}{S_0} \frac{f'(z)}{f(z)}.$$

Broadie and Glasserman (1996) use expression (5.9) in order to compute via Monte Carlo the so-called likelihood ratio (LR) estimate for the (forward) delta, which is unbiased. Similarly, from (5.8), the (forward) gamma is given by

$$\mathbb{E} \left(p(\ln S_0 + Z_T) \tilde{\ell}(Z_T; S_0^2) \right),$$

where

$$\tilde{\ell}(z; S_0^2) = \frac{1}{S_0^2} \frac{f'(z) + f''(z)}{f(z)}. \quad (5.10)$$

In this section, we have derived integral representations for the price sensitivities in the distribution-based setting. These are general in that they can be adapted to a wide range of European payoffs, including payoffs which depend on the geometric average asset price, and models for the asset log-returns from the affine class. In contrast to the geometric average, the true distribution of the arithmetic average is not known, therefore (5.6) and (5.8) cannot be used to calculate the indicated sensitivities for an arithmetic Asian option.

5.3 A convolution approach for the Asian option price sensitivities

In this section, we will apply direct differentiation with respect to parameters of interest to extend the backward price convolution of Černý and Kyriakou (2010) (see also Chapter 4) for discretely sampled arithmetic Asian options to the computation of price sensitivities.

Let $n \in \mathbb{N}^*$ be the number of sampling dates of the Asian option, and $\{Z_k\}_{k=1}^n$ a collection of independent random variables representing the log-returns on asset S , such that

$$S_j = S_0 \exp\left(\sum_{k=1}^j Z_k\right), \quad j = 1, \dots, n, \quad S_0 > 0.$$

Further recall from Section 4.3 the Markov process

$$\begin{aligned}\ln Y_k &= \ln(Y_{k-1} + \lambda_{n+1-k}) + Z_{n+1-k}, \quad 1 < k \leq n, \\ \ln Y_1 &= \ln \lambda_n + Z_n\end{aligned}$$

defined under some risk-neutral measure $\mathbb{P}$ in filtration $\mathbb{G} = \{\mathcal{G}_k\}_{k=1}^n$ with

$$\mathcal{G}_k = \sigma\{Z_n, Z_{n-1}, \dots, Z_{n+1-k}\}.$$

Coefficients $\{\lambda_k\}_{k=0}^n$ are presented in Table 4.1 for different types of Asian options.

Set $\beta = \{S_0, \theta\}$, where $\theta > 0$ is a generic parameter of the probability distribution of $\{Z_k\}$. Also, $\theta \neq S_0$. We define delta, Δ , and gamma, Γ , of the option the first- and second-order price sensitivities with respect to $\beta = S_0$, respectively. We further define theta, ϑ , the first-order price sensitivity with respect to $\beta = \theta$.¹ With the focus on an Asian call option with fixed strike price, we derive in the next theorem backward recursive convolutions for its (forward) price sensitivities.

Theorem 18 *Assume parameter $\beta = \{S_0, \theta\}$, such that $\theta \neq S_0$. Assume that, for all k , Z_{n+1-k} has density $\tilde{f}_k(z; \theta)$. Consider constants $\lambda_k > 0$, $0 < k \leq n$, and $\lambda_0 \in \mathbb{R}$. For $\xi \equiv \{\Delta, \Gamma, \vartheta\}$, define*

$$m = \begin{cases} n, & \xi \equiv \Delta \\ n-1, & \xi \equiv \Gamma \\ n, & \xi \equiv \vartheta, \end{cases}$$

¹Sensitivity theta should not be confused with the sensitivity with respect to the time to maturity of the option, which is sometimes termed this way in the literature.

and functions ξ_k^p , for $0 < k \leq m$, ξ_k^q , for $0 \leq k < m$, and h_k , for $1 \leq k < n$, satisfying

$$\xi_m^p(y; \beta) = \begin{cases} \left(e^y + \frac{\gamma}{n+\gamma}\right) 1_{\{y > \ln(-\lambda_0)\}}, & \xi \equiv \Delta \\ -\frac{\left(\frac{\gamma}{n+\gamma} - \lambda_0\right)^2}{S_0 \lambda_0} \tilde{f}_n\left(\ln\left(-\frac{\lambda_0}{e^y + \lambda_1}\right); \theta\right), & \xi \equiv \Gamma \\ 0, & \xi \equiv \vartheta, \end{cases} \quad (5.11)$$

$$\xi_{k-1}^q(x; \beta) = \int_{\mathbb{R}} \left(\xi_k^p(x+z; \beta) \tilde{f}_k(z; \theta) + p_k(x+z; \theta) \frac{\partial \tilde{f}_k(z; \theta)}{\partial \theta} 1_{\{\xi \equiv \vartheta\}} \right) dz \quad (5.12)$$

$$= \left(\xi_k^p * \tilde{f}_k(-z) \right)(x; \beta) + \left(p_k * \frac{\partial \tilde{f}_k(-z)}{\partial \theta} \right)(x; \beta) 1_{\{\xi \equiv \vartheta\}}, \quad 0 < k \leq m, \quad (5.13)$$

$$h_{k-1}(y) = \ln(e^y + \lambda_{n+1-k}), \quad 1 < k \leq n,$$

$$\xi_{k-1}^p(y; \beta) = \xi_{k-1}^q(h_{k-1}(y); \beta), \quad 1 < k \leq n, \quad (5.14)$$

where $*$ in (5.13) denotes a convolution (see Definition 4).

Then, the forward delta, gamma, and theta of the Asian call option with fixed strike price are given by

$$\xi_0^q(\ln \lambda_n; \beta),$$

for $\xi \equiv \Delta$, $\xi \equiv \Gamma$, and $\xi \equiv \vartheta$ respectively.

Proof. From Černý and Kyriakou ((2010), Theorem 3.1) (see also Section 4.5), the forward price of the Asian call option with fixed strike is provided through the recursion

$$p_n(y; \beta) = S_0(e^y + \lambda_0)^+, \quad (5.15)$$

$$q_{k-1}(x; \beta) = \int_{\mathbb{R}} p_k(x+z; \beta) \tilde{f}_k(z; \theta) dz, \quad 0 < k \leq n, \quad (5.16)$$

$$p_{k-1}(y; \beta) = q_{k-1}(h_{k-1}(y); \beta), \quad 1 < k \leq n. \quad (5.17)$$

The forward option price is given by $q_0(\ln \lambda_n; \beta)$.

Assume that for all $x, z \in \mathbb{R}$, the derivative $\partial(p_k(x+z; \beta) \tilde{f}_k(z; \theta))/\partial \beta$ exists for all $\beta \in \mathcal{B} \subseteq \mathbb{R}$, for all k . If there exist functions F_k , for all k , such that $|\partial(p_k(x+z; \beta) \tilde{f}_k(z; \theta))/\partial \beta| \leq F_k(x, z)$ for all $\beta \in \mathcal{B}$ with $\int_{\mathbb{R}} |F_k(x, z) \tilde{f}_k(z)| dz < \infty$, then, by Talvila ((2001), Corollary 8), we get

$$\frac{\partial}{\partial \beta} \int_{\mathbb{R}} p_k(x+z; \beta) \tilde{f}_k(z; \theta) dz = \int_{\mathbb{R}} \frac{\partial(p_k(x+z; \beta) \tilde{f}_k(z; \theta))}{\partial \beta} dz, \quad (5.18)$$

for all $\beta \in \mathcal{B}$. Define $\xi_k^p(y; \beta) := \frac{\partial p_k(y; \beta)}{\partial \beta}$, $\xi_k^q(x; \beta) := \frac{\partial q_k(x; \beta)}{\partial \beta}$ such that

$$\xi \equiv \begin{cases} \Delta, & \beta = S_0 \\ \vartheta, & \beta = \theta. \end{cases}$$

Given (5.18), differentiation on both sides of (5.16) with respect to $\beta = \{S_0, \theta\}$ yields

$$\begin{aligned} \xi_k^q(x; \beta) &= \frac{\partial q_{k-1}(x; \beta)}{\partial \beta} = \int_{\mathbb{R}} \left(\frac{\partial p_k(x+z; \beta)}{\partial \beta} \tilde{f}_k(z; \theta) + p_k(x+z; \theta) \frac{\partial \tilde{f}_k(z; \theta)}{\partial \theta} 1_{\{\beta=\theta\}} \right) dz \\ &= \int_{\mathbb{R}} \left(\xi_k^p(x+z; \beta) \tilde{f}_k(z; \theta) + p_k(x+z; \theta) \frac{\partial \tilde{f}_k(z; \theta)}{\partial \theta} 1_{\{\xi=\vartheta\}} \right) dz. \end{aligned} \quad (5.19)$$

Also, differentiating both sides of (5.17) with respect to β yields

$$\xi_{k-1}^p(y; \beta) = \frac{\partial p_{k-1}(y; \beta)}{\partial \beta} = \frac{\partial q_{k-1}(h_{k-1}(y); \beta)}{\partial \beta} = \xi_{k-1}^q(h_{k-1}(y); \beta). \quad (5.20)$$

Representations for $\Gamma_k^p(y; S_0) := \frac{\partial \Delta_k^p(y; S_0)}{\partial S_0}$, $\Gamma_k^q(x; S_0) := \frac{\partial \Delta_k^q(x; S_0)}{\partial S_0}$ follow by differentiating (5.19) and (5.20), for $\xi \equiv \Delta$, with respect to S_0 . The forward option sensitivities are then given for each $\xi \equiv \{\Delta, \Gamma, \vartheta\}$ by $\xi_0^q(\ln \lambda_n; \beta)$.

We derive terminal (forward) sensitivity functions as follows. Set $\xi \equiv \Delta$. We have $\lambda_0 = \frac{\gamma}{n+\gamma} - \frac{K}{S_0}$ with derivative $\frac{\partial \lambda_0}{\partial S_0} = \frac{K}{S_0^2} = \frac{1}{S_0} \left(\frac{\gamma}{n+\gamma} - \lambda_0 \right)$. Hence, from (5.15),

$$\Delta_n^p(y) = \frac{\partial p_n}{\partial S_0}(y) = \left(e^y + \lambda_0 + S_0 \frac{\partial \lambda_0}{\partial S_0} \right) 1_{\{y > \ln(-\lambda_0)\}} = \left(e^y + \frac{\gamma}{n+\gamma} \right) 1_{\{y > \ln(-\lambda_0)\}}$$

for $\lambda_0 < 0$. Furthermore, from (5.12),

$$\Delta_{n-1}^q(x) = \int_{\mathbb{R}} \Delta_n^p(x+z) \tilde{f}_n(z) dz = \int_{\alpha(x; \lambda_0)}^{\infty} \left(e^{x+z} + \frac{\gamma}{n+\gamma} \right) \tilde{f}_n(z) dz,$$

where $\alpha(x; \lambda_0) = \ln(-\lambda_0) - x$. Applying Leibniz's integral rule yields the terminal function for

$\xi \equiv \Gamma$

$$\begin{aligned}\Gamma_{n-1}^q(x) &= \frac{\partial \Delta_{n-1}^q}{\partial S_0}(x) = -\frac{\partial \alpha(x; \lambda_0)}{\partial S_0} \tilde{f}_n(\alpha(x; \lambda_0)) \left(e^{x+\alpha(x; \lambda_0)} + \frac{\gamma}{n+\gamma} \right) \\ &= -\frac{\left(\frac{\gamma}{n+\gamma} - \lambda_0 \right)^2}{\lambda_0 S_0} \tilde{f}_n(\alpha(x; \lambda_0)) = -\frac{\left(\frac{\gamma}{n+\gamma} - \lambda_0 \right)^2}{\lambda_0 S_0} \tilde{f}_n(\ln(-\lambda_0) - x).\end{aligned}$$

From (5.14), we get

$$\Gamma_{n-1}^p(y) = \Gamma_{n-1}^q(\ln(e^y + \lambda_1)) = -\frac{\left(\frac{\gamma}{n+\gamma} - \lambda_0 \right)^2}{\lambda_0 S_0} \tilde{f}_n \left(\ln \left(-\frac{\lambda_0}{e^y + \lambda_1} \right) \right).$$

For $\xi \equiv \vartheta$, we get from (5.15)

$$\vartheta_n^p(y) = \frac{\partial p_n}{\partial \theta}(y) = 0,$$

which completes the proof. ■

Note from (5.11) that the delta and theta recursive algorithms start at $k = n$, as opposed to the gamma which starts at $k = n - 1$. This applies to the gamma because of insufficient smoothness of the payoff function, which we can, as usual, compensate by smoothness of the density function. This clearly shows the advantage of using the distribution-based approach for computing sensitivities of non-smooth payoffs.

In proving Theorem 18, we have taken for granted that equality (5.18) holds. With the focus on the option delta, we derive in the following theorem (forward) delta bounds needed to justify the differentiation with respect to S_0 under the integral sign.

Theorem 19 *Set $\xi \equiv \Delta$. Define $\tilde{\mu}_k = \int_{\mathbb{R}} z \tilde{f}_k(z) dz$. Then, there exist positive constants a_k, b_k , for all k , such that*

$$\begin{aligned}0 &\leq \Delta_k^p(y) \leq a_k e^y + b_k, \\ 0 &\leq \Delta_k^q(x) \leq a_k e^x + b_{k+1},\end{aligned}$$

for all $x, y \in \mathbb{R}$, with

$$a_n = 1, b_n = \frac{\gamma}{n+\gamma},$$

$$a_k = a_{k+1}\tilde{\mu}_{k+1},$$

$$b_k = a_k\lambda_{n-k} + b_{k+1}.$$

Proof. The proof of the theorem follows that of Černý and Kyriakou ((2010), Theorem 3.1). ■

From Theorem 19, function $0 \leq \Delta_k^p(y)$ is dominated by an integrable function for S_0 in a compact interval. Therefore, equality (5.18) is valid for $\beta = S_0$. The derivation of integrable bounds for the (forward) gamma and theta is nontrivial, in the first case, by nature of Γ_{n-1}^p and, in the second case, by dependence of the density $\tilde{f}_k$ on the parameter θ which gives rise to an additional (not necessarily positive) term in (5.12). We postpone further investigation on this to a later stage of our research.

By virtue of Theorems 5 and 7, we rephrase (5.13) as

$$\xi_{k-1}^q(x) = \mathcal{F}^{-1}(\mathcal{F}(\xi_{k-1}^q))(x), \quad (5.21)$$

where $\mathcal{F}$ denotes the Fourier transform (see Section 2.2), and

$$\begin{aligned} \mathcal{F}(\xi_{k-1}^q)(u) &= \mathcal{F}(\xi_k^p * \tilde{f}_k(-z; \theta))(u) + \mathcal{F}\left(p_k * \frac{\partial \tilde{f}_k(-z; \theta)}{\partial \theta}\right)(u) 1_{\{\xi \equiv \vartheta\}} \\ &= \mathcal{F}(\xi_k^p)(u) \tilde{\varphi}_k(u; \theta) + \mathcal{F}(p_k)(u) \mathcal{F}\left(\frac{\partial \tilde{f}_k(-z; \theta)}{\partial \theta}\right)(u) 1_{\{\xi \equiv \vartheta\}}. \end{aligned} \quad (5.22)$$

We denote by $\tilde{\varphi}$ the complex conjugate of the characteristic function $\tilde{\phi}$ of Z_{n+1-k}

$$\tilde{\varphi}_k(u; \theta) = \tilde{\phi}_k(-u; \theta) = \mathbb{E}(e^{-iuZ_{n+1-k}}) = \int_{\mathbb{R}} e^{-iuz} \tilde{f}_k(z; \theta) dz.$$

Lemma 20 *If there exist functions g_k , for all k , such that $|\frac{\partial \tilde{f}_k(z; \theta)}{\partial \theta}| \leq g_k(z)$ for all $z \in \mathbb{R}$ and*

all θ in an open neighborhood of θ_0 with $\int_{\mathbb{R}} |g_k(z) \tilde{f}_k(z)| dz < \infty$, then

$$\int_{\mathbb{R}} e^{-iuz} \frac{\partial \tilde{f}_k(z; \theta)}{\partial \theta} dz = \frac{\partial}{\partial \theta} \int_{\mathbb{R}} e^{-iuz} \tilde{f}_k(z; \theta) dz = \frac{\partial \tilde{\varphi}_k(u; \theta)}{\partial \theta}$$

for $\theta = \theta_0$.

Proof. See Talvila ((2001), Corollary 8). ■

Providing that the assumptions of Lemma 20 are satisfied, equation (5.22) then reads

$$\mathcal{F}(\xi_{k-1}^q)(u) = \mathcal{F}(\xi_k^p)(u) \tilde{\varphi}_k(u; \theta) + \mathcal{F}(p_k)(u) \frac{\partial \tilde{\varphi}_k(u; \theta)}{\partial \theta}(u) 1_{\{\xi \equiv \vartheta\}}.$$

In practice, the continuous Fourier transforms in (5.21) and (5.22) are approximated by their discrete analogues, which are implemented as described in Section 4.5.1. Furthermore, as indicated in (5.22), the computation of the theta ($\xi \equiv \vartheta$) requires that we run successive sessions for the price and the theta at each iteration, relatively doubling the total number of operations to carry out. In practice, however, the ultimate output of the scheme comprises computed values for both. Thus, in absolute terms, the computational effort for all the sensitivities remains of the same order as for the price.

5.4 Computation of the price sensitivities via Monte Carlo simulation

A natural way to calculate price sensitivities of contingent claims is via finite difference approximation. As mentioned earlier, this technique requires us to bump the initial value of the parameter in question and recalculate the price of the contingent claim. The finite difference method has been applied mainly in the past to the estimation of price sensitivities via Monte Carlo. For smooth payoff functions, L'Ecuyer and Perron (1994) and Boyle et al. (1997) show that the more positively correlated the two price estimates (obtained on common random number streams), the more efficient the sensitivity estimate obtained via a central finite difference is expected to be. Instead, lack of smoothness causes large bias and poor convergence to the true value.

The need for unbiased sensitivity estimators has motivated the works of Curran (1994b), (1998) and Broadie and Glasserman (1996) and led to the derivation of unbiased estimators for the sensitivities based on direct differentiation of smooth payoff functions (PW technique). For insufficiently smooth payoffs, Broadie and Glasserman (1996) suggest differentiating instead the density of the underlying asset vector which is typically smooth function of its parameters (LR technique). To obtain estimators for second-order sensitivities, one way is to apply the LR technique twice (e.g., see equation (5.10) for the gamma), which is, however, anticipated to have high variance (the same observation applies to all LR estimators compared to the PW estimators, as reported in Glasserman (2004)). To reduce the variance, it is advisable to utilize each of the PW and LR approaches for one order of differentiation to generate mixed PW-LR or LR-PW estimators.

To overcome frequent issues related to non-differentiability of the payoff function and lack of explicit knowledge of the density function of the underlying, Fournié et al. (1999), (2001) introduce a method based on the Malliavin calculus theory. This method allows us to express the estimators for the Greeks as products of the contract payoff with weight functions which are identified with certain Skorokhod integrals. The Malliavin estimators are not unique and subject to different variance, depending on the weight choice. Generating a suitable weight function raises a nontrivial concern in this setup; Benhamou (2003) provides necessary and sufficient conditions for Skorokhod-integrable functions to serve as weight function generators, and finds the weight function with minimal variance. It is proved that the minimal-variance weight is the one given by the LR method, implying that the Malliavin estimators are not superior to the ones given by the LR and PW methods. In terms of efficiency, the Malliavin method suffers in the simulation of Skorokhod integrals, which may become substantially time-consuming due to the simulation of many auxiliary processes, especially for a non-lognormal underlying. Furthermore, to maintain variance at low levels, the payoff must be restricted to small values, implying greater efficiency of the Malliavin formulae for put rather than call options. Selected references include Fournié et al. (1999), (2001) and Benhamou (2000) for results on Asians, Gobet and Kohatsu-Higa (2003) and Bernis et al. (2003) for barriers and lookbacks, and Kohatsu-Higa and Yasuda (2009) for a general review of the method with applications.

Due to large variance of the LR estimators, we limit hereafter their use to mixed estimators for the second-order sensitivities, while we utilize PW estimators for the first-order sensitivities. In Tables 5.1 and 5.2, we present the (forward) delta, vega and gamma estimators for discretely sampled arithmetic and geometric Asian options. Here, vega refers to the first-order price sensitivity with respect to the asset volatility σ in the Black-Scholes-Merton model, i.e., $\theta = \sigma$.² We do not consider the Malliavin estimators any further, whose efficiency fades away as we depart from the continuous monitoring case.

5.4.1 Likelihood ratio estimators coupled with Fourier transforms

In general, the pure LR and mixed estimators require the existence of the joint density function of the values of the underlying asset recorded at discrete points $\{t_k\}_{k=1}^n$ on the time line $[0, T]$. Wherever *independent* increments apply, the joint density simplifies to the product of the transition densities describing the law of successive asset movements. However, it may be the case that these densities are not easy or fast to compute (e.g., normal inverse Gaussian, variance gamma laws), or they are only known through their characteristic functions (e.g., tempered stable law). To retrieve the unknown densities, one must resort to numerical inversion of the associated characteristic functions.

Consider a discretely sampled Asian option written on asset $S_t = S_0 e^{Z_t}$, $S_0 > 0$, where Z_t follows a Lévy process. A basic application of the LR/mixed technique is on the gamma of this option. From Table 5.2, the mixed estimator for the gamma depends on the quantity

$$\ell(z; S_0) = \frac{-1}{S_0} \frac{f'_1(z)}{f_1(z)}, \quad (5.23)$$

where f_1 is the density of $Z_1 = \ln \frac{S_1}{S_0}$ and $f'_1(z) = \frac{\partial f_1(z)}{\partial z}$. Quantity (5.23) is known as the delta score function. For a well-behaved density f_1 in the sense of Theorem 7, we get

$$f_1 = \mathcal{F}^{-1}(\phi_1), \quad (5.24)$$

²Vega is defined originally in a Gaussian economy. In practice, this terminology is used also in non-Gaussian Lévy economies, e.g., NIG economy (see equation (4.14)), to describe sensitivities with respect to parameter σ . Note that in these cases, σ does not represent the asset volatility anymore. Here, we adhere to the original definition.

PW estimators for the forward price sensitivities of Asian options		
“Nonzero” condition	forward delta	forward vega
$\zeta(A_n - K) > 0$	$\zeta \frac{A_n}{S_0}$	$\zeta \frac{1}{(n+\gamma)\sigma} \sum_{k=1}^n S_k (\ln \frac{S_k}{S_0} - (r + \frac{\sigma^2}{2})t_k)$
$\zeta(G_n - K) > 0$	$\zeta \frac{G_n}{S_0}$	$\zeta \frac{G_n}{\sigma} (\ln \frac{G_n}{S_0} - \frac{n+1}{2(n+\gamma)}(r + \frac{\sigma^2}{2})T)$

Table 5.1: PW estimators for the forward delta and vega (vega: Black-Scholes-Merton model) of discretely sampled Asian options (reference: Boyle and Potapchik (2008), Theorems 3.1-3.2). r : risk-free interest rate. “Nonzero”condition: indicates the option payoff providing the condition is satisfied, otherwise both payoff and sensitivity estimators take value 0. Averages: arithmetic $A_n = \frac{1}{n+\gamma}(\gamma S_0 + \sum_{k=1}^n S_k)$, geometric $G_n = (S_0^\gamma \prod_{k=1}^n S_k)^{1/(n+\gamma)}$, where coefficient γ takes value 1 (0) when S_0 is (is not) included in the average, and coefficient ζ takes value 1 (-1) for a call (put) option.

Mixed forward gamma estimators for Asian options		
“Nonzero” condition	PW-LR forward gamma	LR-PW forward gamma
$\zeta(A_n - K) > 0$	$\zeta \frac{A_n}{S_0} (\ell(Z_1; S_0) - \frac{1}{S_0})$	$\zeta \frac{K}{S_0} \ell(Z_1; S_0)$
$\zeta(G_n - K) > 0$	$\zeta \frac{G_n}{S_0} (\ell(Z_1; S_0) - \frac{1}{S_0})$	$\zeta \frac{K}{S_0} \ell(Z_1; S_0)$

Table 5.2: Mixed forward gamma estimators; PW-LR: PW followed by LR, LR-PW: LR followed by PW (reference: Glasserman (2004), Section 7.3.3). $\ell(Z_1; S_0)$ given by equation (5.23). “Nonzero”conditions, A_n , G_n , ζ as in Table 5.1.

where $\phi_1(u) = \mathbb{E}(e^{iuZ_1}) = \int_{\mathbb{R}} e^{iuz} f_1(z) dz$. Score functions applying to other LR estimators (see Glasserman (2004), Section 7.3.1) can be obtained using the same means.

Numerical implementation

Pre-caching the score function. To evaluate numerically the density $f_1 = \mathcal{F}^{-1}(\phi_1)$, we select uniform grids $\mathbf{u} = \{u_0 + j\delta u\}_{j=0}^{N-1}$ (symmetric about zero) and $\mathbf{z} = \{z_0 + l\delta z\}_{l=0}^{N-1}$ with N grid points and spacings δu and δz . The range of values of $\mathbf{u}$ is chosen to ensure that $|\phi_1| < \epsilon$ outside $\mathbf{u}$ for some tolerance level ϵ , e.g., $\epsilon = 10^{-15}$. We denote the values of ϕ_1 on the grid $\mathbf{u}$ by ϕ_1 , and evaluate the inverse transform (5.24) on the grid $\mathbf{z}$ by computing

$$\mathbf{f}_1 = \frac{1}{2\pi} \mathcal{D}(\phi_1, -\mathbf{u}, \mathbf{z}; -N\delta u\delta z/2\pi)\delta u, \quad (5.25)$$

following the conversion (2.24). We refer to Hughett ((1998), Theorem 7) for an explicit bound to the error induced by the discrete transform approximation (5.25). For consistency, one can then inspect on the sign of the numerically retrieved density, check that it integrates to the unity and that the moments calculated by numerical quadrature agree with their true values. Given the values $\mathbf{f}_1$, we approximate $\frac{\partial f_1}{\partial \mathbf{z}}$ on the grid $\mathbf{z}$ by $\frac{\Delta \mathbf{f}_1}{\Delta \mathbf{z}}$ using finite differences³.

Simulating the score. Given the pre-tabulated density function and its first-order derivative, we generate score samples, for m simulation trials, as summarized next:

1. Generate a m -dimensional vector of $\hat{\mathbf{z}}_1$ -draws from the distribution of Z_1 .
2. Approximate $f_1(\hat{\mathbf{z}}_1)$ and $\frac{\partial f_1}{\partial \mathbf{z}}(\hat{\mathbf{z}}_1)$ by fitting a cubic interpolating spline on the nodes $(\mathbf{z}, \mathbf{f}_1)$ and $(\mathbf{z}, \frac{\Delta \mathbf{f}_1}{\Delta \mathbf{z}})$ respectively, using the MATLAB built-in function INTERP1.
3. From (5.23), obtain the score samples $\ell(\hat{\mathbf{z}}_1; S_0)$.

5.5 Numerical study

For the purposes of this study, we opt for an Asian call option (in consistency with the notation in Table 4.1: $\zeta = 1$, $\gamma = 0$) with fixed strike price K , time to maturity $T = 1$ and sampling frequency $n = 50$. To model the log-returns, we employ the calibrated parameters from Section 4.6.1 to achieve target skewness coefficient -0.5 and excess kurtosis 0.7 of the (NIG and CGMY) log-returns for each volatility level $\{0.1, 0.3, 0.5\}$ under the risk-neutral measure. All numerical experiments are coded in MATLAB R2009a on a Sony Vaio VGN-AR31E Intel Core 2 Duo T5500 PC 1.66 GHz with 2.0 GB RAM.

We test our convolution method by computing the delta, gamma, and vega of arithmetic Asian options and comparing with the corresponding Monte Carlo estimates. For the delta and vega, we present only the PW estimates, since the corresponding LR ones have exhibited higher variance; this is consistent with the observation of Boyle and Potapchik (2008). In the case of

³For the numerical differentiation, we have implemented the *central_diff.m* routine for MATLAB which is available to download from <http://www.mathworks.com/matlabcentral/fileexchange/12-centraldiff-m/>. This evaluates the gradient numerically by utilizing a forward difference at the left end of the grid, a backward difference at the right end, and central differences in the interior.

the gamma, the PW-LR and LR-PW estimates have shown closely related variances. All the arithmetic option estimates are obtained in combination with the corresponding estimates for the geometric option, which are used as control variates. For the latter, we only require the characteristic function of the log-geometric average distribution law, which has been derived in Fusai and Meucci ((2008), Appendix A) for general Lévy log-returns. Exact geometric Asian option delta and gamma are obtained by evaluating numerically the integrals (5.5) and (5.7) using the Fourier-cosine series expansion (3.13), while a similar formula for the vega is straightforward to derive. The NIG increments are simulated using standard time-change Brownian representation for the NIG process, while for the CGMY process we utilize the joint Monte Carlo-Fourier transform scheme developed in the follow-up Chapter 7.

5.5.1 Convolution versus Monte Carlo

Tables 5.3 and 5.4 provide speed-accuracy comparisons between the convolution and control variate Monte Carlo (CVMC) methods in calculating the option price sensitivities. Across strikes and asset volatilities, CVMC requires for each of the normal, NIG, and CGMY models 2.3, 7.0, and 27.6 seconds (resp. 4.9, 9.4, and 27.8 seconds) to achieve delta estimates (gamma estimates) with best precision $\pm 10^{-4}$ and worst precision $\pm 10^{-3}$ (at 99% confidence level). For high volatilities ($\text{vol} \in \{0.3, 0.5\}$), the backward convolution method appears clear winner taking for each model at most 0.8, 2.7, and 2.7 seconds to guarantee 6 decimal-place accuracy for the delta and gamma. For $\text{vol} = 0.1$, we obtain accurate results to 4 decimal places in 0.8, 5.3 and 5.3 seconds for each model. CVMC performs worst in the vega requiring 3.0 and 27.0 seconds for precision $\pm 10^{-2}$ and $\pm 10^{-3}$ respectively, whereas our method at most 5.1 seconds for 6 decimal-place precision.

Furthermore, comparing the deltas (resp. gammas) from the convolution scheme for the non-normal (NIG and CGMY) log-returns, we find that these match to at least the third (resp. fourth) decimal place. The normal deltas and gammas match the corresponding non-normal ones to the second and third decimal places respectively. This suggests that the negative skewness and excess kurtosis in the risk-neutral distribution of the non-normal log-returns have only minor impact on these sensitivities. The joint skewness-kurtosis effect becomes more important in the option prices where the normal versus non-normal gap becomes wider, as

discussed in Section 4.6.3.

5.6 Concluding remarks

The present chapter extends the backward price convolution method of Chapter 4 to the calculation of the Asian option price sensitivities. When tested on a range of strikes, log-return models and model parameter values, the numerical scheme exhibits regular convergence in the number of grid points ensuring high-level precision. This further allows for a cautious study on the option delta and gamma, which makes obvious that the effect of the model choice is limited, as opposed to the more significant impact this has on the option price. Furthermore, a numerical comparison against an effective control variate Monte Carlo method illustrates the speed and accuracy of our approach.

model	vol	K			CPU (s)
		90	100	110	
Conv. delta					
normal	0.1	0.966536	0.632629	0.105740	1.4
	0.3	0.770593	0.563355	0.356690	0.8
	0.5	0.695404	0.563825	0.438759	0.4
NIG	0.1	0.956400	0.671022	0.080918	12.3
	0.3	0.799574	0.598537	0.362400	2.7
	0.5	0.731009	0.592983	0.448609	1.3
CGMY	0.1	0.956498	0.670591	0.081181	12.2
	0.3	0.799228	0.598308	0.362797	2.7
	0.5	0.730686	0.592920	0.448945	2.7
Conv. gamma					
normal	0.1	0.0062364	0.0622252	0.0304433	1.4
	0.3	0.0165627	0.0218205	0.0205350	0.8
	0.5	0.0116627	0.0130656	0.0129548	0.4
NIG	0.1	0.0065512	0.0596221	0.0313875	12.3
	0.3	0.0144623	0.0232580	0.0246266	2.7
	0.5	0.0114808	0.0144927	0.0154340	1.3
CGMY	0.1	0.0065464	0.0596070	0.0315333	12.2
	0.3	0.0144825	0.0232080	0.0245690	2.7
	0.5	0.0114704	0.0144549	0.0153935	2.7
Conv. vega					
normal	0.1	2.060207	21.404443	10.847111	5.1
	0.3	16.373056	22.460832	21.894896	1.4
	0.5	19.118652	22.302081	22.905267	1.4

Table 5.3: Fixed-strike Asian call option sensitivities ($\zeta = 1$, $\gamma = 0$, $T = 1$, $n = 50$) for Lévy log-returns: results of the extended backward convolution scheme. “vol”: standard deviation of log-returns for a one-year time horizon. Model parameters: Table 4.2. Other parameters: $r = 0.04$, $S_0 = 100$. CPU times in seconds (s).

model	vol	K						CPU (s)
		90		100		110		
		CVMC	std $\times 10^{-5}$	CVMC	std $\times 10^{-5}$	CVMC	std $\times 10^{-5}$	
CVMC delta								
normal	0.1	0.96651	9	0.63255	17	0.10594	22	2.3
	0.3	0.77060	28	0.56285	30	0.35691	34	
	0.5	0.69580	38	0.56379	40	0.43914	44	
NIG	0.1	0.95641	11	0.67110	17	0.08091	21	7.0
	0.3	0.79941	28	0.59793	29	0.36236	33	
	0.5	0.73059	37	0.59261	38	0.44839	40	
CGMY	0.1	0.95676	11	0.67071	17	0.08130	21	27.6
	0.3	0.79914	28	0.59775	29	0.36309	33	
	0.5	0.73050	37	0.59290	38	0.44853	40	
CVMC gamma								
normal	0.1	0.00625	6	0.06226	13	0.03043	15	4.9
	0.3	0.01651	6.5	0.02199	7.5	0.02051	8	
	0.5	0.01161	5.3	0.01297	5.8	0.01293	6	
NIG	0.1	0.00659	15	0.05947	25	0.03135	33	9.4
	0.3	0.01432	14	0.02319	16	0.02482	17	
	0.5	0.01166	12	0.01460	12	0.01553	13	
CGMY	0.1	0.00669	15	0.05953	27	0.03168	34	27.8
	0.3	0.01437	15	0.02321	16	0.02467	18	
	0.5	0.01150	12	0.01465	13	0.01548	13	
CVMC vega								
normal	0.1	2.068	1200	21.410	530	10.864	1700	3.0
	0.3	16.387	1900	22.475	1200	21.903	1200	
	0.5	19.091	2800	22.266	2300	22.914	2100	

Table 5.4: Fixed-strike Asian call option sensitivities ($\zeta = 1$, $\gamma = 0$, $T = 1$, $n = 50$): results of control variate Monte Carlo (CVMC) strategy with 100,000 trials and Lévy log-returns. “vol”: standard deviation of log-returns for a one-year time horizon, “std”: standard error of CVMC estimator. Model parameters: Table 4.2. Other parameters: $r = 0.04$, $S_0 = 100$. CPU times in seconds (s) (excl. computational time for the true geometric Asian option sensitivities).

Chapter 6

Pricing Asian options under stochastic volatility

6.1 Introduction

Empirical studies of the time series of asset returns and derivatives prices conclude that there are at least three systematic departures from the benchmark lognormal process which describe the asset price dynamics under both the historical and risk-neutral measures. First, asset prices jump, leading to non-Gaussian return innovations. Second, the amplitude of returns is positively autocorrelated in time (volatility clustering). Third, returns and their volatilities are correlated, often negatively for the equity markets (leverage effect). Exponential Lévy models account only for the first stylized feature, and further perform poorly in generating implied volatility patterns across different time scales.

Stochastic volatility models tackle these difficulties at the cost of introducing a second random process, interpreted as the instantaneous volatility of the underlying. In this case, the asset price is no longer a Markov process; to regain a Markov process, one must consider the joint asset-volatility process. This dimensionality increase inevitably affects the computational complexity of various pricing procedures and becomes particularly pronounced for early-exercise and path-dependent derivatives. We consider here the case of Asian options.

In the context of affine jump diffusions with stochastic volatility, Duffie et al. (2000) are the first to deduce an integral representation for arithmetic Asians with continuous sampling of the

underlying price. In fact, their work complements the work of Bakshi and Madan (2000) on the pricing of average-rate interest rate options where the state vector follows a square-root model. Thereafter, a number of approximate PDEs for the option price and Monte Carlo strategies (for continuously sampled Asians) appear primarily in the literature. As we describe in greater detail in Section 6.2, both approaches assume a stochastic volatility diffusion and strongly rely on fast mean-reversion to ensure convergence.

Our contribution to the current state of the literature on Asians is twofold. Firstly, we expand here the original convolution method of Černý and Kyriakou (2010) (see also Chapter 4) based on Lévy log-returns to two dimensions to accommodate non-Lévy log-returns with stochastic volatility, developing on an idea from the presentation of Fang and Oosterlee (2009) applied on barrier and Bermudan vanilla options. The outcome of the algorithm then consists of the option values on a grid of initial variance values. Secondly, we derive the exact distribution law of the log-geometric average of the discrete asset values, and subsequently obtain the price of a geometric Asian in terms of a Fourier transform which we evaluate at high accuracy. Given this, we construct a control variate Monte Carlo strategy which we implement for efficiency testing against the convolution method. In particular, we apply in pricing an Asian put option with floating strike under the Heston (1993) and Bates (1996) stochastic volatility models with parameter values relevant to the equity option markets.

The rest of the chapter is organized as follows. In Section 6.2, we review previous approaches to pricing Asian options under stochastic volatility. In Section 6.3, we present the class of stochastic volatility models and their laws. In Section 6.4, we reconsider the Carverhill-Clelland-Hodges factorization and incorporate stochastic volatility. In Section 6.5, we develop the main theoretical results for the price convolution scheme, and, in Section 6.5.1, detail its implementation via discrete Fourier transform combined with quadrature. In Section 6.6, we obtain the price of the geometric Asian option, as part of our control variate Monte Carlo strategy for the arithmetic Asian option. Section 6.7 presents our numerical study and reports our results, and Section 6.8 concludes the chapter.

6.2 Pricing approaches to Asian options under stochastic volatility

In what follows, we review existing approaches to the computation of Asian option prices when the asset returns are driven by stochastic volatility diffusion models.

Part of the relevant literature is PDE-based and dominated by the work of Fouque and Han on the derivation and approximation of pricing PDEs. Recognizing that the underlying asset, the running arithmetic (continuous) sum of the asset values and the volatility form a joint Markov process, Fouque et al. (2000) originally derive a 3-D (in space) pricing PDE. By applying a singular perturbation asymptotic analysis for fast mean-reverting volatility, they show that the solution of the PDE can be approximated by the sum of two terms which satisfy themselves a pair of 2-D PDEs. Fouque and Han (2003), and later Fouque and Han (2004a) with a two-factor stochastic volatility, further adhere to the state-space reduction employed in Večer (2002) to reduce the original PDE dimension. Their approximate solution shows independence from the volatility and, in fact, consists of two terms which satisfy a pair of 1-D PDEs that are easier to solve. They conclude that the accuracy of their approximation is $O(1/\alpha)$, where $\alpha > 0$ is the mean-reversion rate of the volatility process.

As an alternative to the PDE method, Monte Carlo is often used due to its ability to handle path-dependence with relative flexibility. To reduce the standard error of the Monte Carlo estimator for the price of an arithmetic Asian option, Wong and Cheung (2004) extend the approach of Fouque et al. (2000) to derive a first-order price approximation based on asymptotic analysis in a one-factor Gaussian-OU volatility setup for a continuously sampled geometric Asian for use as control variate. Fouque and Han (2004b) generalize to two-factor volatility and propose a two-step variance reduction combining control variate and importance sampling. Apart from the approximation error in the asymptotic analysis for the non-Monte Carlo price of the geometric Asian option, further issues relate to the simulation of the sample trajectories which are left unexplored in these two works; in general, stochastic volatility models have proved difficult to simulate exactly. Euler discretization (see Schoutens (2003), Section 8.4) has been traditionally employed, subject, however, to a bias that needs to be estimated and also a time grid which is usually much finer, than is strictly necessary for the contract in

question, to keep the bias low. For our purposes here, we refer the reader to Section 6.6.2 for more on the recent advances on the simulation of the Heston model. Moreover, simulating the payoff of a continuously sampled Asian option on a finite set of dates further contributes to the overall magnitude of the bias.

With the focus on discrete monitoring, Albrecher and Schoutens (2005) use comonotonic theory (see Kaas et al. (2000), Dhaene et al. (2002)) to derive a static super-hedge for an arithmetic Asian option with fixed strike in terms of a portfolio of European plain vanilla options maturing at the monitoring dates of the Asian option. The performance of their hedging strategy appears highly dependent on the moneyness of the option, e.g., for out-of-the-money options the gap between their Monte Carlo price estimate and the comonotonic hedge exceeds 60% in the case of the Heston model.

To skip any form of bias, e.g., by discretization of the continuous-time variance process, asymptotic analysis for fast mean-reverting volatility, and comonotonicity, we suggest an accurate recursive pricing algorithm for discretely sampled Asian options based on numerical integration. First, we present popular stochastic volatility models.

6.3 Market models

Fix constant $S_0 > 0$, and define under the risk-neutral measure $\mathbb{P}$ the price process of a risky asset $S = e^X$. Assume constant, continuously compounded interest rate $r > 0$. Moreover, introduce stochastic variance v in the asset dynamics such that (v, X) are modelled by a bivariate affine process. Several stochastic volatility paradigms have been proposed in the option pricing literature: popular is the stochastic volatility Lévy framework developed by Carr et al. (2003). This comprises square-root time-change models of the form

$$X_t = X_0 + (r + \varpi)t + L_{\int_0^t v_s ds} + \varrho(v_t - v_0), \quad (6.1)$$

$$dv_t = \alpha(\beta - v_t)dt + \eta\sqrt{v_t}dW_t, \quad (6.2)$$

where the parameters $\alpha, \beta, \eta, \varpi, \varrho$ are constant, L is a Lévy process and W an independent standard Brownian motion, as well as Lévy-driven OU time-change models of the form

$$X_t = X_0 + (r + \varpi)t + L_{\int_0^t v_{s-} ds} + \tilde{L}_{z_t}, \quad (6.3)$$

$$dv_t = \alpha(\beta - v_{t-})dt + dz_t, \quad (6.4)$$

where α, β, ϖ are constant parameters and $L, \tilde{L}$ and z independent Lévy processes; z is chosen to be a subordinator without drift. Typically, v in (6.4) is driven by an OU process with gamma stationary distribution (Γ -OU) or inverse Gaussian stationary distribution (IG-OU). The mean-adjusting parameter ϖ is chosen to ensure that the martingale condition under $\mathbb{P}$,

$$\mathbb{E}(S_t) = S_0 e^{rt}, \quad (6.5)$$

is satisfied¹.

Special cases of (6.1-6.2) and (6.3-6.4) are respectively the Heston (1993) model with $L_t = \mu t + B_t$, $\mu \in \mathbb{R}$, and standard Brownian motion B , and the Barndorff-Nielsen and Shephard (2001) model with $L_t = \mu t + B_t$, deterministic process $\tilde{L}_t = \rho t$, $\rho \leq 0$, and long-term mean variance $\beta = 0$.² Although the Heston model is able to generate sufficient leverage effect, so as to obtain a skew at long time scales, it cannot give rise to realistic short-term implied volatility patterns. The Barndorff-Nielsen and Shephard model also restricts the skewness of the implied volatility patterns for short and long maturities by letting the same parameter ρ to control both the impact of jumps in the asset returns and the leverage effect. Favourably, Bates (1996) adds Poisson jumps in the asset price process of the original Heston. The existence of jumps which are kept separate from the leverage effect, allows to reproduce strong skews at short maturities. Skews for longer maturities result separately by “Heston’s leverage” effect.

¹The martingale architecture (6.5) is consistent with taking the risk-neutral asset price process by mean-correcting the ordinary exponential of the time-changed Lévy process. Carr et al. (2003) propose also an alternative approach where the martingale model for the discounted asset price is obtained using a stochastic exponential, resulting into a model which is martingale in the joint filtration of the asset price and the stochastic time change. As a result of their survey based on S&P 500 options, the exponential models proved to provide better fit to the actual prices than the stochastic exponential models when the square-root stochastic variance was assumed. For this, we adhere to the ordinary exponential.

²The general model structures (6.1-6.2) and (6.3-6.4) presented here are common in the research papers of Jan Kallsen, e.g., Kallsen and Pauwels (2009).

Duffie et al. (2000) further assume concurrent arrival of jumps in the asset price and variance processes with crosscorrelated sizes.

6.3.1 Laws of affine stochastic volatility models

Duffie et al. (2003) characterize the laws of affine-structure models. For the bivariate affine setup (v, X) , the characteristic function is of the form

$$\mathbb{E} \left(e^{iu_1 v_t + iu_2 X_t} \right) = e^{\Psi_0(u_1, u_2; t) + \Psi_1(u_1, u_2; t)v_0 + iu_2 X_0}, \quad t > 0, \quad (6.6)$$

where $\Psi_0, \Psi_1 : (\mathbb{C}^- \times \mathbb{R}) \times \mathbb{R}^+ \rightarrow \mathbb{C}$ (with $\mathbb{C}^- = \{u \in \mathbb{C} : \operatorname{Re} iu \leq 0\}$) are determined for each individual model as solutions to a system of generalized Riccati differential equations (see Section 3.3 for more). It is possible that Ψ_0, Ψ_1 admit closed-form expressions for certain models, including the Heston model from the square-root time-change class (e.g., Filipović and Mayerhofer (2009)) and the Γ -OU framework (see Kallsen et al. (2009), Example 4.2).

For our purposes here, we state explicitly the functions Ψ_0, Ψ_1 relevant to the Heston model

$$\Psi_0(u_1, u_2; t) = iu_2 \left(r - \frac{\rho}{\eta} \alpha \beta \right) t + \frac{\alpha \beta}{\eta^2} \left((\alpha - \omega_2) t - 2 \ln \left(\frac{\omega_1 e^{-\omega_2 t} - 1}{\omega_1 - 1} \right) \right),$$

where restricting the complex logarithm to its principal branch ensures that the characteristic function remains continuous (see Lord and Kahl (2008), Theorem 3), and

$$\begin{aligned} \Psi_1(u_1, u_2; t) &= \frac{(\alpha - iu_2 \rho \eta - \omega_2) - \omega_1 e^{-\omega_2 t} (\alpha - iu_2 \rho \eta + \omega_2)}{(1 - \omega_1 e^{-\omega_2 t}) \eta^2}, \\ \omega_1(u_1, u_2) &= \frac{\alpha - iu_2 \rho \eta - \omega_2 - iu_1 \eta^2}{\alpha - iu_2 \rho \eta + \omega_2 - iu_1 \eta^2}, \\ \omega_2(u_1, u_2) &= \sqrt{(\alpha - iu_2 \rho \eta)^2 + (u_2^2 + iu_2) \eta^2}, \end{aligned}$$

where the square root in the last equation is operated on the common convention that its real part is nonnegative, with coefficient ρ representing the instantaneous correlation between the variance and asset price processes. As a by-product, in the Bates framework with extra

independent Poisson jumps in the asset price, Ψ_0 modifies to

$$\begin{aligned} \Psi_0(u_1, u_2; t) = & iu_2 \left(r - \frac{\rho}{\eta} \alpha \beta - \lambda_J \left(e^{\mu_J + \frac{1}{2} \sigma_J^2} - 1 \right) \right) t \\ & + \frac{\alpha \beta}{\eta^2} \left((\alpha - \omega_2) t - 2 \ln \left(\frac{\omega_1 e^{-\omega_2 t} - 1}{\omega_1 - 1} \right) \right) + \lambda_J \left(e^{iu_2 \mu_J - \frac{1}{2} \sigma_J^2 u_2^2} - 1 \right) t, \end{aligned}$$

with jump intensity $\lambda_J > 0$ of the time-homogeneous Poisson process and jump size $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$.

6.4 Modelling on reduced state space with stochastic volatility

Consider the set of monitoring dates $\mathcal{T} = \{t_k\}_{k=0}^n$, $n \in \mathbb{N}^*$. Assume these dates are equidistantly spaced, so that $t_k - t_{k-1} = \delta t$ for $0 < k \leq n$, with $t_0 = 0$, $t_n = T > 0$. Let $\{Z_k\}_{k=1}^n$ be the collection of log-asset increments on the complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, where $\mathbb{F} = \{\mathcal{F}_k\}_{k=1}^n$, with $\mathcal{F}_0$ trivial, is the information filtration generated by $\{Z_k\}$, such that

$$S_j = S_0 \exp\left(\sum_{k=1}^j Z_k\right), \quad j = 1, \dots, n, \quad S_0 > 0.$$

The increments $\{Z_k\}$ are identified later with specific non-Lévy increments within the class of stochastic volatility models of Section 6.3.

We recall from Chapter 4 that the pricing of arithmetic Asian options with discrete monitoring amounts to calculating

$$\mathbb{E}(e^{-rt_n} \Lambda_n^+), \tag{6.7}$$

where $\Lambda_k = \sum_{j=0}^k \lambda_j S_j$ for some deterministic process λ (see Table 4.1). Upon assuming stochastic variance v , the expectation (6.7) can be computed recursively on the three-dimensional (excluding time) grid (v, S, Λ) . To reduce dimensionality, we adopt from Section 4.3 the process $\bar{Y}_k = \bar{X}_k - \lambda_k$ in filtration $\mathbb{F}$

$$\ln \bar{Y}_k = \ln(\bar{Y}_{k-1} + \lambda_{k-1}) + Z_k^*, \quad 1 < k \leq n, \tag{6.8}$$

$$\ln \bar{Y}_1 = \ln \lambda_0 + Z_1^*, \tag{6.9}$$

where $Z^* = -Z$ and $\lambda_k > 0$ for $k = 0, \dots, n-1$. This allows us to express the option price as

$$\mathbb{E}(e^{-rt_n} \Lambda_n^+) = \mathbb{E}(e^{-rt_n} (S_n \bar{Y}_n + S_n \lambda_n)^+) = S_0 \bar{\mathbb{E}}((\bar{Y}_n + \lambda_n)^+), \quad (6.10)$$

where the first equality follows from (6.8) by recursive substitution, and the second by homogeneity of the payoff function of degree 1 and a change to the equivalent spot measure $\bar{\mathbb{P}}$ induced by taking the asset price as the numéraire (see Appendix 6.B). $(v, \bar{Y})$ form a joint Markov process in filtration $\mathbb{F}$, so that (6.10) can be evaluated recursively under $\bar{\mathbb{P}}$ on the two-dimensional grid $(v, \bar{Y})$.

6.5 The backward price convolution algorithm

Expectation $\bar{\mathbb{E}}((\bar{Y}_n + \lambda_n)^+)$ can be expressed iteratively in filtration $\mathbb{F}$ as

$$\bar{\mathbb{E}} \left(\bar{\mathbb{E}} \left(\dots \bar{\mathbb{E}} \left((\bar{Y}_n + \lambda_n)^+ \middle| \mathcal{F}_{n-1} \right) \dots \middle| \mathcal{F}_1 \right) \middle| \mathcal{F}_0 \right) \quad (6.11)$$

by virtue of the law of iterated expectations. Let v_k be the level of stochastic variance at t_k . Assume densities $\{\bar{f}_k\}_{k=1}^n$ characterizing the $\bar{\mathbb{P}}$ -law of $\{Z_k^*\}_{k=1}^n$ conditional on the variance levels v_k, v_{k-1} , and $\{\bar{f}_{k-1,k}^v\}_{k=1}^n$ the transition law of the variance v . We take $(v, \ln \bar{Y})$ as the state variables, and express (6.11) in terms of the following recursion for $1 < k \leq n$,

$$\begin{aligned} \bar{p}_n(\ln \bar{Y}_n, v_n) &= (e^{\ln \bar{Y}_n} + \lambda_n)^+ \text{ for } v_n \in \mathbb{R}^+, \\ \bar{h}_{k-1}(\ln \bar{Y}_{k-1}) &= \ln(e^{\ln \bar{Y}_{k-1}} + \lambda_{k-1}), \\ \bar{q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}), v_{k-1}) &= \bar{\mathbb{E}}(\bar{p}_k(\ln \bar{Y}_k, v_k) \middle| \mathcal{F}_{k-1}) \\ &= \bar{\mathbb{E}}(\bar{p}_k(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}) + Z_k^*, v_k) \middle| \mathcal{F}_{k-1}) \\ &= \int_{\mathbb{R}^+} \int_{\mathbb{R}} \bar{p}_k(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}) + z^*, y_v) \bar{f}_k(z^* | y_v, v_{k-1}) \bar{f}_{k-1,k}^v(y_v | v_{k-1}) dy_v dz^*, \\ \bar{p}_{k-1}(\ln \bar{Y}_{k-1}, v_{k-1}) &= \bar{q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}), v_{k-1}). \end{aligned} \quad (6.12)$$

The price of the option is given by

$$S_0 \bar{q}_0(\ln \lambda_0, v_0) \quad (6.13)$$

by virtue of (6.9). The proof of statement (6.13) follows that of Černý and Kyriakou ((2010), Theorem 3.1), hence we omit it here. In this chapter, we apply to the case of a floating-strike

Asian put with coefficients $\lambda_0 = \dots = \lambda_{n-1} = \frac{1}{n+1} > 0$ (see Table 4.1).

For the purposes of the numerical implementation (see Section 6.5.1), we further define the convolution (see Section 2.2.2)

$$\begin{aligned}\bar{Q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}); v_k, v_{k-1}) &= \int_{\mathbb{R}} \bar{p}_k(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}) + z^*, v_k) \bar{f}_k(z^* | v_k, v_{k-1}) dz^* \\ &= (\bar{p}_k * \bar{f}_k(-z^*))(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}); v_k, v_{k-1})\end{aligned}$$

for $v_k, v_{k-1} \in \mathbb{R}^+$, such that

$$\bar{q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}), v_{k-1}) = \int_{\mathbb{R}^+} \bar{Q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}); y_v, v_{k-1}) \bar{f}_{k-1,k}^v(y_v | v_{k-1}) dy_v. \quad (6.14)$$

We then obtain from Theorem 5 that

$$\mathcal{F}(\bar{p}_k * \bar{f}_k(-z^*)) = \mathcal{F}(\bar{p}_k) \mathcal{F}(\bar{f}_k(-z^*)) = \mathcal{F}(\bar{p}_k) \bar{\varphi}_k, \quad (6.15)$$

where

$$\mathcal{F}(\bar{p}_k)(u; v_k) = \int_{\mathbb{R}} e^{iuy} \bar{p}_k(y, v_k) dy$$

with v_k held fixed, and $\bar{\varphi}_k$ is the complex conjugate of the characteristic function $\bar{\phi}_k$ of Z_k^* conditional on the variance levels v_k, v_{k-1} at the endpoints t_k, t_{k-1}

$$\bar{\phi}_k(u; v_k, v_{k-1}) = \mathbb{E} \left(e^{iuZ_k^*} \middle| v_k, v_{k-1} \right) = \int_{\mathbb{R}} e^{iuz^*} \bar{f}_k(z^* | v_k, v_{k-1}) dz^*.$$

$\bar{\varphi}_k$ is derived in Appendix 6.C for the Heston and Bates models. From (6.15), we recover $\bar{Q}_{k-1}$ via

$$\bar{Q}_{k-1}(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}); v_k, v_{k-1}) = \mathcal{F}^{-1}(\mathcal{F}(\bar{p}_k) \bar{\varphi}_k)(\bar{h}_{k-1}(\ln \bar{Y}_{k-1}); v_k, v_{k-1}) \quad (6.16)$$

given the values of v_k, v_{k-1} .

6.5.1 Numerical implementation

Preliminaries. To evaluate numerically $\bar{p}_k$, $\bar{Q}_k$ and $\bar{q}_k$, we select uniform grids $\mathbf{y} = \{y_0 + l\delta y\}_{l=0}^{N-1}$, $\mathbf{x} = \{x_0 + m\delta x\}_{m=0}^{N-1}$ with N points and spacings $\delta y = \delta x$, and uniform grids $\mathbf{y}_v =$

$\{y_{v,0} + l_v \delta y_v\}_{l_v=0}^{N_v-1}$, $\mathbf{x}_v = \{x_{v,0} + m_v \delta x_v\}_{m_v=0}^{N_v-1}$ with N_v points and spacings $\delta y_v = \delta x_v$. To account for the nonlinearity at $y = \ln(-\lambda_n)$ (see equation (6.12)), grid $\mathbf{y}$ is constructed as explained previously in Section 4.5.1.

Moreover, we determine a uniform, symmetric about zero, grid $\mathbf{u}$ with N points and spacing δu , such that $\mathbf{u}_j = (j - N/2)\delta u$ for $j = 0, \dots, N - 1$, and evaluate $\bar{\varphi}_k$ at $\mathbf{u}$ for each of the points $\{\mathbf{y}_{v,l_v}\}_{l_v=0}^{N_v-1}$, $\{\mathbf{x}_{v,m_v}\}_{m_v=0}^{N_v-1}$. We denote the values of $\bar{\varphi}_k$ on the grid $(\mathbf{u}, \mathbf{y}_v, \mathbf{x}_v)$ by $\bar{\varphi} = \{\bar{\varphi}_{j,l_v,m_v}\}_{N \times N_v \times N_v}$. These values are calculated once and pre-stored for later use in the recursive procedure. To speed up the calculation, the components of $\bar{\varphi}_k$, i.e., Φ_0 , Φ_1 , Φ_2 , Φ_3 , Φ_4 and $\tilde{\mu}_J$ (see equations (6.39), (6.40)) which are independent of the variance level, can be pre-computed separately. The characteristic functions (6.39), (6.40) also involve two modified Bessel functions of the first kind $I_\nu(z)$. To calculate these, we use the MATLAB function `besseli`(ν, z) for real order ν and complex argument z . To fix potential overflow problem due to small argument z , it is advisable to compute first the scaled version $e^{-|\operatorname{Re} z|} I_\nu(z)$ available as `besseli`($\nu, z, 1$) and rescale this to retrieve $I_\nu(z)$.

We pre-store the variance density for all possible transitions $\{\mathbf{x}_{v,m_v}\}_{m_v=0}^{N_v-1}$ to $\{\mathbf{y}_{v,l_v}\}_{l_v=0}^{N_v-1}$: for the square-root variance process, this amounts to calculating the transition density

$$\bar{f}_{k-1,k}^v(\mathbf{y}_v | \mathbf{x}_v) = \chi_{d_f}^2(\Phi_5 \mathbf{y}_v, \lambda_{NC})$$

with auxiliary parameters $\Phi_5 = 4\alpha\eta^{-2}(1 - e^{-\alpha\delta t})^{-1}$, $\lambda_{NC} = \Phi_5 e^{-\alpha\delta t} \mathbf{x}_v$, $d_f = 4\alpha\beta\eta^{-2}$, where $\chi_{d_f}^2(\xi, \lambda_{NC})$ is the noncentral chi-square density of the variable ξ with d_f degrees of freedom and non-centrality parameter λ_{NC} . We denote the function values on the grid $(\mathbf{y}_v, \mathbf{x}_v)$ by $\bar{\mathbf{f}}^v = \{\bar{\mathbf{f}}_{l_v,m_v}^v\}_{N_v \times N_v}$.

Given the above preliminaries, we summarize next in three steps the recursive part of the numerical scheme which is applied n times (equal to the number of monitoring dates), starting from maturity and moving backwards to provide the option price at inception.

1. **Swapping between the state and Fourier spaces.** Assume the values approximating $\bar{p}_k$ at the k^{th} monitoring date are given on the grid $(\mathbf{y}, \mathbf{y}_v)$, and denote these by $\bar{\mathbf{p}}_k =$

$\{\bar{\mathbf{p}}_{l,l_v}\}_{N \times N_v}$. Evaluate $\bar{\mathbf{P}}_k = \{\bar{\mathbf{P}}_{j,l_v}\}_{N \times N_v}$, where

$$\bar{\mathbf{P}}_{\cdot,l_v} = \mathcal{D}(\bar{\mathbf{p}}_{\cdot,l_v}^\top, \mathbf{y}, \mathbf{u}; 1) \delta y$$

computed for each $l_v = 0, \dots, N_v - 1$ are the discrete approximations to the transforms $\mathcal{F}(\bar{p}_k)(u; y_v)$ for each y_v . The discrete Fourier transforms are implemented as standard DFTs using the conversion (2.25). We then approximate the inverse Fourier transform (6.16) by computing $\bar{\mathbf{Q}}_{k-1} = \{\bar{\mathbf{Q}}_{m,l_v,m_v}\}_{N \times N_v \times N_v}$ as

$$\bar{\mathbf{Q}}_{\cdot,l_v,m_v} = \frac{1}{2\pi} \mathcal{D}(\bar{\mathbf{P}}_{\cdot,l_v} \bar{\varphi}_{\cdot,l_v,m_v}, -\mathbf{u}, \mathbf{x}; -1) \delta u$$

for each $l_v = 0, \dots, N_v - 1$ and $m_v = 0, \dots, N_v - 1$ using the conversion (2.26).

2. **From $\bar{\mathbf{Q}}$ to $\bar{q}$.** We calculate $\bar{q}_{k-1}$ in (6.14) by means of the trapezoidal sums $\bar{\mathbf{q}}_{k-1} = \{\bar{\mathbf{q}}_{m,m_v}\}_{N \times N_v}$. For this, we utilize the MATLAB routine TRAPZ

$$\bar{\mathbf{q}}_{m,m_v} = \text{trapz}(\bar{\mathbf{Q}}_{m,\cdot,m_v} \bar{\mathbf{f}}_{\cdot,m_v}^\top) \delta y_v$$

for each $m = 0, \dots, N - 1$ and $m_v = 0, \dots, N_v - 1$.

3. **From $\bar{q}$ to $\bar{p}$.** In calculating $\bar{\mathbf{p}}_{k-1} = \bar{q}_{k-1}(\bar{h}_{k-1}(\mathbf{y}), \mathbf{y}_v)$, we approximate $\bar{q}_{k-1}$ inside grid $\mathbf{x}$ by fitting a cubic interpolating spline to the nodes $(\mathbf{x}, \bar{\mathbf{q}}_{\cdot,m_v})$ for each $m_v = 0, \dots, N_v - 1$ by utilizing the MATLAB built-in function INTERP1. Outside the range of $\mathbf{x}$, we approximate $\bar{q}_{k-1}$ by extrapolating linearly in e^x .

The overall computational complexity of the scheme is $O(nN_v^2N \log_2 N)$, as opposed to $O(nN_v^2N^2)$ for a full implementation with quadrature. It also provides us with the option values on a grid of N_v initial variance values v_0 . For greater flexibility on the construction of the grids and higher accuracies for large grid sizes (see Section 4.6.4), one may replace standard (I)DFTs with fractional DFTs as in Section 4.5.1, at the cost, however, of increased CPU demands.

6.6 Asian option pricing via Monte Carlo simulation

In general, Monte Carlo is slow to compete with grid-based methods in a one-dimensional setup. However, in the Asian case this is not a foregone conclusion because geometric Asian options provide a control variate which has proved powerful in reducing the variance of the Monte Carlo estimators of their arithmetic counterparts (see Kemna and Vorst (1990)). Building on the previous results of Fusai and Meucci (2008) for an exponential Lévy underlying and the insights from the presentation of Lord (2006b), we derive accurate formulae for the prices of discretely sampled geometric Asian options when the assumption of independent and stationary log-increments for the underlying is relaxed. In particular, we consider pricing under the stochastic volatility model assumptions exhibited in Section 6.3, though the results can be extended to other members of the affine class with stochastic interest rates.

6.6.1 Geometric Asian options

The payoff of a geometric Asian option depends on the average underlying value

$$G_n = (S_0^\gamma \prod_{k=1}^n S_k)^{1/(n+\gamma)} = \exp \left(\frac{\gamma \ln S_0 + \sum_{k=1}^n \ln S_k}{n + \gamma} \right), \quad (6.17)$$

where γ takes value 1 (0) when S_0 is (is not) included in the average. Further to (6.17), Fusai and Meucci ((2008), Appendix A) show that

$$\ln G_n = \ln S_0 + \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k \quad (6.18)$$

for log-returns $\{Z_k\}$ from the class of affine models. The price of a geometric Asian with fixed strike K then reads

$$\mathbb{E}(e^{-rT}(\zeta(e^{\ln G_n} - K))^+), \quad (6.19)$$

where ζ takes value 1 (-1) for a call (put) option, whereas the price of a floating-strike Asian with coefficient of partiality $\alpha^* > 0$ reads

$$\mathbb{E}(e^{-rT}(\zeta(S_n - \alpha^* e^{\ln G_n}))^+) = \alpha^* S_0 \bar{\mathbb{E}}((\zeta(1/\alpha^* - e^{\ln \frac{G_n}{S_n}}))^+), \quad (6.20)$$

where the last equality follows from a change to the equivalent spot measure $\bar{\mathbb{P}}$. Knowledge of the distribution laws of $\ln G_n$ and $\ln \frac{G_n}{S_n}$ under measures $\mathbb{P}$ and $\bar{\mathbb{P}}$ respectively is necessary for the computation of (6.19) and (6.20).

Proposition 21 *Consider the log-average asset value $\ln G_n$ given by (6.18). For the Markov system (v, X) whose distribution law obeys to (6.6), we have that the characteristic function of $\ln G_n$ under measure $\mathbb{P}$ is given by*

$$\phi_{\ln G_n}(u; v_0, \ln S_0) = e^{\sum_{k=1}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma}u; \delta t) + \Psi_1(\nu_1, \frac{n}{n+\gamma}u; \delta t)v_0 + iu \ln S_0}, \quad (6.21)$$

$$\nu_n = 0, \quad (6.22)$$

$$\nu_{k-1} = -i\Psi_1\left(\nu_k, \frac{n+1-k}{n+\gamma}u; \delta t\right), \quad 1 < k \leq n. \quad (6.23)$$

Similarly, the characteristic function of $\ln \frac{G_n}{S_n}$ is given under measure $\bar{\mathbb{P}}$ by

$$\bar{\phi}_{\ln \frac{G_n}{S_n}}(u; v_0) = e^{-rn\delta t + \sum_{k=1}^n \Psi_0(\bar{\nu}_k, -\frac{k-1+\gamma}{n+\gamma}u - i; \delta t) + \Psi_1(\bar{\nu}_1, -\frac{\gamma}{n+\gamma}u - i; \delta t)v_0}, \quad (6.24)$$

$$\bar{\nu}_n = 0, \quad (6.25)$$

$$\bar{\nu}_{k-1} = -i\Psi_1\left(\bar{\nu}_k, -\frac{k-1+\gamma}{n+\gamma}u - i; \delta t\right), \quad 1 < k \leq n. \quad (6.26)$$

Proof. By induction on j we will prove that

$$\mathbb{E}\left(e^{iu \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k} \middle| \mathcal{F}_{j-1}\right) = e^{\sum_{k=j}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma}u; \delta t) + \Psi_1(\nu_j, \frac{n+1-j}{n+\gamma}u; \delta t)v_{j-1} + iu \sum_{k=1}^{j-1} \frac{n+1-k}{n+\gamma} Z_k 1_{\{j-1>0\}}}, \quad (6.27)$$

for $j = 1, \dots, n$. The statement holds for $j = n$ by $\mathcal{F}_{n-1}$ -measurability of the partial sum $\sum_{k=1}^{n-1} \frac{n+1-k}{n+\gamma} Z_k$ and expressions (6.6) and (6.22). Assume therefore that

$$\mathbb{E}\left(e^{iu \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k} \middle| \mathcal{F}_j\right) = e^{\sum_{k=j+1}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma}u; \delta t) + \Psi_1(\nu_{j+1}, \frac{n-j}{n+\gamma}u; \delta t)v_j + iu \sum_{k=1}^j \frac{n+1-k}{n+\gamma} Z_k 1_{\{j>0\}}}$$

holds for arbitrary $1 < j < n - 1$. By the law of iterated expectations

$$\begin{aligned}
 \mathbb{E} \left(e^{iu \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k} \middle| \mathcal{F}_{j-1} \right) &= \mathbb{E} \left(\mathbb{E} \left(e^{iu \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k} \middle| \mathcal{F}_j \right) \middle| \mathcal{F}_{j-1} \right) \\
 &= \mathbb{E} \left(e^{\sum_{k=j+1}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma} u; \delta t) + \Psi_1(\nu_{j+1}, \frac{n-j}{n+\gamma} u; \delta t) v_j} \right. \\
 &\quad \times e^{iu \sum_{k=1}^j \frac{n+1-k}{n+\gamma} Z_k} \middle| \mathcal{F}_{j-1} \Big) \\
 &\stackrel{\mathcal{F}_{j-1}\text{-meas.}}{=} e^{\sum_{k=j+1}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma} u; \delta t) + iu \sum_{k=1}^{j-1} \frac{n+1-k}{n+\gamma} Z_k} \\
 &\quad \times \mathbb{E} \left(e^{\Psi_1(\nu_{j+1}, \frac{n-j}{n+\gamma} u; \delta t) v_j + iu \frac{n+1-j}{n+\gamma} Z_j} \middle| \mathcal{F}_{j-1} \right) \\
 &= e^{\sum_{k=j}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma} u; \delta t) + \Psi_1(\nu_j, \frac{n+1-j}{n+\gamma} u; \delta t) v_{j-1} + iu \sum_{k=1}^{j-1} \frac{n+1-k}{n+\gamma} Z_k},
 \end{aligned}$$

where the last equality follows from (6.6) and (6.23). By induction, we deduce from (6.27) that for $j = 1$

$$\mathbb{E}(e^{iu \sum_{k=1}^n \frac{n+1-k}{n+\gamma} Z_k}) = e^{\sum_{k=1}^n \Psi_0(\nu_k, \frac{n+1-k}{n+\gamma} u; \delta t) + \Psi_1(\nu_1, \frac{n}{n+\gamma} u; \delta t) v_0},$$

which completes the proof of (6.21) by virtue of (6.18).

Next, for the derivation of the characteristic function of $\ln \frac{G_n}{S_n}$ under measure $\bar{\mathbb{P}}$, we recognize from Propositions 22-23 and result (6.6) that

$$\bar{\mathbb{E}}(e^{iu_1 v_k + iu_2 Z_k} | \mathcal{F}_{k-1}) = \mathbb{E}(e^{iu_1 v_k - r\delta t + i(u_2 - i)Z_k} | \mathcal{F}_{k-1}) = e^{-r\delta t + \Psi_0(u_1, u_2 - i; \delta t) + \Psi_1(u_1, u_2 - i; \delta t) v_{k-1}}. \tag{6.28}$$

Substituting (6.18) for $\ln G_n$ and $\ln S_n = \ln S_0 + \sum_{k=1}^n Z_k$ into $\ln G_n - \ln S_n$ yields

$$\bar{\mathbb{E}}(e^{iu \ln \frac{G_n}{S_n}}) = \bar{\mathbb{E}}(e^{-iu \sum_{k=1}^n \frac{k-1+\gamma}{n+\gamma} Z_k}).$$

Using (6.28) in addition to the arguments for the proof of (6.21), one can show that

$$\bar{\mathbb{E}}(e^{-iu \sum_{k=1}^n \frac{k-1+\gamma}{n+\gamma} Z_k}) = e^{-rn\delta t + \sum_{k=1}^n \Psi_0(\bar{\nu}_k, -\frac{k-1+\gamma}{n+\gamma} u - i; \delta t) + \Psi_1(\bar{\nu}_1, -\frac{\gamma}{n+\gamma} u - i; \delta t) v_0},$$

where the sequence $\bar{\nu}_k$ obeys to (6.25-6.26). This completes the proof. ■

Given the characteristic functions (6.21) and (6.24), the option prices (6.19) and (6.20) can be provided analytically in terms of Fang and Oosterlee's Fourier-cosine series expansions for

European-type contracts (see equation (3.13)).³

Note that, although results (6.21) and (6.24) have been derived on the assumption of a bivariate affine setup, an extension to multivariate structures is straightforward. This is the case, for example, with the trivariate affine Gaussian-OU stochastic volatility model in Stein and Stein (1991) and Schöbel and Zhu (1999), but also the jump diffusion with Vašíček interest rates hybrid setup illustrated later in Chapter 8, and the extended Schöbel-Zhu model to include Hull and White stochastic interest rates in van Haastrecht et al. (2009a).⁴

For the purposes of our study, we focus on the Heston and Bates models. Before moving to option pricing via Monte Carlo, it is essential that we do a critical review of the available schemes for the simulation of the Heston model (the Bates model follows as a straightforward extension), and opt for the one that fits best to the pricing of an Asian option.

6.6.2 Simulation of the Heston model

So far, standard approach to pricing exotic contracts, e.g., path-dependent contracts including barrier, Asian and lookback options, is based on Monte Carlo simulation. Nevertheless, until recently, difficulties with the exact simulation of the Heston model have been encountered due to the lack of explicit knowledge about the distribution law of the time integral of the variance, $\int v_s ds$. For this, Euler discretization has been traditionally employed subject, however, to a bias that has to be estimated, and a time grid which is usually much finer, than is strictly necessary for the contract in question, to minimize the bias. For the Heston model

$$dX_t = (r - v_t/2)dt + \sqrt{v_t}(\rho dW_t + \sqrt{1 - \rho^2}dB_t), \quad (6.29)$$

$$dv_t = \alpha(\beta - v_t)dt + \eta\sqrt{v_t}dW_t \quad (6.30)$$

³Heuristically one can easily show that the price of the discrete geometric Asian option converges at $O(n)$ to the price of its continuously sampled counterpart. This can then serve as “biased” control variate in simulating a continuous arithmetic Asian in the spirit of Fu et al. (1999), who reach that the bias introduced by the use of the continuous geometric Asian control variate offsets the inherent bias due to sampling on a finite set of dates.

⁴For the efficient simulation of the original Schöbel-Zhu model and its extension to include stochastic interest rates, see the recent work of van Haastrecht et al. (2009b).

with independent Brownian motions W, B , the basic Euler scheme takes the form

$$\hat{X}_{t+\delta t} = \hat{X}_t + (r - \hat{v}_t/2)\delta t + \sqrt{\hat{v}_t}(\rho(W_{t+\delta t} - W_t) + \sqrt{1 - \rho^2}(B_{t+\delta t} - B_t)), \quad (6.31)$$

$$\hat{v}_{t+\delta t} = \hat{v}_t + \alpha(\beta - \hat{v}_t)\delta t + \eta\sqrt{\hat{v}_t}(W_{t+\delta t} - W_t), \quad (6.32)$$

where $\hat{X}, \hat{v}$ denote the approximate (by discretization of the SDEs (6.29-6.30)) realizations of X, v . An obvious limitation of the scheme (6.31-6.32) is that the discretized process $\hat{v}$ may become negative with nonzero probability, regardless of the time step size, rendering the computation of $\sqrt{\hat{v}}$ impossible. Various “fixes” to this behaviour have been proposed in the literature; Lord et al. (2010) summarize these “fixes” in the general framework

$$\begin{aligned} \hat{X}_{t+\delta t} &= \hat{X}_t + (r - \xi_2(\hat{v}_t)/2)\delta t + \sqrt{\xi_2(\hat{v}_t)}(\rho(W_{t+\delta t} - W_t) + \sqrt{1 - \rho^2}(B_{t+\delta t} - B_t)), \\ \hat{v}_{t+\delta t} &= \xi_0(\hat{v}_t) + \alpha(\beta - \xi_1(\hat{v}_t))\delta t + \eta\sqrt{\xi_2(\hat{v}_t)}(W_{t+\delta t} - W_t). \end{aligned} \quad (6.33)$$

In particular, Higham and Mao (2005) opt for $\xi_0(x) = \xi_1(x) = x$, $\xi_2(x) = |x|$, allowing therefore for negative variance samples which help to keep the bias low. However, at the same time reflecting large negative variance values at the origin (via ξ_2), causes larger than intended moves in the asset price process. Deelstra and Delbaen (1998) fix the effect from variance reflection by absorbing the variance instead, using $\xi_2(x) = x^+$. This is a “partial truncation” scheme in the sense that only the diffusion coefficient of (6.33) is truncated at zero. With a view to lowering the bias further, Lord et al. (2010) introduce the “full truncation” scheme where the drift of (6.33) is truncated as well, by setting $\xi_1(x) = x^+$. With this modification, they manage to keep the variance samples negative for longer periods of time, effectively lowering the volatility of the underlying which in turn helps in reducing the bias. The “full truncation” scheme has shown to minimize the positive bias amongst the other “fixes” when pricing European and path-dependent options⁵. While the previous schemes aim at controlling the bias from the discretization of the Heston model with the use of appropriate “fixes”, Andersen (2008) proposes a “quadratic-exponential” approximation scheme to simulate the transition of the variance to the next time step given its current position, skipping the use

⁵ Absorption and reflection schemes with “fixes” $\xi_0(x) = \xi_1(x) = \xi_2(x) = x^+$ and $\xi_0(x) = \xi_1(x) = \xi_2(x) = |x|$ respectively, are susceptible to higher bias (see Lord et al. (2010)).

of any “fixes”. In fact, Andersen (2008) demonstrates that his method works efficiently with a smaller number of time steps compared to the “full truncation” technique of Lord et al. (2010). The “quadratic-exponential” technique will be revisited in Section 6.7.3.

Broadie and Kaya (2006) are the first to generate exact sample trajectories for the Heston model without inducing discretization bias. Exact simulation methods preclude the need to use an unnecessarily dense time grid. Key to their method is the sampling from the distribution of the time integral of the variance process conditional on its endpoint values. This distribution law is only known through its characteristic function which they derive and numerically invert to generate samples. In principle, the transform inversion is the most time-consuming part of their method since the characteristic function depends on values of the variance process; thus, transform inversion needs to be repeated for each time step and simulation run. In addressing this issue, Glasserman and Kim (2008a) show that the law of the conditional integrated variance can be analyzed into three different special classes of infinitely divisible distributions. A route towards evaluating these distributions is via inversion of their characteristic functions, two of which are independent of the endpoint variance values so that the associated distributions can be pre-tabulated for fast simulation. Another way is to draw samples from these distributions via truncated gamma series expansions which Glasserman and Kim (2008a) have also developed. Although faster, the second method generates certain amount of bias as a consequence of the series truncation. Finally, summing the three samples obtained either way yields a sample from the conditional integrated variance distribution. Glasserman and Kim (2008a) illustrate that their method reduces substantially the computational burden of Broadie and Kaya (2006) when pricing path-independent plain vanilla options with a single time step to maturity. Building further on their original work, Glasserman and Kim (2008b) exploit the infinite divisibility with respect to $\tilde{v} \equiv v_t + v_{t+\delta t}$ of the third-type distribution, to use bridge sampling and produce a beta approximation for this. With this in hand, Glasserman and Kim (2008b) are able to pre-store all the distributions composing the distribution of the conditional integrated variance.

In addition to the Heston model, in order to simulate the Bates model, one additionally needs to generate independent Poisson jump times within the time horizon and jump sizes, and add these to the log-asset diffusion.

6.7 Numerical study

For the purposes of this study, we opt for an Asian put option with floating strike which matures at a year's time and is subject to monthly observation. The parameters relevant to the contract are $\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$ (see Table 4.1), $T = 1$, $n = 12$, interest rate $r = 0.04$ and $S_0 = 100$. All numerical experiments are coded in MATLAB R2008b on a Dell Optiplex 755 Intel Core 2 Duo PC 2.66 GHz with 2.0 GB RAM.

6.7.1 Models

We evaluate option prices numerically based on three distributions of log-returns from the Lévy class: Gaussian, normal inverse Gaussian and tempered stable, and two distributions from the non-Lévy class with stochastic volatility: Heston and Bates.

We use the three parameter sets in Table 6.1: model parameterizations I (Heston) and III (Bates) are adopted from Duffie et al. (2000) and represent fitted parameters for market option prices for the S&P 500 index on a particular date. Parameter set II (Heston) is taken from Andersen (2008) and is relevant for equity option markets. All three sets have been utilized previously by Broadie and Kaya (2006), Andersen (2008) and Glasserman and Kim (2008a) in the simulation of European vanilla options. Furthermore, we calibrate the three Lévy models for a year's time horizon to achieve volatility $\text{vol} \in \{0.1364, 0.1336\}$, and for the non-Gaussian distributions, skewness coefficient $s \in \{-1.326, -1.235\}$ and excess kurtosis $\kappa \in \{3.483, 2.681\}$ in consistency with the model specifications I and III.⁶ The fitted Lévy parameters are summarized in Table 6.2.

6.7.2 Pricing via convolution

The convolution algorithm we propose in Section 6.5 is free of any approximation bias either due to discretization of the continuous-time variance process (see Monte Carlo methods) or insufficiently large mean-reversion rate (see PDE methods and control variate Monte Carlo strategies in Section 6.2). It is also relevant to pricing Asian options with discrete observation.

⁶Exact cumulants are derived for all models by differentiating the corresponding cumulant generating functions and evaluating at zero, as indicated by equation (2.5).

	set I (Heston)	set II (Heston)	set III (Bates)
α	6.21	1.0	3.99
β	0.019	0.09	0.014
η	0.61	1.0	0.27
ρ	-0.7	-0.3	-0.79
λ_J	n.a.	n.a.	0.11
μ_J	n.a.	n.a.	-0.1391
σ_J	n.a.	n.a.	0.15
v_0	0.010201	0.09	0.008836

Table 6.1: Stochastic volatility model parameters

set	Gaussian	normal inverse Gaussian			tempered stable			
	σ	ν	σ	θ	C	G	M	Y
I	0.1364	0.3786	0.0949	-0.1594	0.1697	6.187	60.86	0.8
III	0.1336	0.2155	0.0617	-0.2552	0.3802	8.291	127.12	0.6

Table 6.2: Lévy model parameters fitted to the models I & III (Table 6.1)

In Table 6.3 we present our results for the Asian option defined above. For fixed $\mathbf{x}_v$ -grid, we observe that the results become more accurate with higher number of $\mathbf{x}$ -grid points N , while the CPU time grows almost linearly. Soon we manage to restore monotone second-order error convergence, which permits a user to gauge the precision of the scheme. This agrees with our conclusion for the original scheme built on the assumption of Lévy log-returns in Chapter 4. High-level accuracy requires that we utilize too many grid points to manage the computations in reasonable time. For this, we exploit the regular quadratic convergence in N to produce prices at high precision using Richardson extrapolation.

Comparing the option prices obtained under Gaussian and non-Gaussian log-returns, we find that the Gaussian option prices are higher as illustrated in Table 6.4. This behaviour stems from the combination of the negative skewness and excess kurtosis effects existing in the non-Gaussian risk-neutral distributions, as opposed to the Gaussian distribution. Also, as deduced in Chapter 4, the prices from the two non-Gaussian Lévy models agree to the penny, suggesting that the skewness and excess kurtosis effects, rather than the Lévy model per se, determine the option price levels. After introducing dependence in the log-returns, the proportional relationship of the skewness coefficient to $t^{-1/2}$ and excess kurtosis to t^{-1} (t : time)

N	set I (Heston)		set II (Heston)		set III (Bates)		CPU (s)
	error	RE error	error	RE error	error	RE error	
2^8	-4.8E-1	n.a.	-3.0E-2	n.a.	-4.8E-1	n.a.	126.7
2^9	2.8E-2	n.a.	-7.4E-3	2.7E-4	-1.0E-2	n.a.	246.9
2^{10}	-1.9E-3	n.a.	-1.8E-3	8.2E-5	-1.9E-3	9.0E-4	469.0
2^{11}	-4.8E-4	-1.3E-5	-4.4E-4	6.2E-6	-4.8E-4	2.1E-6	1029.5

Table 6.3: Floating-strike Asian put option ($\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$, $T = 1$, $n = 12$): precision of the convolution method with increasing $\mathbf{x}$ -grid points N . Fixed $\mathbf{x}_v$ -grid points $N_v = 2^8$. Model parameters: Table 6.1. Other parameters: $r = 0.04$, $S_0 = 100$. “Error” (or “RE error”) computed as the difference between the prices for given N (or Richardson extrapolation prices) and the reference values. Empty cells under the “RE error” headings imply non-smooth convergence. Reference values obtained with the convolution method (precision $\pm 10^{-7}$): 2.0706445 (I), 4.0775434 (II), 2.03099732 (III). CPU times in seconds (s).

set	Heston	Bates	Gaussian	NIG	tempered stable
I	2.07064		2.15309	1.98414	1.99588
III		2.03099	2.09288	1.98772	1.98953

Table 6.4: Lévy versus non-Lévy with stochastic volatility log-returns: comparison of floating-strike Asian put option prices across models ($\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$, $T = 1$, $n = 12$). Model parameters from Tables 6.1 & 6.2. Other parameters: $r = 0.04$, $S_0 = 100$. Lévy prices computed using the convolution algorithm in Chapter 4. Stochastic volatility prices computed using the extended algorithm in Section 6.5.

for a Lévy model is no longer valid, but instead different types of time-dependence apply for different log-return models with stochastic volatility; this effect causes slight departures from the Lévy prices.

6.7.3 Monte Carlo pricing

We test the performance of the convolution method by comparing with the outcome from the QE method (with $\tilde{n}$ time steps) of Andersen (2008).⁷ A thorough description of this can be found in Appendix 6.A. This scheme has shown to produce the least bias compared to the alternative biased methods delineated in Section 6.6.2, and be a nontrivial competitor to the exact schemes of Broadie and Kaya (2006) and Glasserman and Kim (2008a) in pricing path-independent European vanilla contracts. For this type of contract, the gamma expansion

⁷As we discuss next, due to time discretization of the variance process in the QE scheme, it may be necessary to use a larger number of time steps $\tilde{n}$ than the actual monitoring points n of the option. In this case, we employ $\tilde{n}$ as integer multiple of n .

$\tilde{n}$	set I		set II		set III	
	bias	std error	bias	std error	bias	std error
12	0.0117	0.00002	0.0062	0.00011	0.0015	0.000021
24	0.0027	0.00002	0.0023	0.00011	0.0014	0.000021
48	0.0011	0.00002	0.0009	0.00011	0.0006	0.000021

Table 6.5: Floating-strike Asian put option ($\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$, $T = 1$, $n = 12$): estimated biases (with standard errors) of the QE method for number of time steps $\tilde{n}$ for each parameter set in Table 6.1. Other parameters: $r = 0.04$, $S_0 = 100$. True option prices: 2.0706445 (I), 4.0775434 (II), 2.03099732 (III).

method with a single time step ($n = 1$) outperforms the QE method with, e.g., $\tilde{n} = 48$, time steps by a factor of 3–4.5. This however suggests that generating $n = 12$ intermediate samples for the underlying when pricing the Asian option would take at least 3 times as long using the gamma expansion. From a practical point of view, we therefore conclude that QE is the most suitable and fairest speed-accuracy competitor to the convolution method. In the attempt to ameliorate the convergence of the simulation error, we further apply the control variate technique where as control for the price of the arithmetic Asian option we employ its geometric counterpart with price computed as explained in Section 6.6.1.

In our implementation of the QE approximation, we set the parameters $\epsilon_1 = \epsilon_2 = 0.5$ and $\psi_c = 1.5$ (see Appendix 6.A), following Andersen (2008). The theoretical convergence rate of the QE method is not known, still we can estimate the order of the bias due to time discretization, $O(\tilde{n}^{-\varepsilon})$. If π_0 is the true option price, we define the bias of the Monte Carlo price estimate $\hat{\pi}_0$ generated by the approximate QE method as $\text{bias} = \mathbb{E}(\hat{\pi}_0 - \pi_0)$, and its variance $\mathbb{E}((\hat{\pi}_0 - \mathbb{E}(\hat{\pi}_0))^2)$. The root mean square error is then defined as $\text{RMSE} = (\text{bias}^2 + \text{variance})^{1/2}$. To obtain an accurate estimate for the bias we utilize 50 million sample trajectories for each $\tilde{n} \in \{12, 24, 48\}$, together with the geometric Asian control variate to speed up convergence⁸. The true price required for the estimation of the bias is deduced at high precision via the convolution method of Section 6.5. The results are summarized in Table 6.5. The ordinary least-squares estimates ε are then 1.7, 1.37 and 0.66 for the parameter sets I, II and III respectively. Furthermore, Duffie and Glynn (1995) show that, for first-order time-discretization schemes,

⁸Note that variance reduction techniques reduce the standard error of the Monte Carlo estimates, but do not fix the discretization bias.

i.e., $O(\tilde{n}^{-1})$, it is optimal to vary $\tilde{n}$ proportional to the square root of the number of simulation trials m . Then, the RMSE of the Monte Carlo estimator converges at $O(w^{-\frac{\varepsilon}{1+2\varepsilon}})$, where the total workload w encompasses both $\tilde{n}$ and m . Following this rule, the ε estimates for the sets I-III yield estimates for the optimal convergence rates 0.38, 0.37 and 0.29 respectively⁹.

Figure 6-1 plots the RMSE versus CPU time (on a log-log scale) tradeoff of the QE scheme with $\tilde{n} \in \{12, 24, 48\}$ in pricing an Asian put option with $n = 12$ for the three model specifications in Table 6.1. Overall, for low accuracies, the crude (without control variate) QE with $\tilde{n} = 12$ produces faster results than the QE with $\tilde{n} \in \{24, 48\}$. However, as the number of simulation trials (also the CPU time) increases, the bias eventually dominates the RMSE slowing down its rate of decrease. Raising $\tilde{n}$ reduces the bias and increases the accuracy range, but also the CPU effort. The QE method with $\tilde{n} = 48$ achieves a steeper constant slope close to the optimum -0.5 of an unbiased scheme, raising the potential for high-level precision; this is more obvious for models II and III which induce smaller biases than I. Additionally, supplying the arithmetic Asian option simulation with the geometric Asian control variate leads to nontrivial standard error reductions (up to 30 times), speeding up the decay of the RMSE to the bias level for each $\tilde{n}$. The simulation results are summarized in Table 6.6.

We compare the results reported in Tables 6.3 and 6.6 in terms of accuracy and CPU time demands. Over all the three parameter sets, the control variate QE (CVQE) strategy with $\tilde{n} = 48$ achieves highest RMSE 7.8E-3 (set II) within 5 seconds of CPU time, which is less than the highest absolute error reported for the convolution method. Nevertheless, the convolution method converges at a higher rate with the potential of significant improvement in precision, in contrast to the CVQE whose rate of convergence is damped by the existence of bias. In particular, for cases I and II the CVQE reaches RMSE level 1E-3 in excess of 800 seconds, while the convolution scheme achieves absolute error of the same magnitude in not more than 700 seconds for case I and 250 seconds for case II. In case III, the CVQE requires more than 1000 seconds to attain RMSE 6E-4, whereas the convolution scheme slightly more than 1000 seconds for absolute error 2E-6.

⁹Note that a standard Euler scheme converges at $O(w^{-\frac{1}{3}})$ (for smooth payoffs and Lipschitzian asset dynamics – not the case for the Heston model), whereas an unbiased crude (without variance reduction) scheme achieves $O(w^{-\frac{1}{2}})$.

set	$m \times 10^3$	$\tilde{n}$	crude QE std error	CPU (s)	crude QE RMSE	$\tilde{n}$	CVQE std error	CPU (s)	CVQE RMSE
I	10	12	0.04346	1.85	0.04501	48	0.00144	3.42	0.00181
	40	12	0.02197	5.03	0.02490	48	0.00083	12.94	0.00138
	160	24	0.01104	29.67	0.01137	48	0.00040	52.28	0.00117
	640	24	0.00552	118.73	0.00615	48	0.00020	207.45	0.00112
	2,560	48	0.00276	834.89	0.00297	48	0.000099	834.89	0.00110
	10,240	48	0.00138	3315.89	0.00176	48	0.000049	3315.89	0.00110
II	10	12	0.08742	1.70	0.08762	48	0.00773	3.55	0.00778
	40	12	0.04283	4.83	0.04325	48	0.00403	13.94	0.00413
	160	12	0.02120	17.28	0.02203	48	0.00199	55.52	0.00218
	640	24	0.01062	116.13	0.01080	48	0.00099	221.92	0.00134
	2,560	48	0.00531	887.67	0.00538	48	0.00049	887.67	0.00103
	10,240	48	0.00265	3293.65	0.00280	48	0.00025	3293.65	0.00093
III	10	12	0.04263	2.27	0.04265	48	0.00165	4.80	0.00176
	40	12	0.02134	6.65	0.02139	48	0.00078	18.91	0.00098
	160	12	0.01072	24.20	0.01083	48	0.00038	75.43	0.00071
	640	24	0.00539	144.67	0.00548	48	0.00019	301.20	0.00063
	2,560	48	0.00268	1205.09	0.00275	48	0.000094	1205.09	0.00061
	10,240	48	0.00134	4779.58	0.00147	48	0.000047	4779.58	0.00060

Table 6.6: Floating-strike Asian put option ($\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$, $T = 1$, $n = 12$): QE simulation results for parameter sets I-III. Columns 3-6 (columns 7-10): crude QE (CVQE: QE with geometric Asian control variate) output. Geometric Asian ref. prices: 1.9929986 (I), 3.6981218 (II), 1.9569746 (III). RMSEs based on biases for each set of parameters and number of time steps $\tilde{n}$ in Table 6.5 and the standard errors corresponding to m simulation trials. CPU times in seconds (s).




Figure 6-1: Floating-strike Asian put option ($\zeta = -1$, $\gamma = 1$, $\alpha^* = 1$, $T = 1$, $n = 12$): accuracy versus speed comparisons for the crude QE method and the QE method equipped with geometric Asian control variate for $\hat{n} \in \{12, 24, 48\}$.

6.8 Concluding remarks

In this chapter, we have presented the backward convolution method for pricing discretely sampled arithmetic Asian options under the Heston and Bates stochastic volatility models. The same approach can be applied to other stochastic volatility models within the bivariate affine class, providing that the variance transition law and the variance-conditional log-return distribution (see Appendix 6.C) are known.

Furthermore, building on the original Fusai and Meucci (2008) for Lévy log-returns, we extend their approach to derive exact prices for geometric Asians under non-Lévy log-returns by means of Fourier transforms, for use as control variates in the simulation of arithmetic Asians. Utilizing the geometric Asian as a control for the price of the arithmetic Asian option in the stochastic volatility market setup, results in substantial variance reduction.

Numerical examples run using both the convolution and control variate Monte Carlo schemes illustrate that, by regular second-order convergence in the number of grid points, the convolution method performs better at high levels of precision. Control variate Monte Carlo is faster for smaller accuracies, still it requires that we pre-estimate the inherent bias of common biased simulation schemes in the literature; a usually time-consuming procedure. For measuring the bias, it is necessary to know in advance the option price with sufficient accuracy which is possible by virtue of the convolution method¹⁰. Moreover, implementing the convolution algorithm by utilizing FFT routines provides us with the option prices on a grid of initial variances in a single run. In the spirit of Chapter 5, an extension to obtain the price sensitivities is straightforward.

¹⁰Using instead unbiased simulation schemes (see Section 6.6.2) to price a path-dependent option, e.g., an Asian option, would impact substantially the computing time.

Appendix 6.A—The truncated Gaussian and quadratic-exponential schemes

It is known that the true $v_{t+\delta t}$ is proportional to a noncentral chi-square distribution with $d_f = 4\alpha\beta\eta^{-2}$ degrees of freedom and non-centrality parameter $\lambda_{NC} = v_t \frac{4\alpha e^{-\alpha\delta t}}{\eta^2(1-e^{-\alpha\delta t})}$ (e.g., Cox et al. (1985)). Also, the noncentral chi-square distribution converges asymptotically to a normal distribution as λ_{NC} increases to infinity, while this becomes asymptotically proportional to a (central) chi-square distribution with d_f degrees of freedom as λ_{NC} approaches zero. Thus, for sufficiently large v_t a Gaussian random variable serves as a reasonable proxy for $v_{t+\delta t}$. On the other hand, for small v_t , it is necessary to provide an approximation for the distribution of $v_{t+\delta t}$ that mimics the features of the chi-square density; that is, concentration of the probability mass at the origin combined with an upper density tail. To adapt to these properties for sufficiently small/large v_t , Andersen (2008) provides the so-called “truncated Gaussian” (TG) approximation

$$\hat{v}_{t+\delta t} = (\mu_{TG} + \sigma_{TG}Z^v)^+,$$

where $Z^v \sim \mathcal{N}(0, 1)$ and the constants μ_{TG} , σ_{TG} are determined by matching the first two moments of $\hat{v}_{t+\delta t}$ to the true moments of $v_{t+\delta t}$ conditional on $v_t = \hat{v}_t$ (see Andersen (2008), Proposition 4). It is observed however that the tail of the TG distribution goes to zero much faster than the chi-square distribution implies.

To address the tail issue related to the TG random variable, Andersen (2008) alternatively introduces a “quadratic-exponential” (QE) approximation: the noncentral chi-square distribution with moderate/high λ_{NC} is represented by a power function applied to a normal random variable, e.g.,

$$\hat{v}_{t+\delta t} = (a_{QE} + b_{QE}Z^v)^2, \tag{6.34}$$

where the constants a_{QE} , b_{QE} are determined by moment matching as with the TG scheme. In fact, from Andersen ((2008), Proposition 5) for $m = \mathbb{E}(v_{t+\delta t} | v_t = \hat{v}_t)$, $s^2 = \text{Var}(v_{t+\delta t} | v_t = \hat{v}_t)$

and $\psi = s^2 m^{-2} \leq 2$

$$\begin{aligned} b_{QE}^2 &= 2\psi^{-1} - 1 + \sqrt{2\psi^{-1}} \sqrt{2\psi^{-1} - 1} \geq 0, \\ a_{QE} &= \frac{m}{1 + b_{QE}^2}, \end{aligned}$$

such that $\hat{v}_{t+\delta t}$ in (6.34) satisfies $\mathbb{E}(\hat{v}_{t+\delta t}) = m$ and $\mathbb{V}\text{ar}(\hat{v}_{t+\delta t}) = s^2$. On the other hand, for small λ_{NC} the density of $\hat{v}_{t+\delta t}$ is approximated explicitly by

$$p_{QE} \delta(0) + \beta_{QE} (1 - p_{QE} e^{-\beta_{QE} \hat{v}_{t+\delta t}}), \quad (6.35)$$

where δ is the Dirac delta, and p_{QE} , β_{QE} are nonnegative constants to be determined. Density (6.35) permits one to specify explicitly the probability mass p_{QE} at the origin, as opposed to the TG case. Furthermore, the mass at the origin is supplemented with an exponential tail which reflects the chi-square distribution. From (6.35) we derive the cumulative distribution function and subsequently invert this to yield

$$\hat{v}_{t+\delta t} = \begin{cases} 0, & 0 \leq u \leq p_{QE} \\ \beta_{QE}^{-1} \ln(\frac{1-p_{QE}}{1-u}), & p_{QE} < u \leq 1. \end{cases} \quad (6.36)$$

We use (6.36) to produce samples from the distribution of $\hat{v}_{t+\delta t}$. Regarding the choice on p_{QE} and β_{QE} , we have from Andersen ((2008), Proposition 6) for $\psi \geq 1$ that

$$\begin{aligned} p_{QE} &= \frac{\psi - 1}{\psi + 1} \in [0, 1], \\ \beta_{QE} &= \frac{2}{m(\psi + 1)} \geq 0, \end{aligned}$$

such that $\hat{v}_{t+\delta t}$ in (6.36) satisfies $\mathbb{E}(\hat{v}_{t+\delta t}) = m$ and $\mathbb{V}\text{ar}(\hat{v}_{t+\delta t}) = s^2$.

Note that the “quadratic” part (6.34) can be moment-matched for $\psi \leq 2$, whilst the “exponential” part (6.36) for $\psi \geq 1$. To deal with the applicability overlap for $\psi \in [1, 2]$, Andersen (2008) suggests $\psi_c = 1.5$ as the switching rule between the two sampling schemes, though a different choice $\psi_c \in [1, 2]$ appears to have a minor effect on the performance of the algorithm during simulation.

Finally, combining (6.30) to (6.29) and integrating over $[t, t + \delta t]$ yields

$$\begin{aligned} X_{t+\delta t} &= X_t + (r - \rho\alpha\beta\eta^{-1})\delta t + (\rho\alpha\eta^{-1} - 1/2) \int_t^{t+\delta t} v_s ds + \rho\eta^{-1}(v_{t+\delta t} - v_t) \\ &\quad + \sqrt{1 - \rho^2} \int_t^{t+\delta t} \sqrt{v_s} dB_s, \end{aligned} \quad (6.37)$$

which, by approximation of the time integral of the variance using a linear combination of its realizations at t and $t + \delta t$, i.e.,

$$\left(\int_t^{t+\delta t} v_s ds \middle| v_t = \hat{v}_t, v_{t+\delta t} = \hat{v}_{t+\delta t} \right) \stackrel{d}{\approx} (\epsilon_1 \hat{v}_t + \epsilon_2 \hat{v}_{t+\delta t}) \delta t,$$

generates the following approximate realization for the log-underlying at time $t + \delta t$

$$\begin{aligned} \hat{X}_{t+\delta t} &= \hat{X}_t + (r - \rho\alpha\beta\eta^{-1})\delta t + (\epsilon_1 \delta t (\rho\alpha\eta^{-1} - 1/2) - \rho\eta^{-1}) \hat{v}_t \\ &\quad + (\rho\eta^{-1} + \epsilon_2 \delta t (\rho\alpha\eta^{-1} - 1/2)) \hat{v}_{t+\delta t} + \sqrt{\epsilon_1 \delta t (1 - \rho^2) \hat{v}_t + \epsilon_2 \delta t (1 - \rho^2) \hat{v}_{t+\delta t}} Z^X, \end{aligned}$$

where $Z^X \sim \mathcal{N}(0, 1)$ is independent of $\hat{v}$. Also, the use of a central discretization for the integrated variance, i.e., $\epsilon_1 = \epsilon_2 = 0.5$, in the numerical tests of Andersen (2008) suffices to induce a fairly small amount of bias for a small number of time steps, in other words, relatively large step size δt .

Appendix 6.B—Spot measure change for Lévy and time-changed Lévy processes

Recall from Section 3.2 that the characteristic function of a Lévy process L_t with characteristics (μ, σ^2, Π) under some probability measure P has form

$$E(e^{iuL_t}) = e^{\psi_L(u)t}, \quad (6.38)$$

where the Lévy exponent ψ_L admits representation

$$\psi_L(u) = iu\mu - \frac{1}{2}\sigma^2 u^2 + \int_{\mathbb{R} \setminus \{0\}} (e^{iul} - 1 - iul1_{\{|l| \leq 1\}}) \Pi(dl)$$

with drift parameter $\mu \in \mathbb{R}$, diffusion parameter $\sigma > 0$ and Lévy density Π on $\mathbb{R}$ satisfying standard hypotheses.

Proposition 22 *Assume asset price process $S_t = S_0 e^{rt+L_t}$ with $\mu = -\frac{1}{2}\sigma^2 - \int_{\mathbb{R} \setminus \{0\}} (e^l - 1 - l1_{\{|l| \leq 1\}}) \Pi(dl)$ such that $\psi_L(-i) = 0$, hence $\mathbb{E}(S_t) = S_0 e^{rt}$ under the risk-neutral measure $\mathbb{P}$. Then, there exists an equivalent probability measure $\bar{\mathbb{P}}$ defined by its Radon-Nikodým derivative with respect to $\mathbb{P}$,*

$$\bar{\gamma}_t := \left. \frac{d\bar{\mathbb{P}}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = e^{L_t}.$$

The characteristic function of $Z_t = \ln \frac{S_t}{S_0}$ under $\bar{\mathbb{P}}$ is given as

$$\bar{\mathbb{E}}(e^{iuZ_t}) = e^{-rt} \mathbb{E}(e^{i(u-i)Z_t}).$$

Proof. From Geman et al. ((1995), Definition 2), the asset price S_t satisfies the definition of a numéraire. Following Geman et al. ((1995), Theorem 1), we construct the Radon-Nikodým derivative of the equivalent measure $\bar{\mathbb{P}}$ with respect to $\mathbb{P}$ as

$$\bar{\gamma}_t = \left. \frac{d\bar{\mathbb{P}}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \frac{S_t}{S_0 e^{rt}} = \frac{e^{Z_t}}{e^{rt}}.$$

Substituting $Z_t = rt + L_t$ yields

$$\bar{\gamma}_t = e^{L_t},$$

such that the process $\bar{\gamma}_t$ is a $\mathbb{P}$ -martingale and has $\mathbb{P}$ -expectation equal to 1. From Bayes' change-of-law rule

$$\bar{\mathbb{E}}(e^{iuZ_t}) = \mathbb{E}(\bar{\gamma}_t e^{iuZ_t}) = \mathbb{E}(e^{Z_t - rt} e^{iuZ_t}) = e^{-rt} \mathbb{E}(e^{i(u-i)Z_t}),$$

where the last equality follows from (6.38). ■

Proposition 23 *Define $\Upsilon_t = \int_0^t v_{s-} ds$ for activity process v . Suppose Υ_t is L -continuous, i.e., L is constant on all intervals $[\Upsilon_{t-}, \Upsilon_t]$, $t > 0$.*

We identify the local characteristics of the subordinated process L_{Υ_t} under $\bar{\mathbb{P}}$ in terms of the Lévy characteristics (μ, σ^2, Π) of L_t under $\mathbb{P}$ as

$$\left(\left(\mu + \sigma^2 - \int_{\{|l| \leq 1\} \setminus \{0\}} l(1 - e^l) \Pi(dl) \right) \Upsilon_{t-}, \sigma^2 \Upsilon_{t-}, e^l \Pi(dl) d\Upsilon_{t-} \right).$$

Proof. The proof follows from Küchler and Sørensen ((1997), p. 230) (see also Carr and Wu (2004), Proposition A.1). ■

Proposition 23 implies that the local characteristics of the subordinated process L_{Υ_t} under $\bar{\mathbb{P}}$ are found from those of L_t by applying the random time transformation Υ_{t-} in place of the deterministic time t .

Appendix 6.C—Characterization of the log-return distribution conditional on the variance at the endpoints of a time interval

Consider the log-return process given by (6.37), and define $Z_{t+\delta t} = X_{t+\delta t} - X_t$. Propositions 22-23 (with the Heston asset price process viewed as a subordinated Brownian motion) permit

$$\bar{\mathbb{E}} \left(e^{iuZ_{t+\delta t}} \middle| v_t, v_{t+\delta t} \right) = e^{-r\delta t} \mathbb{E} \left(e^{i(u-i)Z_{t+\delta t}} \middle| v_t, v_{t+\delta t} \right).$$

By virtue of the tower property of expectations we have

$$\begin{aligned} & \mathbb{E} \left(e^{i(u-i)Z_{t+\delta t}} \middle| v_t, v_{t+\delta t} \right) \\ = & \mathbb{E} \left(\mathbb{E} \left(e^{i(u-i)Z_{t+\delta t}} \middle| \int_t^{t+\delta t} v_s ds \right) \middle| v_t, v_{t+\delta t} \right) \\ = & e^{i(u-i)((r-\rho\alpha\beta\eta^{-1})\delta t + \rho\eta^{-1}(v_{t+\delta t} - v_t))} \mathbb{E} \left(\exp \left(i(u-i)(\rho\alpha\eta^{-1} - 1/2) \int_t^{t+\delta t} v_s ds \right) \right. \\ & \times \mathbb{E} \left(\exp \left(i(u-i)\sqrt{1-\rho^2} \int_t^{t+\delta t} \sqrt{v_s} dB_s \right) \middle| \int_t^{t+\delta t} v_s ds \right) \middle| v_t, v_{t+\delta t} \right) \\ = & e^{i(u-i)((r-\rho\alpha\beta\eta^{-1})\delta t + \rho\eta^{-1}(v_{t+\delta t} - v_t))} \\ & \times \mathbb{E} \left(\exp \left(i(u-i) (\rho\alpha\eta^{-1} - 1/2 + i(u-i)(1-\rho^2)/2) \int_t^{t+\delta t} v_s ds \right) \middle| v_t, v_{t+\delta t} \right) \end{aligned}$$

Broadie and Kaya ((2006), Appendix) derive the characteristic function of $\left(\int_t^{t+\delta t} v_s ds \middle| v_t, v_{t+\delta t} \right)$

$$\mathbb{E} \left(e^{iu \int_t^{t+\delta t} v_s ds} \middle| v_t, v_{t+\delta t} \right) =: \Phi(u; v_t, v_{t+\delta t}),$$

where

$$\begin{aligned}
 \Phi(u; v_t, v_{t+\delta t}) &= \Phi_1(u) e^{(v_t + v_{t+\delta t})\Phi_2(u)} \frac{I_{2\alpha\beta\eta^{-2}-1}(\Phi_3(u) \sqrt{v_t v_{t+\delta t}})}{I_{2\alpha\beta\eta^{-2}-1}(\Phi_4 \sqrt{v_t v_{t+\delta t}})}, \\
 \Phi_0(u) &= \sqrt{\alpha^2 - 2iu\eta^2}, \\
 \Phi_1(u) &= \Phi_0(u) \frac{e^{-\frac{1}{2}(\Phi_0(u)-\alpha)\delta t}}{1 - e^{-\Phi_0(u)\delta t}} \frac{1 - e^{-\alpha\delta t}}{\alpha}, \\
 \Phi_2(u) &= \frac{1}{\eta^2} \left(\frac{\alpha(1 + e^{-\alpha\delta t})}{1 - e^{-\alpha\delta t}} - \frac{\Phi_0(u)(1 + e^{-\Phi_0(u)\delta t})}{1 - e^{-\Phi_0(u)\delta t}} \right), \\
 \Phi_3(u) &= \frac{2\Phi_0(u)}{\eta^2 \sinh(\frac{1}{2}\Phi_0(u)\delta t)}, \\
 \Phi_4 &= \frac{2\alpha}{\eta^2 \sinh(\frac{1}{2}\alpha\delta t)},
 \end{aligned}$$

and $I_\nu(\cdot)$ denotes the ν^{th} -order modified Bessel function of the first kind. Conditions to ensure continuity of the characteristic function Φ have been derived in Lord and Kahl ((2008), Theorem 6).

Hence, we obtain

$$\begin{aligned}
 \bar{\mathbb{E}}(e^{iuZ_{t+\delta t}} | v_t, v_{t+\delta t}) &= e^{(iu(r-\rho\alpha\beta\eta^{-1})-\rho\alpha\beta\eta^{-1})\delta t + i(u-i)\rho\eta^{-1}(v_{t+\delta t}-v_t)} \\
 &\quad \times \Phi((u-i)(\rho\alpha\eta^{-1} - 1/2 + i(u-i)(1-\rho^2)/2); v_t, v_{t+\delta t}) \quad (6.39)
 \end{aligned}$$

In addition to the Heston model, we obtain for the Bates model

$$\begin{aligned}
 \bar{\mathbb{E}}(e^{iuZ_{t+\delta t}} | v_t, v_{t+\delta t}) &= e^{(iu(r-\rho\alpha\beta\eta^{-1}-\lambda_J\tilde{\mu}_J(-i))-\rho\alpha\beta\eta^{-1}-\lambda_J\tilde{\mu}_J(-i))\delta t + i(u-i)\rho\eta^{-1}(v_{t+\delta t}-v_t)} \\
 &\quad \times \Phi((u-i)(\rho\alpha\eta^{-1} - 1/2 + i(u-i)(1-\rho^2)/2); v_t, v_{t+\delta t}) e^{\lambda_J\tilde{\mu}_J(u-i)\delta t}; \\
 \tilde{\mu}_J(u) &= e^{iu\mu_J - \frac{1}{2}\sigma_J^2 u^2} - 1. \quad (6.40)
 \end{aligned}$$

Chapter 7

Monte Carlo option pricing coupled with Fourier transformation

7.1 Introduction

Within stock option pricing applications, we have seen that it is common to model the stock log-returns by Lévy models. Popular are the variance gamma process of finite variation with infinite activity of jumps, the normal inverse Gaussian and the generalized hyperbolic processes of infinite variation. For greater flexibility, the family of tempered stable models has been introduced with a Lévy measure that allows for processes with finite activity, infinite activity and finite variation, and infinite variation. A 5-parameter version (also known as KoBoL) originates from Koponen (1995), and has been considered in financial applications in Boyarchenko and Levendorskiĭ (2002). Carr et al. (2002) study a 4-parameter subclass known as the CGMY model, which itself generalizes the original VG framework.

With a view to enhancing the applicability of the CGMY model in asset modelling, Madan and Yor (2008) develop the representation of this model as a time-changed Brownian motion with drift. Their construction reflects a time change which is absolutely continuous with respect to a one-sided stable subordinator. Given this, they propose to simulate the increments of the time-change process by simulating the big jumps and replacing the small ones with their expectation. This approach yields a compound Poisson approximation to the distribution of the random time increments, where the jump-size random variables are sampled using the rejection

method by Rosiński (2001) (see also Cont and Tankov (2004a)), leading to biased simulation of the CGMY trajectories. To skip any approximation error, Poirot and Tankov (2006) construct a new probability measure under which the original tempered stable process reduces to a stable process whose exact simulation is well-established. Although faster in execution, the Poirot-Tankov method does not provide access to the entire trajectory of the process, prohibiting the pricing of path-dependent products.

Motivated by the previous concerns, we set up a general, efficient Monte Carlo simulation scheme coupled with Fourier transformation which is simple and fast to implement. Its efficiency and generality are attributed to its ability to generate unbiased sample trajectories for any stochastic process which admits a closed-form characteristic function. This encompasses mainly Lévy models, but also the Heston stochastic volatility model whose unbiased simulation coupled with Fourier transforms has been the area of concern for Broadie and Kaya (2006) and Glasserman and Kim (2008a). Alternatively, the VG and NIG frameworks allow for well-structured time-change representations which render their exact simulation (see Glasserman (2004), Section 3.5.2) straightforward and faster. This opposes the tempered stable process whose deficient simulation so far allows for a good reason to pay particular attention to it.

After we develop the theoretical framework for the joint Monte Carlo-Fourier transform scheme in Section 7.2, we lead through its efficient numerical implementation. In Section 7.3 we present the market model and the pricing of contingent claims via Monte Carlo, while in Section 7.4 we focus on the CGMY model. In Section 7.5 we price European plain vanilla and discretely sampled Asian options under the CGMY assumption, by applying the simulation algorithm of Section 7.2.1 directly on the increments of the CGMY process and, alternatively, on the random time increments in the Madan-Yor time-change representation of the CGMY process. We then compare them with the results from the Poirot-Tankov and Madan-Yor methods. Section 7.6 concludes the chapter.

7.2 Monte Carlo simulation coupled with Fourier transform

Denote by F_X the cumulative distribution function of some random variable X . A random sample from the distribution of X can be drawn via the inverse distribution function $F_X^{-1}(U)$,

where $U \sim \text{Unif}[0, 1]$. In general, standard Lévy distribution laws do not admit a distribution function in closed form, still one can resort to numerical inversion techniques to retrieve this via the characteristic function. This is not a trivial task since $F_X(x)$ does not decay to zero as $x \rightarrow \infty$, hence does not satisfy sufficient integrability condition for the existence of its Fourier transform (see Definition 1). To solve this issue, one can make use of the auxiliary function $\tilde{F}_X(x)$ proposed by Hughett ((1998), Lemma 9),

$$\tilde{F}_X(x) = F_X(x) - \frac{1}{2}F_X(x - \xi) - \frac{1}{2}F_X(x + \xi), \quad \xi > 0,$$

which is well-behaved in the sense that both itself and its Fourier transform $\mathcal{F}(\tilde{F}_X)$ decay rapidly to zero. Given the analytical expression (2.8) for $\mathcal{F}(\tilde{F}_X)$, we can recover $\tilde{F}_X$ via

$$\tilde{F}_X = \mathcal{F}^{-1}(\mathcal{F}(\tilde{F}_X)) \quad (7.1)$$

by virtue of Theorem 7. For sufficiently large $\xi > 0$,

$$F_X(x) \approx \tilde{F}_X(x) + \frac{1}{2} \quad (7.2)$$

for $|x| \leq \frac{1}{2}\xi$. Approximating F_X through $\tilde{F}_X$ which satisfies certain regularity conditions, i.e., fast decay to zero for both $\tilde{F}_X$ and $\mathcal{F}(\tilde{F}_X)$, generates the so-called regularization error. We observe that as ξ increases, the regularization error decreases.

Next, we lead through the efficient implementation of the simulation scheme equipped by Fourier transform presented above.

7.2.1 Numerical implementation

Computation of the distribution function via Fourier transform inversion

To evaluate numerically $\tilde{F}_X = \mathcal{F}^{-1}(\mathcal{F}(\tilde{F}_X))$, we select evenly spaced, symmetric about zero, grids $\mathbf{u} = \{u_0 + j\delta u\}_{j=0}^{N-1}$ and $\mathbf{x} = \{x_0 + l\delta x\}_{l=0}^{N-1}$ with N grid points and spacings δu and $\delta x = \xi/N$. The range of values of $\mathbf{u}$ is determined to ensure that $|\mathcal{F}(\tilde{F}_X)| < \epsilon'$ outside $\mathbf{u}$ for some tolerance level ϵ' , e.g., $\epsilon' = 10^{-15}$. We denote the function values $\mathcal{F}(\tilde{F}_X)$ on grid $\mathbf{u}$ by

$\tilde{\mathbf{D}}_X$, and evaluate the inverse Fourier transform (7.1) on the grid $\mathbf{x}$ by computing

$$\tilde{\mathbf{F}}_X = \frac{1}{2\pi} \mathcal{D}(\tilde{\mathbf{D}}_X, -\mathbf{u}, \mathbf{x}; -N\delta u \delta x / 2\pi) \delta u \quad (7.3)$$

using the conversion (2.24). Then, from (7.2), we obtain

$$\mathbf{F}_X \approx \tilde{\mathbf{F}}_X + \frac{1}{2}.$$

The finite Fourier series¹ (7.3) that approximates the auxiliary distribution function generates two sources of error (excluding any round-off error): the discretization error induced from the evaluation of the integrand at specific points only, and the truncation error by truncating the Fourier series above and below. The truncation error is controlled by the (finite) number of points taken, whereas the discretization error by the interval between successive points. To determine N and ξ such that the overall approximation error (including the error from regularizing F_X) is below a pre-specified level ϵ , one may consider the error bound developed in Hughett ((1998), Theorem 10) for continuous distributions with finite variance, based solely on knowledge of the associated characteristic function.

Standard consistency checks on the computed distribution include inspections on the minimum and maximum values, which should lie within ϵ of 0 and 1 respectively, and the intermediate pattern, which should be nonnegative and monotonically increasing. Comparison of the numerically computed moments (using the approximated distribution) against the true ones indicates the existence of any approximation error (see Section 7.5.1).

Simulation procedure

As explained earlier, key to the method presented here is to generate a uniform random variable $U \sim \text{Unif}[0, 1]$, and find x such that $F_X(x) = U$; the $x = F_X^{-1}(U)$ value returned is a sample from F_X . We identify X with a Lévy increment $L_{\delta t}$ over interval δt . For time horizon T divided

¹Implementing the discrete Fourier transform via FFT-based routines allows us to pre-evaluate the distribution function fast on a grid with a single inversion, and store this for later use in the simulation. This is a substantial CPU power saving over the popular inversion formula (and variations of it) of Abate and Whitt (1992), but also the recent Fourier-cosine expansion method of Fang and Oosterlee (2008a), which evaluate the distribution at a single point per transform inversion.

into $n > 0$ equidistant time steps of length $\delta t = T/n$, we sample n Lévy increments $L_{\delta t}$ per simulation trial (in total $m > 0$ trials) as follows²:

1. Generate a $n \times m$ matrix $\mathbf{U}$ of uniform random variables using the MATLAB pseudorandom number generator RAND.
2. Given the pre-tabulated distribution $\mathbf{F}_{L_{\delta t}}$ evaluated on the uniform grid $\mathbf{x}$, we need to find the grid values $\{\mathbf{F}_{L_{\delta t}, l}\}_{l=0}^{N-1}$ where the elements of $\mathbf{U}$ lie in-between, for the purpose of approximating the corresponding $\hat{\mathbf{x}}$ -samples in the next step. To speed up the search through the tabulated $\{\mathbf{F}_{L_{\delta t}, l}\}_{l=0}^{N-1}$ values, we employ the cutpoint method (e.g., Fishman (1996), Section 3.3).
3. We obtain the $\hat{\mathbf{x}}$ -samples, such that $F_{L_{\delta t}}(\hat{\mathbf{x}}) = \mathbf{U}$, by interpolating linearly on the nodes $(\mathbf{F}_{L_{\delta t}}, \mathbf{x})$ using the MATLAB function INTERP1. For greater accuracy, though at higher computational cost, one may consider fitting a cubic interpolating spline instead (see Section 7.5.1).

In the next section, we define the problem of contingent claims pricing in the Monte Carlo context.

7.3 Market model setup and option pricing

Fix a terminal time $T > 0$. Assume constant continuously compounded interest rate $r > 0$. Fix constant $S_0 > 0$ and define the price process of a risky asset as

$$S_t = S_0 e^{(r+\varpi)t + L_t}, \quad 0 < t \leq T,$$

where L is a Lévy process. The mean-adjusting parameter ϖ is chosen such that the martingale condition $\mathbb{E}(S_t) = S_0 e^{rt}$ applies under a risk-neutral measure $\mathbb{P}$.

²Since we are sampling independent and stationary Lévy increments, we compute (7.3) once, and pre-cache the distribution for use in all simulation trials and all time steps (for an equidistant time grid). This opposes to Broadie and Kaya's exact simulation scheme for the (non-Lévy) Heston model, which requires repeated transform inversion, depending on the number of time steps employed, significantly raising the CPU timing.

Consider the problem of pricing a contingent claim maturing at T with terminal payoff given by $\pi_T = p(\{S_t; 0 \leq t \leq T\})$. By the fundamental theorem of asset pricing, the arbitrage-free price of the contract at inception is given by

$$\pi_0 = \mathbb{E} \left(e^{-rT} \pi_T \right).$$

For a discretely monitored path-dependent contract, e.g., Asian, barrier, lookback, the process S is observed at fixed discrete points in time (assumed equidistant, though this is not a requirement) $0 = t_0 < t_1 < \dots < t_n = T$. To estimate π_0 , it is necessary that we sample the process S at the monitoring dates $\{t_k\}_{k=1}^n$, $n \in \mathbb{N}^*$, using Monte Carlo. This generates a collection of samples $\{\hat{\pi}_T^j\}_{j=1}^m$, for m simulation runs. The Monte Carlo estimate of the contract price, $\hat{\pi}_0$, is obtained as

$$\hat{\pi}_0 = e^{-rT} \frac{\sum_{j=1}^m \hat{\pi}_T^j}{m}.$$

7.4 The tempered stable framework

Before applying the joint Monte Carlo-Fourier transform method in option pricing, let us first briefly review the properties and important results about the class of tempered stable processes, with particular focus on the CGMY subclass (see Carr et al. (2002)), that we make critical use of in our numerical study in Section 7.5.

7.4.1 Properties

The one-dimensional tempered stable process is constructed by taking a one-dimensional stable process and multiplying the Lévy measure by exponentially decaying factors on each half of the real axis. Thus, we obtain a Lévy measure associated to the tempered stable process of the form

$$\Pi(x) = c_+ \frac{e^{-\lambda_+ |x|}}{|x|^{1+\alpha}} 1_{\{x>0\}} + c_- \frac{e^{-\lambda_- |x|}}{|x|^{1+\alpha}} 1_{\{x<0\}}, \quad (7.4)$$

with parameters $c_+ > 0$, $c_- > 0$, $\lambda_+ \geq 0$, $\lambda_- \geq 0$, and $\alpha < 2$. For greater flexibility, the 5-parameter version (7.4) can be extended to allow for different values of α on the two sides of the real axis. This yields the generalized 6-parameter tempered stable model with Lévy

measure

$$\Pi(x) = c_+ \frac{e^{-\lambda_+|x|}}{|x|^{1+\alpha_+}} 1_{\{x>0\}} + c_- \frac{e^{-\lambda_-|x|}}{|x|^{1+\alpha_-}} 1_{\{x<0\}}, \quad (7.5)$$

with $\alpha_+ < 2$ and $\alpha_- < 2$. Parameters $\lambda_{\pm}$ determine the tail behaviour of the Lévy measure, i.e., how far the process may jump, $c_{\pm}$ tell us about the arrival rate of jumps of given size, while $\alpha_{\pm}$ determine the local behaviour of the process between big jumps. When $\alpha_+ \geq 1$ and/or $\alpha_- \geq 1$, the process exhibits infinite variation (many small oscillations observed between big jumps), whilst if $\alpha_+ < 1$, $\alpha_- < 1$ and $\alpha_+ \geq 0$ and/or $\alpha_- \geq 0$, the process has trajectories of infinite activity and finite variation (relative calmness observed between big jumps) (see Cont and Tankov (2004a), Section 4.5). Common in the option pricing literature is the 4-parameter CGMY process, where $c_+ = c_-$ and $\alpha_+ = \alpha_-$ apply such that

$$\Pi(x) = C \frac{e^{-M|x|}}{|x|^{1+Y}} 1_{\{x>0\}} + C \frac{e^{-G|x|}}{|x|^{1+Y}} 1_{\{x<0\}}.$$

Explicit knowledge of the Lévy measure allows one to derive the characteristic function describing the law of the tempered stable distribution via the Lévy-Khintchine formula. In particular, the characteristic function of the CGMY process L_t^{CGMY} is

$$E\left(e^{iuL_t^{\text{CGMY}}}\right) = \exp(tC\Gamma(-Y)((M-iu)^Y - M^Y + (G+iu)^Y - G^Y))$$

(see Carr et al. (2002), Theorem 1). For a derivation of the characteristic function applying in the generalized tempered stable case, the reader is referred to Poirot and Tankov ((2006), Section 2.5).

7.4.2 CGMY as time-changed Brownian motion

Madan and Yor (2008) construct the CGMY process by randomly changing the time in a Brownian motion with drift, i.e.,

$$L_t^{\text{CGMY}} = \theta_1 Z_t + W_{Z_t}, \quad (7.6)$$

for $Y \in (0, 2)$, $\theta_1 = \frac{G-M}{2}$ and an increasing zero-drift time-change process Z independent of the Brownian motion W . The process Z has Lévy measure

$$\zeta(\tau) = \chi(\tau) v(\tau),$$

where

$$\begin{aligned} \chi(\tau) &= \frac{2^{\frac{Y}{2}} \Gamma\left(\frac{Y}{2} + \frac{1}{2}\right) e^{\tau \frac{\theta_1^2}{2} - \tau \frac{\theta_2^2}{4}}}{\sqrt{\pi}} D_{-Y}(\theta_2 \sqrt{\tau}) \\ &= \frac{\Gamma\left(\frac{Y}{2} + \frac{1}{2}\right) e^{\tau \frac{\theta_1^2}{2} - \tau \frac{\theta_2^2}{2}}}{\sqrt{\pi}} U\left(\frac{Y}{2}, \frac{1}{2}, \tau \frac{\theta_2^2}{2}\right) \leq 1, \end{aligned}$$

$\theta_2 = \frac{G+M}{2}$, $D(\cdot)$ denotes Whittaker's parabolic cylinder function, $U(\cdot, \cdot, \cdot)$ the confluent hypergeometric function of the second kind (see Abramowitz and Stegun (1968)), and

$$v(\tau) = \frac{\Theta}{\tau^{\frac{Y}{2}+1}} 1_{\{\tau>0\}}; \quad \Theta = \frac{2^{-\frac{Y}{2}} \sqrt{\pi}}{\Gamma\left(\frac{Y}{2} + \frac{1}{2}\right)} C$$

is the Lévy density of the one-sided $\frac{Y}{2}$ -stable subordinator.

By exploiting the absolute continuity of Z to the $Y/2$ -stable subordinator, Madan and Yor (2008) suggest approximating Z with a drifted compound Poisson subordinator where small jumps (of size smaller than ε) are replaced by their expected value. Z is then approximated using the rejection method by Rosiński (2001), which amounts to accepting every jump of the compound Poisson subordinator with size τ_i for which $\chi(\tau_i)$ is greater than an independent random variable $U_i \sim \text{Unif}[0, 1]$. This yields

$$Z_t \approx Z_t^\varepsilon = b_\varepsilon t + \sum_i \tau_i 1_{\{\Psi_i \leq t\}} 1_{\{\chi(\tau_i) > U_i\}}, \quad (7.7)$$

where $b_\varepsilon = \frac{\Theta \varepsilon^{1-Y/2}}{1-Y/2}$. $\{\Psi_i\}$ denote the jump times of a compound Poisson process with intensity $\lambda_\varepsilon = \frac{\Theta \varepsilon^{-Y/2}}{Y/2}$, and independent jump sizes $\{\tau_i\}$ with cumulative distribution function $G(\tau; \varepsilon) = 1 - (\tau/\varepsilon)^{-Y/2}$. In Section 7.5, we examine the impact of generating CGMY random variates using the construction (7.6-7.7).

Moreover, Madan and Yor (2008) provide us with the Laplace transform of the subordinator

Z

$$E(e^{-\eta Z_t}) = \exp(tC\Gamma(-Y)(2(2\eta + GM)^{\frac{Y}{2}} \cos(l(\eta)Y) - M^Y - G^Y)); \quad (7.8)$$

$$l(\eta) = \arctan\left(\frac{\left(2\eta - \left(\frac{G-M}{2}\right)^2\right)^{\frac{1}{2}}}{\frac{G+M}{2}}\right).$$

Result (7.8) is considered in our practical application of the joint Monte Carlo-Fourier transform scheme in the next section.

7.5 Numerical study

In order to illustrate the performance of the joint Monte Carlo-Fourier transform method in option pricing, we assume an underlying driven by an exponential CGMY model whose simulation so far has proved difficult. For comparison, we consider two separate implementations: in the first case, we directly simulate the increments of the CGMY process (implem. I), whereas in the second we exploit the Madan-Yor time-change representation (7.6), and simulate separately the stochastic time increments described by (7.8) and the arithmetic Brownian motion subject to a time change (see Glasserman (2004), Section 3.5.2, for the simulation of subordinated Brownian motions) (implem. II). We apply to pricing a (path-independent) plain vanilla put option with terminal payoff

$$\pi_T = (K - S_T)^+,$$

where K is the strike price, and a discretely monitored (path-dependent) fixed-strike Asian call option with payoff

$$\pi_T = \left(\frac{1}{n + \gamma} \left(\gamma S_0 + \sum_{k=1}^n S_{t_k}\right) - K\right)^+, \quad (7.9)$$

where n denotes the number of monitoring dates and coefficient γ takes value 1 (0) when S_0 is (is not) included in the average. The two options can be priced numerically (without Monte Carlo) at high precision using the Fourier-cosine algorithm of Fang and Oosterlee (2008a) and the convolution algorithm of Chapter 4 respectively, and therefore can serve as benchmarks to the Monte Carlo estimates.

In this exercise, we test the exact scheme in Section 7.2.1 by comparing with the results from

set	C	G	M	Y
I	0.5	2.0	3.5	0.5
II	0.1	2.0	3.5	1.5

Table 7.1: Model parameters

the Poirrot-Tankov (PT) exact scheme for path-independent contracts, and the approximate Madan-Yor (MY) method. Descriptions of both schemes can be found in Appendix 7 and Section 7.4.2. All pricing comparisons are run in MATLAB R2008b on a Dell Optiplex 755 Intel Core 2 Duo PC 2.66GHz with 2.0GB RAM.

The reference parameter sets in Table 7.1 are from Poirrot and Tankov (2006). Set I corresponds to a model of finite variation and infinite activity, while for set II the CGMY process exhibits infinite variation. The reason for two parameterizations is to illustrate the effect of a change from a finite to an infinite variation model on the bias and computational complexity of the MY scheme (for details, see Section 7.5.2).

7.5.1 Distribution function tests

Given that the distribution function is computed numerically as illustrated in Section 7.2.1, we need a way of assessing its accuracy. This is a nontrivial task, since both the cumulative distribution and density functions of the CGMY model are not available in closed form. For this, as an indication of the quality of the distribution function approximation, we choose to compute numerically the first four moments and compare with their true values for $\delta t = T/n$ (T : time to maturity, n : number of sampling points). For example, in the implem. I (direct simulation of the CGMY increments) for the parameter set I and $\delta t = 1/4$, the numerical mean, variance, skewness coefficient, and excess kurtosis agree with the true values at -0.03823712808, 0.056084217, -1.6657160, and 13.319987 respectively, while for $\delta t = 1/12$ they agree at -0.012745737603, 0.018694739, -2.8851047, and 39.959960. The same accuracy level is also reached with the parameter set II.

Moreover, relative to the third step of the simulation procedure in Section 7.2.1, our study on the tabulated distribution $\mathbf{F}_{L_{\delta t}^{\text{CGMY}}}$ evaluated on the symmetric uniform grid $\mathbf{x}$ with $\xi = 24$, $N = 2^{19}$ and $\delta x = 0.000045$ has shown that the use of linear interpolation yields error of the

order 10^{-9} and takes 0.49 seconds of CPU time for 10^6 trials. Instead, fitting a cubic spline reduces the order of the error to 10^{-15} , while significantly raises the CPU timing to 1.4-1.5 seconds for 10^6 trials. To avoid raising unnecessarily the computational burden, we adhere to linear interpolation.

7.5.2 Simulation tests

We first compare the results from the two suggested implementations (implem. I & II) of the exact scheme in Section 7.2.1 with the ones from the exact PT scheme in pricing European plain vanilla put options. Table 7.2 illustrates the Monte Carlo price estimates for a range of strikes when the parameter set I applies. In general, we observe that the more in-the-money the put option is, the higher the simulation error generated. However, implem. I produces estimators with smaller standard error than the PT estimators and, for deep in-the-money options, the error reduces to almost half levels. We observe similarly for the second parameterization.

Implem. II achieves same standard errors as implem. I, nevertheless for 10^6 trials it requires higher CPU time (implem. II: 9.5 seconds versus implem. I: 1.3 seconds). This is due to the oscillatory decay of the characteristic function of the random time increment Z (see equation (7.8)) which affects the convergence of the Fourier series expansion (7.3) of the distribution function F_Z . The order of convergence is exacerbated further for small time intervals δt . To ensure accurate evaluation of the distribution, it is essential that we employ a wide and finely refined Fourier grid $\mathbf{u}$ at the cost of increased computational effort.

Due to higher CPU requirements of implem. II, we focus solely on implem. I and compare against the PT method in terms of efficiency: we define the efficiency ratio

$$\mathcal{E}_{A|B} := \frac{t_B \sigma_B^2}{t_A \sigma_A^2},$$

where $\sigma^2 = \mathbb{E}((\hat{\pi}_0 - \mathbb{E}(\hat{\pi}_0))^2)$ is the variance of the $\hat{\pi}_0$ estimator obtained in t seconds via the indicated method. When $\mathcal{E}_{A|B} > 1$ we say that method A is more efficient than method B , and vice versa if $\mathcal{E}_{A|B} < 1$; for more on the study of efficiency we refer to Glynn and Whitt (1992). We investigate for $T = \delta t \in \{0.25, 1.0\}$, $K \in \{80, 100, 120\}$, $Y \in \{0.5, 1.5\}$, as shown in Table 7.3. Given T , the efficiency gains of the joint Monte Carlo-Fourier transform method become

K	ref. prices	implem. I (10^6 trials)	std error	implem. II (10^6 trials)	std error	PT (10^6 trials)	std error
80	1.7444	1.748	0.006	1.743	0.006	1.751	0.005
85	2.3926	2.396	0.007	2.391	0.007	2.403	0.006
90	3.2835	3.288	0.008	3.280	0.008	3.299	0.008
95	4.5366	4.532	0.010	4.544	0.010	4.560	0.010
100	6.3711	6.373	0.011	6.360	0.011	6.397	0.012
105	9.1430	9.142	0.012	9.135	0.012	9.168	0.015
110	12.7631	12.767	0.013	12.764	0.013	12.789	0.019
115	16.8429	16.847	0.014	16.838	0.014	16.867	0.023
120	21.1855	21.180	0.015	21.183	0.015	21.207	0.028

Table 7.2: European plain vanilla put price estimates (with standard errors) computed for the parameter set I (see Table 7.1) using the joint Monte Carlo-Fourier transform method (implem. I & II) and the PT method. Reference prices computed via the Fourier-cosine series expansion (3.13). Fixed parameters: $S_0 = 100$, $r = 0.04$, $T = 0.25$.

T	K	param. I	param. II
0.25	80	0.2	0.7
	100	0.4	1.1
	120	1.2	1.7
1.0	80	4.0	4.9
	100	8.3	6.7
	120	14.7	8.6

Table 7.3: Efficiency gains for European plain vanilla put options. Efficiency ratios computed as $\mathcal{E}_{\text{implem.I|PT}}$: CPU timings and variances used are for 10^6 trials. Model parameters: Table 7.1. Fixed parameters: $S_0 = 100$, $r = 0.04$.

existent following increases in the moneyness of the option and the Y parameter value. The efficiency gains are more significant for higher T . The improvement in efficiency across strikes and times to maturity are attributed to increases in the variance of the PT estimator, while across Y to increases in the CPU timing of the stable random number generator (see Appendix 7). The CPU timings of the joint Monte Carlo-Fourier transform technique remain unaffected by changes in T , K and Y .

A remarkable advantage of the joint Monte Carlo-Fourier transform scheme is in pricing path-dependent options for which the intermediate values of the underlying are needed but cannot be accessed through the PT method. Table 7.4 illustrates the performance of the

simulation scheme in pricing Asian call options with payoff (7.9), subject to quarterly and monthly monitoring over a year's time to maturity. We consider only the case where we simulate directly the CGMY increments as this has shown to be faster. With 10^7 sample trajectories and quarterly observation, implem. I takes 127 seconds to produce price estimates with penny accuracy (at 99% confidence level). With monthly monitoring, the CPU timing per price triples while the increase in the standard error is slight. With 10^6 trials the running time reduces by a factor of 10, nevertheless only 1 decimal place of accuracy is then guaranteed, demonstrating the deteriorating effect of path-dependence (frequent sampling) on the quality of the simulation outcome compared to the path-independent case (sampling only at maturity) (see fourth column of Table 7.2). Moreover, as with the plain vanilla option, the standard error increases as the Asian call moves deeper into the money, whilst it remains almost unaffected by the model parameters choice.

The other candidate for the simulation of the CGMY trajectories is the biased MY method³. From Section 7.4.2, ε is the threshold below which we approximate the jumps of the subordinator Z in the CGMY representation (7.6) by their expected value. An approximation bias is then induced, which depends on ε and the model parameter Y . In general, if $\hat{\pi}_0$ is the Monte Carlo estimator for the derivative's price today and π_0 the true price, then the bias of the estimator for given ε and Y is computed as $\text{bias} = \mathbb{E}(\hat{\pi}_0 - \pi_0)$. The root mean square error (RMSE) of the estimator $\hat{\pi}_0$ with variance σ^2 is given by

$$\text{RMSE} = (\text{bias}^2 + \sigma^2)^{1/2}. \quad (7.10)$$

Following Poirot and Tankov (2006), we use 10^7 trials to estimate the bias resulting from different ε and Y . The results are presented in Table 7.5 for at-the-money options. We observe that the bias decreases as $\varepsilon \rightarrow 0$, while the computational burden is approximately inversely proportional to $\varepsilon^{Y/2}$. In particular, for set I ($Y = 0.5$), the bias reduces roughly by factors in the range 3–20 (plain vanilla put) and 2–14 (Asian call) following successive reductions of ε by factors of 10, while the CPU timing increases roughly by a factor of $10^{1/4} \approx 1.77$ per ε reduction by 10. The approximation is exacerbated in the infinite variation model (set II, $Y = 1.5$) where

³For the construction of the MY scheme in MATLAB, we have used as our basis the C++ code of Peter Tankov which is downloadable from <http://people.math.jussieu.fr/~tankov/>.

param.	n	K	ref. prices	implem. I (10^7 trials)	std error	implem. I (10^6 trials)	std error
I	4	80	22.8629	22.8704	0.0073	22.869	0.023
		90	15.3053	15.2950	0.0068	15.256	0.021
		100	9.4395	9.4339	0.0061	9.425	0.019
		110	5.6734	5.6774	0.0053	5.681	0.017
		120	3.5537	3.5537	0.0046	3.555	0.014
	12	80	23.0533	23.0410	0.0074	22.986	0.023
		90	15.5249	15.5190	0.0069	15.530	0.022
		100	9.6433	9.6413	0.0062	9.610	0.019
		110	5.8405	5.8409	0.0055	5.844	0.017
		120	3.6888	3.6961	0.0048	3.681	0.016
II	4	80	22.9249	22.9275	0.0075	22.967	0.024
		90	15.9385	15.9374	0.0068	15.947	0.022
		100	10.6216	10.6103	0.0060	10.638	0.019
		110	6.8788	6.8705	0.0051	6.898	0.016
		120	4.3898	4.3833	0.0042	4.417	0.013
	12	80	23.1589	23.1556	0.0077	23.176	0.024
		90	16.2348	16.2337	0.0070	16.255	0.022
		100	10.9196	10.9179	0.0061	10.910	0.019
		110	7.1342	7.1312	0.0052	7.136	0.016
		120	4.5865	4.5836	0.0044	4.583	0.014

Table 7.4: Fixed-strike Asian call price estimates (with standard errors) computed using the joint Monte Carlo-Fourier transform method (implem. I). Reference prices computed via the convolution algorithm of Chapter 4. Model parameters: Table 7.1. Fixed parameters: $S_0 = 100$, $r = 0.04$, $T = 1.0$, $\gamma = 1$.

option	parameterization I				parameterization II			
	ε	bias	std error	CPU (s)	ε	bias	std error	CPU (s)
vanilla put	10^{-1}	2.0915	0.0040	508	10^{-1}	0.1612	0.0041	106
	10^{-2}	0.2104	0.0038	863	10^{-2}	0.3778	0.0038	541
	10^{-3}	0.0105	0.0038	1676	10^{-3}	0.0673	0.0037	2942
	10^{-4}	-0.0013	0.0038	2938	10^{-4}	0.0155	0.0037	16441
Asian call (n=4)	10^{-1}	1.0868	0.0066	1994	10^{-1}	1.0369	0.0069	575
	10^{-2}	0.0774	0.0061	3291	10^{-2}	0.2222	0.0062	2333
	10^{-3}	-0.0019	0.0060	5583	10^{-3}	0.0305	0.0060	12095
	10^{-4}	-0.0088	0.0060	9661	10^{-4}	0.0178	0.0060	66825
Asian call (n=12)	10^{-1}	1.1797	0.0066	2520	10^{-1}	1.1076	0.0070	597
	10^{-2}	0.0797	0.0062	4151	10^{-2}	0.2413	0.0063	2340
	10^{-3}	0.0083	0.0062	7038	10^{-3}	0.0341	0.0062	12017
	10^{-4}	-0.0043	0.0062	12144	10^{-4}	0.0092	0.0061	66141

Table 7.5: Estimated biases (with standard errors) of the Madan-Yor approximation for different threshold ε and parameter Y based on 10^7 trials. Model parameters: Table 7.1. Fixed parameters: $S_0 = K = 100$, $r = 0.04$. Left-Top (Right-Top) panel: European put option; additional parameters: $T = 0.25$; ref. price 6.3711 (8.3014) computed via the Fourier-cosine series expansion (3.13). Left-Mid (Right-Mid) panel: Asian call option; additional parameters: $T = 1.0$, $n = 4$, $\gamma = 1$; ref. price 9.4395 (10.6216) computed via the convolution algorithm of Chapter 4. Left-Bottom (Right-Bottom) panel: Asian call option; additional parameters: $T = 1.0$, $n = 12$, $\gamma = 1$; ref. price 9.6433 (10.9196). CPU timings computed in seconds (s).

the bias reduces by factors of 4–5 (plain vanilla put) and 2–7 (Asian call), while the CPU timing increases roughly by a factor of $10^{3/4} \approx 5.62$ as ε becomes smaller. Moreover, while the MY method exhibits smaller (absolute) bias in the case of the Asian option with quarterly sampling when $\varepsilon \in \{10^{-1}, 10^{-2}, 10^{-3}\}$, this changes abruptly in favour of the monthly sampling when $\varepsilon = 10^{-4}$. This confirms that the approximation bias is indeed hard to quantify.

Table 7.6 focuses on the outcome from the biased MY method with 10^7 simulation trials when applied to the at-the-money Asian call option. The RMSE measure (7.10) reflects both bias and variance. Based on the bias estimates in Table 7.5, we see in Table 7.6 that for $\varepsilon > 10^{-4}$ the bias dominates the standard error in the RMSE. This phenomenon becomes even more pronounced in the infinite variation model. We observe that under parameterization II the smallest RMSE achieved is closely comparable to the standard error of the corresponding exact price estimates ($K = 100$) in Table 7.4 obtained using only 10^6 trials. In the finite variation with infinite activity model the overall performance improves and we can attain similar accuracy

n	parameterization I				parameterization II			
	ε	MY	std error	RMSE	ε	MY	std error	RMSE
4	10^{-1}	10.5	0.0066	1.087	10^{-1}	11.6	0.0069	1.037
	10^{-2}	9.517	0.0061	0.078	10^{-2}	10.84	0.0062	0.22
	10^{-3}	9.437	0.0060	0.0063	10^{-3}	10.652	0.0060	0.031
	10^{-4}	9.431	0.0060	0.011	10^{-4}	10.639	0.0060	0.019
12	10^{-1}	10.8	0.0066	1.180	10^{-1}	12.0	0.0070	1.108
	10^{-2}	9.723	0.0062	0.080	10^{-2}	11.16	0.0063	0.24
	10^{-3}	9.651	0.0062	0.010	10^{-3}	10.953	0.0062	0.035
	10^{-4}	9.639	0.0062	0.0076	10^{-4}	10.928	0.0061	0.011

Table 7.6: Fixed-strike Asian call price estimates (with standard errors and RMSEs) computed using the biased Madan-Yor scheme for 10^7 trials. RMSEs computed as in (7.10) based on reported standard errors and corresponding bias estimates from Table 7.5. Model parameters: Table 7.1. Fixed parameters: $S_0 = K = 100$, $r = 0.04$, $T = 1.0$, $\gamma = 1$. Left-Top (Right-Top) panel: ref. price 9.4395 (10.6216) computed with the convolution algorithm of Chapter 4. Left-Bottom (Right-Bottom) panel: ref. price 9.6433 (10.9196).

levels (depending on ε) for the same number of trials, though the bias is not easy to control and the CPU demands grow significantly (see Table 7.5).

7.6 Concluding remarks

We have presented a Monte Carlo simulation scheme coupled with Fourier transform which is simple and fast to implement. A key feature is its ability to generate exact sample trajectories for stochastic models which admit closed-form characteristic functions. In financial applications, this leads to unbiased price estimators. We have focused on the tempered stable process, in particular the CGMY subclass, which has hitherto shown hard to simulate. Numerical examples on plain vanilla and Asian options for two model parameterizations illustrate the speed-accuracy merits of the proposed technique against existing methods in the literature. More specifically, the efficiency gains of the joint Monte Carlo-Fourier transform over the exact method of Poirrot and Tankov (2006) in pricing applications become existent for CGMY models of infinite variation for the log-returns and in-the-money path-independent options with long time to maturity (e.g., in excess of one year). Our scheme also deals with the limitations of the methods by Poirrot and Tankov (2006) and Madan and Yor (2008), i.e., can be used to price

path-dependent contracts and does not involve any approximation bias.

Appendix 7—Simulation of the CGMY process using change of measure

Assume payoff function

$$\pi_T = p(S_T) = p(S_0 e^{(r+\varpi)T+L_T}),$$

where L_T is a tempered stable random variable, $r > 0$ is the (constant) continuously compounded interest rate, and parameter ϖ is chosen such that the martingale condition $\mathbb{E}(S_T) = S_0 e^{rT}$ is satisfied under a risk-neutral measure $\mathbb{P}$. From Poirot and Tankov ((2006), Theorem 3.1), there exists an equivalent probability measure $\mathbb{Q}$ such that

$$\begin{aligned} \mathbb{E}(\pi_T) &= \mathbb{E}\left(p(S_0 e^{(r+\varpi)T+L_T})\right) \\ &= E^{\mathbb{Q}}\left(p(S_0 e^{rT+\tilde{L}_{1,T}+\tilde{L}_{2,T}})e^{-\lambda_+\tilde{L}_{1,T}+\lambda_-\tilde{L}_{2,T}-\tilde{c}T}\right), \end{aligned} \quad (7.11)$$

where constant

$$\begin{aligned} \tilde{c} &= \lambda_+ c_+ \Gamma(-\alpha_+) ((\lambda_+ - 1)^{\alpha_+} - \lambda_+^{\alpha_+}) + \lambda_+^{\alpha_+} c_+ \Gamma(-\alpha_+) \\ &\quad - \lambda_- c_- \Gamma(-\alpha_-) ((\lambda_- + 1)^{\alpha_-} - \lambda_-^{\alpha_-}) + \lambda_-^{\alpha_-} c_- \Gamma(-\alpha_-) \end{aligned}$$

for $0 < \alpha_{\pm} < 1$ or $1 < \alpha_{\pm} < 2$. The random variables $\tilde{L}_{1,T}$, $\tilde{L}_{2,T}$ follow the α_+ -stable distribution $S_{\alpha_+}(\sigma_+, 1, \mu_+)$ and α_- -stable distribution $S_{\alpha_-}(\sigma_-, -1, \mu_-)$ respectively, with parameters

$$\begin{aligned} \sigma_{\pm} &= (-c_{\pm} \Gamma(-\alpha_{\pm}) \cos(\pi\alpha_{\pm}/2)T)^{\alpha_{\pm}^{-1}}, \\ \mu_{\pm} &= -c_{\pm} \Gamma(-\alpha_{\pm}) ((\lambda_{\pm} \mp 1)^{\alpha_{\pm}} - \lambda_{\pm}^{\alpha_{\pm}})T. \end{aligned}$$

The equivalence result (7.11) enables us to transform the tempered stable variable to a sum of two stable variables whose simulation is straightforward (see Chambers et al. (1976) and Weron (1996)).

Chapter 8

A backward convolution algorithm for convertible bonds in a jump diffusion setting with stochastic interest rates

8.1 Introduction

The aim of this chapter is to introduce a Fourier transform approach for the pricing of convertible bonds (CBs) under a jump diffusion market model with correlated stochastic interest rates. In contrast with the previous literature, the proposed numerical pricing technique can accommodate a number of risk factors and contract-design features, and is shown to be efficient and accurate.

CBs are hybrid instruments which represent a pricing challenge because of their complex design. Firstly, they depend on variables related to the underlying firm value (or stock), the fixed income part, which includes both interest rates and default risk, and the interaction between these components. Secondly, CBs usually carry call options giving the issuer the right to demand premature redemption in exchange for the current call price. Put option features, which allow the investor to force the issuing firm to prematurely repurchase the CB for a

pre-specified price, are also sometimes met.

The early-exercise features that CBs present imply that the pricing problem of these contracts shares strong analogies with the one of American/Bermudan options. Closed-form solutions for the price of the CB in a Black-Scholes-Merton economy have been obtained by Ingersoll (1977a) for the case of non-callable/callable products; however, the introduction in the valuation model of a more realistic specification including, for instance, discretely payable coupons, dividends on the underlying stock, soft call provisions (which preclude the issuer from calling the CB until the firm value rises above a specified level), and a call notice period prevent the derivation of explicit pricing formulae. For these reasons, various numerical techniques have been employed in order to evaluate CBs. The literature mainly distinguishes among three types of approach: (i) numerical schemes for partial differential equations/inequalities (PDE/Is) (see Brennan and Schwartz (1977), (1980), Carayannopoulos (1996), Tsiveriotis and Fernandes (1998), Zvan et al. (1998), (2001), Takahashi et al. (2001), Barone-Adesi et al. (2003), Bermúdez and Webber (2004)), (ii) lattice methods (see Goldman Sachs (1994), Ho and Pfeffer (1996), Takahashi et al. (2001), Davis and Lischka (2002)) and (iii) Monte Carlo simulation (see Lvov et al. (2004), for an approach based on the joint simulation-regression technique by Longstaff and Schwartz (2001), and Ammann et al. (2008), for an approach based on the optimization method by García (2003)).

Contributions using the PDE/I approach rely on the finite difference (Brennan and Schwartz (1977), (1980)) and finite volume (Zvan et al. (2001)) schemes; a more recent development is the so-called joint characteristics-finite elements method suggested by Barone-Adesi et al. (2003), which aims at overcoming previously reported challenges originated by complex boundary conditions, the existence of spurious oscillations due to convection dominance, and the slow convergence. On the contrary, the popularity of lattices is frequently attributed to their intuitiveness and simplicity; lattice methods suffer, though, from an increasing number of spatial nodes at each time step, especially for long maturities. This issue becomes even more noticeable in the case of stochastic interest rates, as this requires the generation of a 2-D lattice. Furthermore, Geske and Shastri (1985) demonstrate that lattices tend to lose efficiency when dealing with discrete payments and early-exercise options.

One significant problem with the traditional PDE/I and lattice methods is the so-called

curse of dimensionality. This is about the limited number of dimensions their grids can hold effectively, and therefore the number of risk factors that the pricing model can actually include. For example, in the attempt to provide a more realistic representation of the firm's value behaviour in the CB context, Bermúdez and Webber (2004) adopt a firm value approach in which they assume the arrival of a single jump with fixed (non-random) jump size; thereafter the firm value is assumed to evolve as a pure diffusion. Although this assumption facilitates the implementation of the numerical scheme, it still remains simplistic and inadequate to the effective modelling of credit risk, as we discuss in Section 8.2.2. In this respect, Monte Carlo simulation turns out to be the preferred alternative when multiple state variables (especially more than two) and Bermudan features are considered. For example, the adaptation of Monte Carlo methods by Longstaff and Schwartz (2001) to accommodate early-exercise features is based on the approximation of the continuation value by a linear combination of suitably chosen basis functions, and the estimation of the corresponding coefficients by regression. Nevertheless, Broadie and Detemple (2004) argue that results converge slowly, demanding an increasing number of basis functions and simulation runs. In the case of CBs, additional care is required on splitting the spatial domain beforehand into regions where the CB behaves differently (likely to be called/put/continue existing), otherwise unnecessary approximation of the continuation value over the unified domain is anticipated to be poor (see Lvov et al. (2004) for a more detailed discussion of this point). Finally, Monte Carlo methods suffer a slow and non-monotone convergence, preventing the application of convergence-accelerating techniques like Richardson extrapolation.

In the light of the previous discussion, our contribution to the current state of the literature on CBs is threefold. Firstly, we propose a Fourier transform pricing technique built on martingale theory, which aims at handling effectively any real-world CB specification, including discrete cash flows, and conversion which is either forced by a call on notice from the issuer, or takes place voluntarily at the holders' choice before a dividend payment. The method belongs to the class of backward price convolutions, similar in spirit to Lord et al. (2008) described in Section 3.4; in general terms, the approach we suggest works by evaluating the convertible bond going backwards from maturity, while allowing for the early-exercise features and discrete payments at relevant time points. Secondly, we use a market model which comprises four risk

factors: an underlying evolving as a diffusion augmented by jumps, subject to random arrival and size, and stochastic interest rates. We consider both the cases of the Merton jump diffusion and the double exponential (Kou) jump diffusion for the log-increments of the underlying; these reflect the affine diffusion model (3.6-3.7) presented in Section 3.3 additionally equipped with independent random jumps. To the best of our knowledge, such a setup has not been implemented earlier in the convertible bonds' literature due to dimensionality issues. The proposed numerical pricing scheme is shown to be flexible enough to handle the dimensionality imposed by the abovementioned market model, while remaining smoothly convergent and precise. Thirdly, we show that the bivariate log-firm value-interest rate process falls within the class of affine models in the spirit of Duffie et al. (2000) and Duffie et al. (2003) (see Section 3.3) and, using the results on numéraire pair changes from Geman et al. (1995), we derive its characteristic function in closed form. Characterizing the law of the underlying model is pivotal for the implementation of the convolution algorithm.

The suggested pricing methodology is then tested using different parameterization of the adopted market framework. In particular, we examine the discrepancy between the prices generated by the two jump diffusion models under consideration as a function of the model parameter values and the moneyness (measure of the likelihood of conversion) of the convertible bond. We explore the effects of coupons payable to the CB holders and dividends distributed to the current stock holders, as well as the impact of varying call policy on the computed prices.

The remainder of this chapter is organized as follows. In Section 8.2 we introduce the basic notation and our assumptions for the firm value and interest rate processes. We then justify our choice based on the empirical evidence available on the credit-spread term structure, and provide intuition on how to overcome significant impracticalities related to the calibration of the firm value model. In Section 8.3 we describe the CB design under consideration, with particular emphasis on the optimal call strategy assumed for the issuing firm; we also derive the payoff to the CB holders after a with-notice call by the firm. In Section 8.4 we develop the theoretical ground for the Fourier transform-based backward price convolution scheme and discuss its implementation via discrete Fourier transform. Section 8.5 demonstrates the proposed numerical scheme in practice. Section 8.6 concludes the chapter.

8.2 Market model

From a valuation perspective, a pricing model for CBs requires assumptions on the term structure of interest rates, the dynamic followed by the asset underlying the conversion option, and the firm's default-driving mechanism. We adopt here a structural approach to model the underlying of the contract and the default-triggering event; in particular, we assume that the dynamic of the firm value is driven by a jump diffusion. The detailed assumptions of our model are presented in the following sections, together with the rationale of our choice. Finally, we use the Vašíček (1977) model for the term structure of interest rates.

8.2.1 The firm value-interest rate setup

Let $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ be a complete filtered probability space, where $\mathbb{P}$ is some risk-neutral probability measure. We assume that the firm value V is given by

$$V_t = e^{Y_t},$$

where Y follows the affine diffusion model (3.6-3.7) augmented by jumps, to yield the jump diffusion process

$$Y_t = Y_0 + \int_0^t (r_s - \sigma^2/2 - \lambda(\phi_L(-i) - 1)) ds + \sigma W_t + \int_{\mathbb{R}} l N_t(dl), \quad (8.1)$$

with $Y_0 = \ln V_0$, W a $\mathbb{F}$ -adapted standard Brownian motion in $\mathbb{R}$, N a time-homogeneous Poisson process with constant intensity λ and L the random jump size; L is modelled by a sequence of independent and identically distributed random variables with $\mathbb{E}(L) = \mu_L$, $\mathbb{Var}(L) = \sigma_L^2$ and characteristic function ϕ_L , while W , N and L are assumed to be mutually independent. As far as the distribution governing L is concerned, two popular choices in the literature are the double exponential distribution (Kou (2002)) and the normal distribution (Merton (1976)). Specifically, in the case of the double exponential jump diffusion process (DEJD), L has characteristic function

$$\phi_L(u) = \frac{p\eta_1}{\eta_1 - iu} + \frac{q\eta_2}{\eta_2 + iu}, \quad (8.2)$$

where $p, q \geq 0$, $\eta_1 > 1$, $\eta_2 > 0$, and $p + q = 1$ as they represent the (risk-neutral) probabilities of an upward and a downward jump respectively. In the case of the Merton jump diffusion model (MJD), instead, we assume that L follows a normal distribution; therefore, the characteristic function is

$$\phi_L(u) = e^{i\mu_L u - \sigma_L^2 u^2 / 2}. \quad (8.3)$$

The short rate process r is assumed to evolve according to the Vašíček (1977) model; hence, the log-price $\ln P_t(v)$ at $t > 0$ of a pure-discount bond maturing at $v \geq t$ satisfies

$$\ln P_t(v) = \ln P_0(v) + \int_0^t (r_s - m_s^2(v)/2) ds + \int_0^t m_s(v) dW_{r,s}; \quad (8.4)$$

$$|m_t(v)| = \frac{\sigma_r}{\kappa}(1 - e^{-\kappa(v-t)}); \quad \kappa, \sigma_r > 0, \quad (8.5)$$

where W_r is a standard Brownian motion, such that W and W_r have constant correlation ρ , whereas W_r is independent of both N and L . Alternatively,

$$\ln P_t(v) = A_t(v) - B_t(v) r_t; \quad (8.6)$$

$$A_t(v) = \frac{1}{\kappa^2} (B_t(v) - v + t) (\mu_r \kappa^2 - \sigma_r^2 / 2) - \frac{\sigma_r^2 B_t^2(v)}{4\kappa}; \quad \mu_r > 0, \quad (8.7)$$

$$B_t(v) = \frac{1}{\kappa} (1 - e^{-\kappa(v-t)}). \quad (8.8)$$

The results (8.4-8.8) can be found, for example, in Hull (2003) and Vašíček (1977).

8.2.2 Stock versus firm value and real-world considerations

Generally speaking, the available approaches to model credit risk can be classified in two main categories: the structural and intensity-based (reduced-form) models¹.

The main feature of the structural methods is the fact that the credit events are triggered by movements of the firm value below some boundary. Thus, a key aspect of this framework is the modelling of the firm value process. Structural default has been first introduced by Merton (1974), who considers the possibility of bankruptcy of a risky bond only at maturity. Various modified versions of the original Merton (1974) model have been proposed, including Black

¹Apart from the purely structural and intensity-based models, hybrid approaches combining elements from both techniques also exist. More about these can be found in Bielecki and Rutkowski (2002).

and Cox (1976), Longstaff and Schwartz (1995) (with stochastic interest rate), Leland (1994) and Leland and Toft (1996), amongst others. The main difference between these methods and Merton (1974) is the inclusion of a stopping time, which signifies the default time upon the breaching of the benchmark level by the firm value trajectory at any time over the term of the contract. In the CB context, Ingersoll (1977a), (1977b) follows Merton (1974), while Brennan and Schwartz (1977), (1980) proceed one step further by additionally allowing for default prior to maturity.

All the abovementioned contributions use a drifted Brownian motion to describe the dynamic of the log-firm value, in the spirit of Black and Scholes (1973) and Merton (1973). Nevertheless, by nature of the diffusion process, this particular assumption precludes a sudden, unexpected drop of the firm value process below the default-triggering threshold level. Consequently, the firm can go bankrupt only when its value reaches exactly that level after a smooth decline. According to Zhou (1997), for a firm which is subject to such a “predictable” default and is not in financial distress, the probability of default in the short-run is negligible, although the credit risk becomes more significant for longer maturities. Therefore, these models imply a flat term structure of credit spread at zero level for short maturities with an increasing slope at longer maturities. Unfortunately, such a shape for the credit-spread curve is inconsistent with the empirical results of Fons (1994) and others. According to these contributions, the curve for certain corporate bonds may be observed to be not only upwards-sloping, but also flat or even downwards-sloping. A possible route to face this matter of “predictability” is to include unforeseeable jumps into the dynamics of the firm value evolution, using for example Poisson jumps, as in Zhou (1997), Hilberink and Rogers (2002), Chen and Kou (2009) and Dao and Jeanblanc (2006), or by completely discarding the diffusion component and replacing it with pure jumps, as in Madan (2000). In both cases, bankruptcy takes place in the form of a jump, i.e., by crossing the critical default boundary without exactly touching it. As Zhou (1997) points out, in a jump diffusion structural approach, the diffusion component generates conceptual insights on default behaviour, since the default events can be associated to the smooth decline of the firm’s capital structure, whilst the additional existence of jumps allows for likely external impacts and enables a more flexible fitting to the observed credit spreads. Another way to produce high short-term spreads is by incorporating a stochastic barrier level, as in the CreditGrades (2002)

technical document; this feature proves to raise the likelihood for the firm's assets being at a level which is closer to the bankruptcy point than otherwise believed. In the CB context, we note that Bermúdez and Webber (2004) resort to the firm value technique by implementing a jump-augmented geometric Brownian motion, where the exogenous default event coincides with the jump time of a time-inhomogeneous Poisson counter with “semi-stochastic” intensity (see Lando (1998)).

The alternative approach to modelling default is known as the intensity-based technique. The distinguishing feature of this framework is the unpredictability of the default time, which is totally inaccessible (i.e., it comes as a “surprise”). Such a default is said to be exogenous, exactly because it occurs in a sudden manner and is related to an external cause. The concept behind intensity-based default is simple: the instantaneous probability of default is exogenously specified by means of some intensity (hazard rate), which may be treated either “semi-stochastically”, as a function of the underlying stock, or directly stochastically. In the CB context, the reduced-form technique of Duffie and Singleton (1999) is mostly popular (see Takahashi et al. (2001), Davis and Lischka (2002), Andersen and Buffum (2003), Carayannopoulos and Kalimipalli (2003)). According to this specification, pre-default prices can be reasonably assumed to be driven by a diffusion process. It can be argued though that these (semi-) stochastic intensity-based models unnecessarily penalize the default-free equity component of the convertible bond, as the default intensity appears in the drift part of the stock process. In general, a company's ability to issue stock is not strongly influenced by its credit rating and it can always deliver that stock. On the contrary, coupon and principal payments depend on the issuer's timely access to the required amounts. Inability to access these payments at the right time induces credit risk. On the same grounds, Tsiveriotis and Fernandes (1998) choose to split the CB into artificial debt-only and equity-only elements. The debt-only part is discounted at a higher rate, subject to a constant spread over the short rate, to reflect the default risk associated to it. Then, the two components are added to provide the overall CB price². On the other hand, Takahashi et al. (2001) claim that the assumption of the stock not being subject to default risk is likely to result into model inconsistency with the market. Furthermore, Takahashi et al.

²The early works of McConnell and Schwartz (1986), Cheung and Nelken (1994), and Ho and Pfeffer (1996) also consider a constant credit spread-adjusted discount rate, which is applicable, however, to the entire CB.

(2001), Davis and Lischka (2002) and Carayannopoulos and Kalimipalli (2003) presume that, upon default, the stock instantaneously jumps to zero. Based on empirical results, Ayache et al. (2003) consider this assumption as extreme and controversial. For this reason, they apply a proportional reduction to the pre-default stock price at the time of default; the optimal choice of this reduction adds to the limitations of the stock-based models.

In the light of the previous discussion, in this note we follow Bermúdez and Webber (2004) and adopt a firm value approach to credit risk in order to avoid a disputable treatment of equity. Nevertheless, we emphasize on the simple assumption of their approach that default occurs only once; thereafter the firm value is assumed to evolve as a pure diffusion, hence any possibility of future exogenous default events to occur is eliminated. This compromise is necessary by limitations of the PDI numerical scheme they employ. Here, however, we manage to overcome this modelling weakness and adopt an exponential jump diffusion firm value approach, as in Merton (1976) and Kou (2002), so that default can be reached following a number of consecutive shocks in the value of the firm. To the best of our knowledge, this is the first time that these two jump diffusion processes are utilized in the context of CBs valuation.

8.2.3 Calibration issues

Despite its appealing implications in the credit risk context, a model based on the value of the firm, which is not directly market-observable, traditionally suffers in calibration. The lack of this information poses crucial impracticalities especially in an incomplete market, such as the one proposed here, which we need to handle as efficiently as possible. As King (1986) explains, many of the firm's liabilities are not traded in organized exchanges or have limited trading activity, as opposed to the highly liquid stock, prohibiting their synchronous observation in many instances and, hence, their simultaneous estimation. The first contributions offering a solution to the estimation problem include Carayannopoulos (1996), who suggests the volatility of the common stock (obtained from the market) as proxy for the volatility of the firm value (the former actually forms an upper bound to the latter) and King (1986), who proposes a leverage-adjusted stock volatility for the firm value. As a consequence, in the case of Carayannopoulos (1996), some overpricing effects have been reported, especially for deep in-the-money CBs, due to the overstated firm value volatility, while King's version appears to work even less efficiently.

A more recent candidate for the firm's assets volatility is the one which recovers, as consistently as possible, the market CDS (Credit Default Swap) spread on that firm (e.g., CreditGrades (2002)).

Although our intention here is not to actually calibrate the firm value and interest rate processes, we brief on alternative promising calibration routes, whose application is postponed to a future stage of our research.

Starting with the short rate, calibration can be operated for the Hull and White interest rate model (which generalizes the original Vašíček model): following Barone-Adesi et al. (2003), the asymptotic mean interest rate level, μ_r , can be inferred from the market prices of zero-coupon bonds, as at a given reference date, while the remaining parameters, κ and σ_r , are obtained from the minimization of the root mean square error between the theoretical (closed-form formula) and market prices of actively traded interest rate options, e.g., caps, as at the reference date.

As far as the firm value is concerned, the nonparametric calibration methodology by Cont and Tankov (2004b) for Lévy processes guarantees consistency with the observed market prices via minimization of a well-defined model-market prices distance functional. Further, it allows for dependence on the information gained since the previous calibration, via an entropic measure of the closeness between the current market martingale measure and the prior measure (the outcome from the previous calibration). Cont and Tankov (2004b) test successfully the performance of their algorithm on European plain vanilla options on stocks in a DEJD setup. An extensive discussion on this procedure and the associated technicalities can be found in Cont and Tankov (2004a), (2004b). Therefore, instead of seeking to infer the unknown firm parameter values from stock data, we may use the information contained in the historical prices of ordinary and convertible bonds from the same issuer. However, because of our model's structural nature, a complication arises due to the need to infer simultaneously all claims. To eliminate this complication, we should ideally restrict our sample to firms with simple capital structures consisting only of common stock, senior debt and subordinated convertible debt. Then, based on the derived guesses for the parameters, we could test the out-of-sample forecasting power of our CB model. Zabolotnyuk et al. (2009) have set up the structural model of Brennan and Schwartz (1977), (1980), where riskless senior debt has been easily incorporated to be deduced as part of the calibration procedure. In this way, they have managed to calibrate effectively

and produce price forecasts, which are comparable to the Tsiveriotis and Fernandes (1998) stock-based model predictions for the same sample of firms.

8.3 Convertible bonds: contract features

A convertible bond is an ordinary bond which additionally offers the investors the option to exchange it for a predetermined number of shares at certain points in time. In this respect, the conversion rights originate a Bermudan option. In the case of conversion, each investor receives the conversion value γV_t , where $\gamma \in (0, 1)$ denotes the dilution factor, i.e., the fraction of common stock possessed by each CB holder post-conversion. The CB issue usually offers regular aggregate coupon payments C_{t_j} at time $t_j \in [0, T]$ and, for m outstanding CBs, this corresponds to $c_{t_j} = C_{t_j}/m$ payment per bond. In the case the issue is kept alive to its expiration at time T , it is redeemed for a total face value mF . The firm's stock holders receive, instead, a discrete aggregate dividend D_{t_i} at the dividend date $t_i \in [0, T]$, such that $t_j \neq t_i$.

Furthermore, CBs contain a call option allowing the issuer to redeem it prematurely in exchange for the current call price; the issuer is in general obliged to announce his/her decision to call the bond a certain period in advance (call notice period). Once the CB is called, the investor needs to consider if it is the case to exercise the conversion option at the end of the call notice period, in order to convert instead of receiving the call price. Put option provisions entitling the investor to force a premature repurchase of the CBs by the issuing firm, are another feature which is sometimes met. We currently ignore the putability provision, as it has been shown to cause minor effects on the CB price (see Bermúdez and Webber (2004)).

The existence of a callability provision implies that the CB payoff depends on the optimal exercise strategy adopted by the issuer. This is discussed in the next section.

8.3.1 The optimal call strategy

Under the assumption of a market not subject to any imperfections, in which the Modigliani-Miller theorem holds and no call notice applies, Ingersoll (1977a) proves that the optimal call policy for a callable convertible issue is to call as soon as the firm value V_t reaches the critical level K_t/γ , for a deterministic call price K_t which usually is either fixed by the firm at the

issue of the contract, or a piecewise constant function (see Ammann et al. (2008)). This feature endows the CB with path-dependence and, consequently, implies the need for frequent monitoring. Despite his original result, Ingersoll (1977b) observes empirically that firms tend to follow different call strategies; they choose, in fact, to call when the conversion value is in excess of the call price. Forcing conversion by a call at the earliest opportunity, instead, leads to undervalued CBs, as shown in Carayannopoulos (1996) and Carayannopoulos and Kalimipalli (2003).

The extensive empirical analysis carried out by Asquith and Mullins (1991) and Asquith (1995) shows that the observed call delays can be attributed mainly to three factors: a call notice period, the existence of significant cash flows advantages, and a safety premium on the given call price. In details, the call notice period, which prohibits the CB to be called for as long as this period is active, in fact proves to be the main reason for the delayed calls; for those CBs that are not called at the end of this period, the firm might be saving cash by delaying the call if, for example, the after-corporate tax coupons on the CB are less than the dividends payable post-conversion. Another important reason for delaying is linked to the existence of a safety premium imposed by the issuing firm prior to the call announcement, in the attempt to guarantee that the conversion value will still exceed the call price at the end of the call notice period and, hence, avoid the bond redemption in cash.

In this work, we build on these findings and formulate the optimal call policy for the CB as follows. Let $\vartheta \in (0, 1)$ denote the safety premium mentioned above; then, the firm's optimal call announcement is given by the stopping time

$$\tau^c := \inf \left\{ t : V_t \geq \frac{(1 + \vartheta) K_t}{\gamma} \right\}.$$

Assume the call notice period is s^c and define the accrued interest $AccIR = \frac{\tau^c + s^c - t_j}{t_{j+1} - t_j} c_{t_{j+1}}$, $t_j \leq \tau^c + s^c < t_{j+1}$, such that the call price at the end of this period is $K_{\tau^c + s^c} = K_{\tau^c} + AccIR$. Then, the investor's payoff upon the call of the CB by the issuer is $\bar{K}_{\tau^c + s^c}(V_{\tau^c + s^c}) = \max(\gamma V_{\tau^c + s^c}, K_{\tau^c + s^c})$, and its no-arbitrage price at the time of the call is

$$\tilde{K}_{\tau^c}(V_{\tau^c}, r_{\tau^c}) = \mathbb{E} \left(e^{-\int_{\tau^c}^{\tau^c + s^c} r_s ds} \bar{K}_{\tau^c + s^c}(V_{\tau^c + s^c}) \middle| \mathcal{F}_{\tau^c} \right). \quad (8.9)$$

8.3.2 The payoff function and pricing considerations

Because of the early-exercise rights embedded in the CB, we define the contract payoff (per bond) function $\tilde{H}_t$ at any possible decision time $t \in (0, T]$ as follows.

At maturity T , the investors can choose between converting to common stock (see Brennan and Schwartz (1977), Lemma 1) and receiving the face value and the last coupon, providing that the firm can afford the total of this payment. Otherwise, they recover the outstanding firm value at that time. Hence,

$$\tilde{H}_T(V_T, r_T) = \begin{cases} \gamma V_T, & V_T \geq (F + c_T) / \gamma \\ F + c_T, & mF + C_T \leq V_T < (F + c_T) / \gamma \\ V_T / m, & V_T < mF + C_T. \end{cases} \quad (8.10)$$

At a date where neither coupon nor dividend payments are due, the CB may be forced by a call to conversion, or continue to exist at least until the next monitoring point, i.e.,

$$\tilde{H}_t(V_t, r_t) = \begin{cases} \tilde{K}_t(V_t, r_t), & V_t \geq \frac{(1+\vartheta)K_t}{\gamma}, \quad 0 < t < T, \quad t \neq t_i, t_j \\ H_t(V_t, r_t), & V_t < \frac{(1+\vartheta)K_t}{\gamma}, \quad 0 < t < T, \quad t \neq t_i, t_j, \end{cases} \quad (8.11)$$

where H_t denotes the no-arbitrage (continuation) value of the CB, and $\tilde{K}_t$ is given by equation (8.9).

At a coupon date, t_j , the payoff of the CB depends on whether the firm has enough funding to meet the claim. If $V_{t_j-} \leq C_{t_j}$, the CB defaults, its value is $H_{t_j} = 0$, since $0 \leq H_{t_j} \leq V_{t_j}$ by limited liability and the Modigliani-Miller theorem, and $C_{t_j} = V_{t_j-}$, i.e., the investor sizes the available assets. If, instead, $V_{t_j-} > C_{t_j}$ and for as long as the CB is uncalled, the contract remains in force and the coupon is paid in full. On the other hand, if the CB is called, its holders receive both the call payoff and the coupon. Hence,

$$\tilde{H}_{t-}(V_{t-}, r_t) = \begin{cases} V_{t-} / m, & V_{t-} \leq C_{t_j}, & 0 < t < T, \quad t = t_j \\ H_t(V_t, r_t) + c_t, & C_t < V_{t-} < \frac{(1+\vartheta)K_t}{\gamma}, & 0 < t < T, \quad t = t_j \\ \tilde{K}_t(V_t, r_t) + c_t, & V_{t-} \geq \frac{(1+\vartheta)K_t}{\gamma}, & 0 < t < T, \quad t = t_j. \end{cases} \quad (8.12)$$

Finally, at a dividend date, t_i , the investors may find optimal to convert prior to the dividend

payment³ (voluntary conversion). The following condition, which is proved in Brennan and Schwartz ((1977), Lemma 1), applies

$$\tilde{H}_{t-}(V_{t-}, r_t) = \max(H_t(V_t, r_t), \gamma V_{t-}), \quad 0 < t < T, \quad t = t_i. \quad (8.13)$$

The payoff function defined by equations (8.10-8.13) highlights the Bermudan style and high path-dependency of the CB; these features imply that the no-arbitrage price of the CB, H_0 , can only be recovered by numerical approximation.

8.4 The backward price convolution algorithm

Lord et al. (2008) utilize a backward recursive integration scheme to produce accurate prices for Bermudan vanilla options. Their method relies on the property of independent increments shown by the log-returns in their market model. We adapt this approach to the pricing of CBs; however, the straightforward extension of the method is not possible due to the fact that in our model the increments of the log-firm value are not independent (see equation (8.1)). Moreover, the contract under consideration presents a higher degree of complexity due to the presence of intermediate discrete payments, exotic features, like call provision with attached call notice, and additional risk factors, such as stochastic interest rates.

We consider the partition $\mathcal{T} = \{t_k\}_{k=0}^n$, $n \in \mathbb{N}^*$, of the contract's term $[0, T]$ signifying the set of the decision dates. For ease of exposition we assume that these dates are equally spaced so that $t_k - t_{k-1} = \delta t$ for $0 < k \leq n$, with $t_0 = 0$, $t_n = T$. With these assumptions in mind, the price of the CB is the solution to the dynamic programming problem described next.

³At a dividend date, the existing stock holders are entitled to receive dividends for as long as the firm can afford their payment, providing that it has already met all the other claims ranking above them.

We define functions g , g_r as follows

$$g_{k-1-}(y, y_r) = \begin{cases} \ln(e^y - D_{t_{k-1}}) - A_{t_{k-1}}(t_k) + B_{t_{k-1}}(t_k) y_r, & 1 \leq k \leq n, \ k-1 = i \\ & y > \ln D_{t_{k-1}} \\ \ln(e^y - C_{t_{k-1}}) - A_{t_{k-1}}(t_k) + B_{t_{k-1}}(t_k) y_r, & 1 \leq k \leq n, \ k-1 = j \\ & y > \ln C_{t_{k-1}} \\ y - A_{t_{k-1}}(t_k) + B_{t_{k-1}}(t_k) y_r, & 1 \leq k \leq n, \ k-1 \neq i, j, \end{cases}$$

$$g_{r,k-1}(y_r) = y_r e^{-\kappa(t_k - t_{k-1})}, \quad 1 \leq k \leq n,$$

where $A_{t_{k-1}}(t_k)$ and $B_{t_{k-1}}(t_k)$ are given by equations (8.7) and (8.8) respectively. We further define the pairs

$$(Z_k, Z_{r,k}) = (Y_{t_{k-}} - g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}), r_{t_k} - g_{r,k-1}(r_{t_{k-1}})), \quad 1 \leq k \leq n, \quad (8.14)$$

and denote by f^* their joint $\mathbb{P}^*$ -density function for all k .⁴

Based on the fundamental theorem of asset pricing, we write in filtration $\mathbb{F}$ for $1 \leq k \leq n$ the iteration

$$H_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}) = \mathbb{E} \left(e^{-\int_{t_{k-1}}^{t_k} r_s ds} \tilde{H}_{k-}(Y_{t_{k-}}, r_{t_k}) \middle| \mathcal{F}_{t_{k-1-}} \right) = P_{t_{k-1}}(t_k) \mathbb{E}^* \left(\tilde{H}_{k-}(Y_{t_{k-}}, r_{t_k}) \middle| \mathcal{F}_{t_{k-1-}} \right), \quad (8.15)$$

where the second equality follows by a change to the t_k -forward measure $\mathbb{P}^*$, induced by taking as numéraire the price $P_\tau(t_k) = \exp(A_\tau(t_k) - B_\tau(t_k) r_\tau)$ (see equation (8.6)) of a pure-discount bond maturing at t_k as at time $t_{k-1} \leq \tau \leq t_k$ (see Appendix 8.A).⁵ From (8.14),

$$\begin{aligned} \mathbb{E}^* \left(\tilde{H}_{k-}(Y_{t_{k-}}, r_{t_k}) \middle| \mathcal{F}_{t_{k-1-}} \right) &= \mathbb{E}^* \left(\tilde{H}_{k-}(g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}) + Z, g_{r,k-1}(r_{t_{k-1}}) + Z_r) \middle| \mathcal{F}_{t_{k-1-}} \right) \\ &= \int_{\mathbb{R} \times \mathbb{R}} \tilde{H}_{k-}(g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}) + z, g_{r,k-1}(r_{t_{k-1}}) + z_r) f^*(z, z_r) d(z, z_r). \end{aligned}$$

⁴As shown in Appendix 8.B, equation (8.32), the pair (Z, Z_r) forms a sequence of identically distributed random variables; hence, we may drop the time-subscripts from $(Z_{k-}, Z_{r,k-})$.

⁵For completeness, impose time-subscripts $t_{0-} = t_0$, $t_{n-} = t_n$.

We define

$$\begin{aligned} & \hat{H}_{k-1-}(g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}), g_{r,k-1}(r_{t_{k-1}})) \\ &= \int_{\mathbb{R} \times \mathbb{R}} \tilde{H}_{k-}(g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}) + z, g_{r,k-1}(r_{t_{k-1}}) + z_r) f^*(z, z_r) d(z, z_r), \end{aligned} \quad (8.16)$$

such that, from (8.15), the no-arbitrage price of the CB at t_{k-1} is

$$H_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}) = P_{t_{k-1}}(t_k) \hat{H}_{k-1-}(g_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}}), g_{r,k-1}(r_{t_{k-1}})). \quad (8.17)$$

Result (8.17) is then used to compute the new payoff $\tilde{H}_{k-1-}(Y_{t_{k-1-}}, r_{t_{k-1}})$, in accordance with (8.10-8.13), to be applied in the subsequent iteration. Ultimately, the price of the CB at inception is $H_0(Y_{t_0}, r_{t_0})$.

Moreover, since (8.16) forms a convolution, we can express the Fourier transform of $\hat{H}_{k-1-}$ as

$$\mathcal{F}(\hat{H}_{k-1-}) = \mathcal{F}(\tilde{H}_{k-} * f^*(-z, -z_r)) = \mathcal{F}(\tilde{H}_{k-}) \mathcal{F}(f^*(-z, -z_r)) = \mathcal{F}(\tilde{H}_{k-}) \varphi^*,$$

where $\mathcal{F}$ denotes the Fourier transform (see equation (2.3)), and φ^* is the complex conjugate of the characteristic function ϕ^* of the pair (Z, Z_r)

$$\varphi^*(u, u_r) = \phi^*(-u, -u_r) = \mathbb{E}^*(e^{-iuZ - iu_r Z_r}) = \int_{\mathbb{R} \times \mathbb{R}} e^{-iuz - iu_r z_r} f^*(z, z_r) d(z, z_r). \quad (8.18)$$

ϕ^* is given by the closed analytical form (8.32) derived in Appendix 8.B. By Fourier inversion, we recover

$$\hat{H}_{k-1-} = \mathcal{F}^{-1}(\mathcal{F}(\hat{H}_{k-1-})) = \mathcal{F}^{-1}(\mathcal{F}(\tilde{H}_{k-}) \varphi^*).$$

8.4.1 The call payoff

Upon a call of the CB by the issuing firm, it is required that we compute the payoff to the CB holders (see equations (8.11), (8.12)). From (8.9), the call payoff as at the call announcement date is given by

$$\tilde{K}_{\tau^c}(Y_{\tau^c}, r_{\tau^c}) = \mathbb{E} \left(e^{-\int_{\tau^c}^{\tau^c+s^c} r_s ds} \bar{K}_{\tau^c+s^c}(Y_{\tau^c+s^c}) \middle| \mathcal{F}_{\tau^c} \right) = P_{\tau^c}(\tau^c + s^c) \mathbb{E}^* \left(\bar{K}_{\tau^c+s^c}(Y_{\tau^c+s^c}) \middle| \mathcal{F}_{\tau^c} \right),$$

following a change to the $\mathbb{P}^*$ measure.

We define the function

$$h_{\tau^c}(y, y_r) = \begin{cases} y - A_{\tau^c}(\tau^c + s^c) + B_{\tau^c}(\tau^c + s^c)y_r, & \tau^c \neq t_j \\ \ln(e^y - C_{\tau^c}) - A_{\tau^c}(\tau^c + s^c) + B_{\tau^c}(\tau^c + s^c)y_r, & \tau^c = t_{j-}, \end{cases}$$

and the random variable

$$Z_{\tau^c+s^c} = Y_{\tau^c+s^c} - h_{\tau^c}(Y_{\tau^c}, r_{\tau^c}) \quad (8.19)$$

for which we assume marginal $\mathbb{P}^*$ -density function f_Z^* .⁶ Then, from (8.19),

$$\begin{aligned} \mathbb{E}^*(\bar{K}_{\tau^c+s^c}(Y_{\tau^c+s^c}) | \mathcal{F}_{\tau^c}) &= \mathbb{E}^*(\bar{K}_{\tau^c+s^c}(h_{\tau^c}(Y_{\tau^c}, r_{\tau^c}) + Z_{\tau^c+s^c}) | \mathcal{F}_{\tau^c}) \\ &= \int_{\mathbb{R}} \bar{K}_{\tau^c+s^c}(h_{\tau^c}(Y_{\tau^c}, r_{\tau^c}) + z) f_Z^*(z) dz = \hat{K}_{\tau^c}(h_{\tau^c}(Y_{\tau^c}, r_{\tau^c})). \end{aligned}$$

Based on Theorem 5, we write

$$\mathcal{F}(\hat{K}_{\tau^c}) = \mathcal{F}(\bar{K}_{\tau^c+s^c} * f_Z^*(-z)) = \mathcal{F}(\bar{K}_{\tau^c+s^c})\mathcal{F}(f_Z^*(-z)) = \mathcal{F}(\bar{K}_{\tau^c+s^c})\varphi_Z^*,$$

where $\varphi_Z^* = \mathbb{E}^*(e^{-iuZ}) = \varphi^*(u, 0)$ (see equation (8.18)). Subsequently, we recover

$$\hat{K}_{\tau^c} = \mathcal{F}^{-1}(\mathcal{F}(\hat{K}_{\tau^c})) = \mathcal{F}^{-1}(\mathcal{F}(\bar{K}_{\tau^c+s^c})\varphi_Z^*), \quad (8.20)$$

and compute

$$\tilde{K}_{\tau^c} = P_{\tau^c}(\tau^c + s^c) \hat{K}_{\tau^c}. \quad (8.21)$$

8.4.2 Numerical implementation

Preliminaries. To evaluate numerically the price functions $\tilde{H}_{k-}$ and $\hat{H}_{k-1-}$ we select uniform grids $\mathbf{y}$, $\mathbf{x}$ with N grid points and spacings $\delta y = \delta x$ (log-firm value dimension), and $\mathbf{y}_r$, $\mathbf{x}_r$ with N_r grid points and spacings $\delta y_r = \delta x_r$ (short rate dimension). Moreover, we determine uniform, symmetric about zero, grids $\mathbf{u}$, $\mathbf{u}_r$ with N and N_r grid points and spacings δu and δu_r respectively. We evaluate the joint characteristic function φ^* on $\mathbf{u}^* = (\mathbf{u}, \mathbf{u}_r)$ and denote the

⁶The density f_Z^* is function only of the time length s^c , and not the actual call announcement time τ^c (see equation (8.32), Appendix 8.B); hence, any time index can be removed.

grid function values by φ^* . Since φ^* does not depend on the state variable values, it can be calculated and stored outside the recursive part of the algorithm for later use.

Recursive part. Starting from the contract maturity time and moving backwards along the time axis split into n subintervals, we summarize next the three repeating steps which characterize the operations undertaken in each subinterval.

1. **Swapping between the state and Fourier spaces.** Assume the values approximating $\tilde{H}_{k-}$ at the k^{th} decision date are given on the grid $\mathbf{y}^* = (\mathbf{y}, \mathbf{y}_r)$, and denote these by $\tilde{\mathbf{H}}_{k-}$. Evaluate then the discrete approximation to the Fourier transform $\mathcal{F}(\tilde{H}_{k-})$ on $\mathbf{u}^*$ as

$$\mathcal{D}(\tilde{\mathbf{H}}_{k-} \cdot \mathbf{w}, \mathbf{y}^*, \mathbf{u}^*) \delta y \delta y_r,$$

using the conversion (2.27). We denote by $\cdot$ the element-wise matrix multiplication. Define the trapezoidal weights

$$\begin{aligned} \mathbf{w}_{l,l_r} = & 1 - (\delta_l + \delta_{l_r} + \delta_{N-1-l} + \delta_{N_r-1-l_r})/2 + (\delta_l \delta_{l_r} + \delta_l \delta_{N_r-1-l_r} \\ & \delta_{N-1-l} \delta_{l_r} + \delta_{N-1-l} \delta_{N_r-1-l_r})/4, \end{aligned}$$

where $l = 0, \dots, N-1$, $l_r = 0, \dots, N_r-1$ and the Kronecker delta δ_ϵ takes value 1 (0) for $\epsilon = 0$ ($\epsilon \neq 0$). The inverse transform $\hat{H}_{k-1-} = \mathcal{F}^{-1}(\mathcal{F}(\tilde{H}_{k-})\varphi^*)$ is provided in terms of the discrete approximation

$$\hat{\mathbf{H}}_{k-1-} = \frac{1}{(2\pi)^2} \mathcal{D}(\mathcal{D}(\tilde{\mathbf{H}}_{k-} \cdot \mathbf{w}, \mathbf{y}^*, \mathbf{u}^*) \delta y \delta y_r \cdot \varphi^*, -\mathbf{u}^*, \mathbf{x}^*) \delta u \delta u_r$$

on the grid $\mathbf{x}^* = (\mathbf{x}, \mathbf{x}_r)$, where the outer IDFT is operated according to the conversion (2.28).

2. **From $\hat{H}_{k-1-}$ to H_{k-1-} .** Calculate

$$\mathbf{H}_{k-1-} = \exp(A_{t_{k-1}}(t_k) - B_{t_{k-1}}(t_k) \mathbf{Y}_r) \cdot \hat{H}_{k-1-}(g_{k-1-}(\mathbf{y}^*), g_{r,k-1}(\mathbf{y}_r)),$$

where $\mathbf{Y}_r$ has l^{th} row $\mathbf{Y}_{r,l} = \mathbf{y}_r$ for $l = 0, \dots, N-1$. We approximate $\hat{H}_{k-1-}$ at $g_{k-1-}(\mathbf{y}^*) \subseteq \mathbf{x}$, $g_{r,k-1}(\mathbf{y}_r) \subseteq \mathbf{x}_r$ by fitting a surface to the nodes $(\mathbf{x}^*, \hat{\mathbf{H}}_{k-1-})$ using

cubic interpolation, whereas for $g_{k-1-}(\mathbf{y}^*) \not\leq \mathbf{x}$ we extrapolate linearly in e^x .

3. **From H_{k-1-} to $\tilde{H}_{k-1-}$.** Given $\mathbf{H}_{k-1-}$, we determine $\tilde{\mathbf{H}}_{k-1-}$ according to the applicable payoff function (8.10-8.13) at the decision date t_{k-1} .

The call payoff. Upon the call announcement date, the approximate values for $\tilde{K}_{\tau^c}$, as by (8.20-8.21), are obtained via implementation of the steps 1-3, by making, instead, use of one-dimensional transforms. Furthermore, in step 3, $\tilde{K}_{\tau^c}$ is computed at $h_{\tau^c}(\mathbf{y}^*)$ only twice (for $\tau^c \neq t_{j-}$ and $\tau^c = t_{j-}$) for any number of time steps, and stored for later use in the recursive part of the scheme.

Ultimately, the numerical procedure outlined above provides us with the CB values at the inception of the contract on the $N \times N_r$ grid of initial firm values and short rate values.

8.5 Numerical study

8.5.1 Black-Scholes-Merton model

In order to illustrate the performance of our algorithm, we choose as our benchmark the prices obtained from the exact analytical formula of Ingersoll (1977a) for continuously callable CBs in the Black-Scholes-Merton economy under the assumption of constant interest rates. For testing purposes, we focus on CBs which mature in 5 years (typical) and 2 years. In Table 8.1, we present our results for finite and infinite (continuous) observation, alongside the prices obtained from the closed-form solution. Table 8.1 reports both the cases of constant and time-dependent call prices.

Under constant interest rates, we can achieve results, subject to finite monitoring, which are precise to 7 decimal places. Furthermore, precision to 4 decimal places is attainable in 1.1 seconds when $T = 2$, $n = 500$ (i.e., daily sampling), and in 2.2 seconds when $T = 5$, $n = 1250$.⁷ By regular convergence of our scheme, we can additionally extrapolate the discretely monitored CB prices to approximate the price of an otherwise equivalent, continuously monitored CB (infinite sampling). We then generate results subject to less than 0.0005% error. For given

⁷Hereafter, all CPU times reported are for MATLAB R2007b on an Intel Core 2 Duo processor T5500 1.66GHz with 2.0GB of RAM.

K_t	T	n	Backward convolution: constant interest rates	Ingersoll: constant interest rates	% error
40	2	250	36.9477708	36.9522338	0.00017
		500	36.9490640		
		∞	36.95217		
	5	500	33.1045746	33.1444004	0.00018
		1250	33.1191553		
		∞	33.14434		
$40e^{-0.02(T-t)}$	2	250	36.9306700	36.9331542	0.00012
		500	36.9313884		
		∞	36.93311		
	5	500	32.8570091	32.8813609	0.00046
		1250	32.8659022		
		∞	32.88121		

Table 8.1: Callable CB prices in the Black-Scholes-Merton model. Callable CB specification: $F = 40$, $C = D = 0$, $m = 1$, $\gamma = 0.2$, $\vartheta = 0$, $s^c = 0$. Firm value parameters: $V_0 = 100$, $\sigma = 0.25$. Constant interest rate: $r = 0.04$. Error expressed as a percentage of the exact price obtained using the result by Ingersoll (1977a).

time to maturity, this precision can be improved if we raise the sampling frequency. We reach the same conclusions on the assumptions of constant call price and call price as a function of time. The case with stochastic interest rates is investigated in Section 8.5.3.

Here, our recursion proves competitive with standard numerical techniques, given also the number of risk factors they can flexibly accommodate. In particular, Ammann et al. (2008) simulate callable CB prices ($T = 2$, daily sampling) in a two-factor setting with stochastic interest rates. Monte Carlo price estimates are reported up to the second decimal place, subject to standard error of order 10^{-1} (stochastic interest rates) and 10^{-2} (constant interest rates). Standard errors of variable order 10^{-1} - 10^{-2} are also common in the two-factor joint simulation-regression application in Lvov et al. (2004) (subject to 16 exercise times per year). In the PDE/I context with two-factors, Zvan et al. (2001) obtain monotone convergence in the number of grid and time points, which they attribute, nevertheless, to the conversion and call boundary conditions forcing the solution to be closely linear over large parts of the spatial domain. They report callable CB prices ($T = 10$, $n = 320$) with precision up to three decimal places. Barone-Adesi et al. (2003) and Bermúdez and Webber (2004) employ a joint characteristics-finite elements scheme to price callable CBs ($T = 5$, $n = 400$), which converges at first order in the

number of grid and time steps. Although they report accuracy up to 3 decimal places, their PDI method suffers from increasing dimensionality when random jumps are included into the firm value dynamics. To maintain the 2-D structure of their PDI, they resort to the simplifying assumption of a single jump of fixed size.

8.5.2 Jump diffusion setup

In this section, we examine the impact of including jumps in the original Gaussian log-return diffusion and how variations of the jump intensity λ , mean μ_L and variance σ_L of the jump size affect the callable CB prices. To this end, we ignore for convenience and without loss of generality the case of stochastic interest rates, due to their independence from the jump component, and calibrate the DEJD and MJD risk-neutral models to match mean and variance of the log-return distribution as well as λ , μ_L and σ_L . The base values for these quantities are consistent with the assumptions of Dao and Jeanblanc (2006). The exact moments of the log-return distribution $I_t = \ln(V_t/V_0)$ are derived by differentiating the cumulant generating function and evaluating at zero, as indicated by equation (2.5). On the assumption of constant interest rates, the risk-neutral cumulant generating function of I_t is

$$\psi_I(u)t = (i(r - \chi_I(-i))u + \chi_I(u))t$$

with $\chi_I(u) = -\sigma^2 u^2/2 + \lambda(\phi_L(u) - 1)$ and $\phi_L(u)$ given by (8.2) and (8.3) for the DEJD and MJD models respectively. The fitted parameters and moments, including the resulting skewness coefficient and excess kurtosis, are summarized in Table 8.2.

The accuracy of the convolution algorithm has already been explored in the Lévy and non-Lévy with stochastic volatility context in Chapters 4 and 6, in pricing discretely sampled Asian options. For the MJD and DEJD setups considered here, the numerical method shows similar robustness across different levels of moneyness of the convertible bond and model parameter values.

In Table 8.3, we study the average price deviation between the two paradigms, as a function of the parameter values λ , μ_L , σ_L and the moneyness of the CB. Moneyness is calculated as the ratio between the conversion and investment values, where the latter is defined as the

case	MJD & DEJD					MJD		DEJD				
	λ	μ_L	σ_L	$\mathbb{E}(I_1)$	$\mathbb{V}\text{ar}(I_1)$	$s(I_1)$	$\kappa(I_1)$	η_1	η_2	p	$s(I_1)$	$\kappa(I_1)$
base	3	-0.0150	0.0357	0.0178	0.2110	-0.0194	0.0101	50	33.3333	0.3	-0.0316	0.0224
λ_I	5	-0.0150	0.0357	0.0163	0.2179	-0.0293	0.0147	50	33.3333	0.3	-0.0478	0.0327
λ_{II}	1	-0.0150	0.0357	0.0193	0.2037	-0.0072	0.0039	50	33.3333	0.3	-0.0117	0.0086
$\mu_{L,I}$	3	-0.0300	0.0357	0.0168	0.2157	-0.0424	0.0174	13.4374	31.4533	0.0169	-0.0442	0.0507
$\mu_{L,II}$	3	-0.0075	0.0357	0.0180	0.2097	-0.0095	0.0082	38.8943	38.6981	0.3558	-0.0099	0.0165
$\sigma_{L,I}$	3	-0.0150	0.0714	0.0121	0.2366	-0.0527	0.0813	19.6218	19.2543	0.3590	-0.0575	0.1629
$\sigma_{L,II}$	3	-0.0150	0.0179	0.0192	0.2041	-0.0063	0.0014	26.0657	62.9170	0.0165	-0.0064	0.0041

Table 8.2: Calibrated model parameters. Parameters $r = 0.04$, $\sigma = 0.2$ remain fixed in all cases. Assume process $I_t = \ln(V_t/V_0)$ with mean $\mathbb{E}(I_t)$, variance $\mathbb{V}\text{ar}(I_t)$, skewness coefficient $s(I_t)$, and excess kurtosis $\kappa(I_t)$. These quantities are calculated via differentiation of the cumulant generating functions, as explained in the text.

hypothetical bond value in the absence of the conversion option and the credit risk.

Several comments are in order. In all cases, the MJD model prices are in excess of the prices generated by the DEJD model. This is due to the constantly stronger negative skewness and leptokurtosis of the DEJD distribution, which together guarantee higher and lower likelihoods of default and call respectively. The reduction in the value caused by the default effect is strong enough to overshadow the raise in the CB value caused by the call effect. The marked asymmetry of the DEJD distribution and the corresponding excess kurtosis are due to the fact that the event of a downward jump is more likely under every parameter combination considered here. In fact, as a result of imposing the same rate of arrival λ , the same mean and variance for both the jump size L and the log-return I across the two models, the parameter p , i.e., the probability assigned to an upward jump, is always less than 0.5 regardless of the mean size of the up/downward jump (controlled by η_1 , η_2 respectively). The observed skewness and excess kurtosis also explain the higher prices generated by the Gaussian model, as this underestimates the probability of default.

Further, Table 8.2 shows that, in the λ_{II} , $\mu_{L,II}$, $\sigma_{L,II}$ cases, the effect of the jump component is negligible; Table 8.3 confirms in fact that the prices originated by the two jump diffusion processes and the Gaussian model coincide to penny accuracy. On the contrary, in the λ_I , $\mu_{L,I}$, $\sigma_{L,I}$ cases, when the presence of the jump part is more significant, the MJD versus the DEJD price deviation reaches up to five pence, whilst the non-leptokurtic versus leptokurtic deviation

case	MJD – DEJD			non-leptokurtic – leptokurtic		
	moneyness			moneyness		
	0.6-1	1-1.2	1.2-1.4	0.6-1	1-1.2	1.2-1.4
base	0.006	0.004	0.001	0.016	0.012	0.006
λ_I	0.009	0.006	0.002	0.025	0.017	0.007
λ_{II}	0.002	0.001	0.001	0.006	0.006	0.004
$\mu_{L,I}$	0.010	0.021	0.012	0.028	0.023	0.020
$\mu_{L,II}$	0.003	0.006	0.008	0.010	0.016	0.019
$\sigma_{L,I}$	0.023	0.050	0.026	0.101	0.150	0.090
$\sigma_{L,II}$	0.001	0.003	0.004	0.005	0.005	0.004

Table 8.3: MJD versus DEJD, non-leptokurtic (Gaussian) versus leptokurtic distribution. Estimated average price difference as function of λ , μ_L , σ_L and CB moneyness. Benchmark for the leptokurtic case: DEJD model with stronger departure from the Gaussian case. Callable CB specification: $T = 5$, $n = 1250$, $F = 40$, $K = 50$, $C = D = 0$, $m = 1$, $\gamma = 0.2$, $\vartheta = 0$, $s^c = 0$. Prices (accurate to 5 decimal places) computed for 17,150 equidistant values $\ln V_0$ in $[\ln 100, \ln(K/\gamma)]$. Moneyness ranges from 0.6 to 1.5. Price differences obtained and averaged piecewise for moneyness regions $[0.6, 1)$, $[1, 1.2)$, $[1.2, 1.4)$, $[1.4, 1.5]$. For deep in-the-money CBs (top moneyness slice), the average price difference tends practically to zero level due to the firm's highly likely call (excluded from the table).

can be up to 15 pence. Changes in σ_L appear to have the most noticeable impact on the price discrepancy among the three parameter-type modifications we consider here, due to the higher impact that this parameter has on the overall skewness and excess kurtosis of the log-returns.

Moreover, in all cases, for deep in-the-money CBs, all the models' prices converge to the call price since the CB is then forced-by-call converted. For μ_L which is well below zero, the MJD versus DEJD price difference is observed to peak from an early stage, when the CB is close to the money, while for μ_L closer to zero, the peak delays until the CB is in the money. This behaviour is attributed to the different level of skewness and excess kurtosis originated by the two different combinations of parameters associated to $\mu_{L,I}$ and $\mu_{L,II}$, and, therefore, the different impact of the default and call effects, as previously discussed. Similar pattern, with higher-level peak though, is spotted for the two cases of σ_L considered here.

8.5.3 Effects of discrete coupon and dividend payments

We explore the consequences of adding discrete coupons and dividends into the valuation framework. In Tables 8.4 and 8.5, we report prices for 5-year callable CBs on a daily sampling basis ($n = 1250$), subject to both constant and stochastic interest rates. In the case of stochastic

interest rate, we select $r_0 = \mu_r = 0.04$, $\rho = 0.2$ and $\kappa = 0.858$, $\sigma_r = 0.047$, as estimated by Aït-Sahalia (1996) for the Vašíček model, whereas we set $r = r_0$ for the constant interest rates assumption. We employ the base parameter set as in Table 8.2, and, additionally, assume $V_0 = 100$, $F = 40$, $K = 50$, $C = 1$ (payable at the middle and end of the year), $D = 2$ (payable at the first and third quarters of the year), $m = 1$, $\gamma = 0.2$, $s^c = 0$, $\vartheta = 0$. We adopt here the scaled constant F , C and D values, as in the example of Brennan and Schwartz ((1977), Section V). Under stochastic interest rates, the precision of the reported numbers is up to the third decimal place. When the interest rate is constant, we acquire higher CPU power and produce results precise to the fifth decimal place. Higher accuracies (up to 7 decimal places) are possible via Richardson extrapolation, due to the smooth linear convergence of the numerical scheme in the number of grid points. As it becomes obvious from Tables 8.4 and 8.5, the CPU timings rise from the constant to the stochastic interest rates setup, and from the simply callable CB to the coupon-bearing one and to another CB with associated dividend-paying stock. In fact, the increase originated by the introduction of stochastic interest rates is due to the change from 1-D to 2-D Fourier transforms, whereas the additional computational times required by a callable CB with coupons, and with both coupons and dividends, are due to the need to approximate the CB values at the relevant time points on three different costly 2-D grids (see step 2 of the numerical implementation). In any case, we do not exceed the typical 6700 seconds Fortran execution time, independent of the contract specification, of the PDI implementation reported in Bermúdez and Webber (2004).

Few comments are in order. Adding coupons in the bond indenture raises substantially the payoff to the investors and, consequently, the CB value. At the same time, the firm value and, consequently, the chances for a call reduce, increasing in this way the value of the CB, whilst the default event becomes more likely, negatively affecting the CB value. Nevertheless, the first two effects beat the third one, justifying the overall increase in the CB value observed. Moreover, for a dividend-paying common stock, a decline in the contract's price is noticed. This occurs because the dividends are not payable to the CB holders pre-conversion and, at the same time, they affect the rate at which the firm value appreciates, boosting, in this way, the chances of future default.

Furthermore, the discrepancy between the MJD and the DEJD model prices remains positive

CB specification	model		price differences $\times 10^{-4}$	CPU (s)		
	MJD	DEJD		prec $\pm 10^{-5}$	prec $\pm 10^{-4}$	prec $\pm 10^{-3}$
call	33.40333	33.39807	53	220	2.2	1.1
call, coupons	41.91545	41.91100	45	250	2.4	1.2
call, coupons, dividends	40.76062	40.75578	48	265	4.8	2.5

Table 8.4: Constant interest rates: comparison between MJD and DEJD prices for different CB specifications. CPU timings (in seconds (s)) correspond to accuracy up to 5, 4, 3 decimal places.

CB specification	model		price differences $\times 10^{-4}$	CPU (s)	
	MJD	DEJD		prec $\pm 10^{-3}$	prec $\pm 10^{-2}$
call	33.655	33.650	50	1520	410
call, coupons	42.089	42.085	40	3530	680
call, coupons, dividends	40.702	40.697	50	5740	930

Table 8.5: Stochastic interest rates: comparison between MJD and DEJD prices for different CB specifications. CPU timings (in seconds (s)) correspond to accuracy up to 3, 2 decimal places.

and smaller than the average computed for the out-of-the-money CBs, reported in Table 8.3 for the base case parameters. This price difference reduces in the case of a coupon-bearing CB, since the coupons have a primary positive upshot on the value of the bond, reducing the impact of the stronger negative skewness and fatter tails of the DEJD, as compared to the MJD.

8.5.4 Effects of call policy

In order to investigate the consequences of the adopted call strategy, we assume a typical call notice period of a month ($s^c = 1/12$) and safety premium $\vartheta = 0.2$, and examine how deviations from these initial values affect the model prices produced under the base case and the $\sigma_{L,I}$ parameter set in Table 8.2. We generate prices for callable CBs with 5 years to maturity subject to daily sampling ($T = 5, n = 1250$), $V_0 = 100$, $F = 40$, $K = 50$, $m = 1$, $\gamma = 0.2$. We ignore discrete coupons and dividends in this section. For the interest rates model we assume the same parameters as in Section 8.5.3. The precision of the reported numbers is 5 decimal places (achieved in 270 seconds) and 3 decimal places (achieved in 1570 seconds) for constant and stochastic interest respectively.

In general, increasing ϑ and/or s^c raises the chances for a successful forced-by-call conversion

(s^c, ϑ)	model		$s^c = 1/12$ vs $1/24$		$\vartheta = 0.20$ vs 0.25	
	MJD	DEJD	MJD	DEJD	MJD	DEJD
(1/12, 0.20)	33.20939	33.20351				
(1/12, 0.25)	33.14899	33.14321			0.06040	0.06030
(1/24, 0.20)	33.20935	33.20345	0.00004	0.00006		
(1/24, 0.25)	33.14899	33.14319	$< 10^{-5}$	0.00002	0.06036	0.06026

Table 8.6: Base parameter set, constant interest rates: callable CB prices for varying call specification (s^c, ϑ) . Precision up to 5 decimal places.

(s^c, ϑ)	model		$s^c = 1/12$ vs $1/24$		$\vartheta = 0.20$ vs 0.25	
	MJD	DEJD	MJD	DEJD	MJD	DEJD
(1/12, 0.20)	33.40978	33.40891				
(1/12, 0.25)	33.33957	33.33896			0.07021	0.06995
(1/24, 0.20)	33.40940	33.40846	0.00038	0.00045		
(1/24, 0.25)	33.33946	33.33877	0.00011	0.00019	0.06994	0.06969

Table 8.7: $\sigma_{L,I}$ parameter set, constant interest rates: callable CB prices for varying call specification (s^c, ϑ) . Precision up to 5 decimal places.

at the end of the call notice period and, hence, reduces the CB value. According to the indications in Tables 8.6-8.8, the call notice has a minor effect on the CB values. This is because of the short call notice lengths (30 days, and 15 days) as compared to the longer time to maturity of the CB issue (5 years).⁸ Furthermore, when we change from the base case to the $\sigma_{L,I}$ parameter set, which originates strongest variance, skewness and excess kurtosis features (see Table 8.2), we observe at best an increase in the MJD and DEJD differences across the two s^c choices by a factor of 10, although the actual deviation does not exceed 5×10^{-4} .

On the contrary, the choice of the safety premium, which relates to the firm's decision on the date of the call announcement, appears to be a primary factor driving the CB values. The MJD and DEJD differences across the two ϑ values under consideration are about 6×10^{-2} and 6.5×10^{-2} for the base parameter set, and 7×10^{-2} and 7.3×10^{-2} for the $\sigma_{L,I}$ parameter set, in the case of constant and stochastic interest rates respectively. The effect caused by ϑ is more pronounced when the jump sizes are more volatile, due to the increase in the skewness and excess kurtosis, similarly to what observed in Section 8.5.2.

⁸We have reached the same conclusion for $T = 2$.

(s^c, ϑ)	model		$\vartheta = 0.20$ vs 0.25		(s^c, ϑ)	model		$\vartheta = 0.20$ vs 0.25	
	MJD	DEJD	MJD	DEJD		MJD	DEJD	MJD	DEJD
$(1/12, 0.20)$	33.451	33.445			$(1/12, 0.20)$	33.633	33.634		
$(1/12, 0.25)$	33.385	33.380	0.066	0.065	$(1/12, 0.25)$	33.560	33.561	0.073	0.073
$(1/24, 0.20)$	33.451	33.445			$(1/24, 0.20)$	33.633	33.634		
$(1/24, 0.25)$	33.385	33.380	0.066	0.065	$(1/24, 0.25)$	33.560	33.561	0.073	0.073

Table 8.8: Left panel: base parameter set; right panel: $\sigma_{L,I}$ parameter set, stochastic interest rates: callable CB prices for varying call specification (s^c, ϑ) . Precision up to 3 decimal places.

8.6 Concluding remarks

We have developed and implemented a backward price convolution scheme for convertible bonds. Supported by Fourier transforms techniques, the proposed method is shown to be efficient and accurate, and to flexibly accommodate a number of contract-design features such as callability provisions, dividends and coupon payments. The procedure has also been shown capable of coping with up to four risk factors, allowing a market setup based on a jump diffusion-driven underlying asset for the CB, and stochastic interest rates.

The proposed pricing methodology has been tested using several parameter sets of the market model. As a benchmark to the Fourier transform algorithm, we have used the closed-form solution obtained by Ingersoll (1977a) for the case of a continuously callable CB in the Black-Scholes-Merton framework with constant interest rate. The numerical results have indicated accuracies up to 7 decimal places, for varying monitoring frequency. The analysis has been then extended first to a jump diffusion setting with constant interest rates, then to the case of a stochastic term structure of interest rates.

As a jump diffusion market setup for pricing CBs is new in the literature, we have also used the proposed algorithm to analyze the behaviour of the contract price under this more complex representation of the firm value. The main results of the numerical analysis show that the jump diffusion setup originates lower values of the CB when compared to the classical Black-Scholes-Merton framework. This is essentially due to the higher probability of default generated by the inclusion of market shocks in the model, i.e., by the negatively skewed and leptokurtic distribution of the log-firm value resulting from the jump diffusion setup.

An issue that is not dealt with in this work is the calibration of the market model. In fact,

as discussed in Section 8.2.2, we adopt a structural approach to the credit risk in the same spirit as Merton (1974). However, the fact that the firm value is not directly observable in the market, leaves open the problem of model calibration. In Section 8.2.3 we discuss a few possible solutions to this issue, the implementation and testing of which is left for future research.

Finally, we note that the CB valuation scheme we present here is general enough to accommodate, for example, a stock-based setting, as opposed to the current firm value-based model, with jumps and stochastic interest rates, for suitably modified CB payoff functions (8.10-8.13) (e.g., Goldman Sachs (1994), Barone-Adesi et al. (2003)). Also, if necessary, the put provision can be flexibly incorporated into the payoff function (see Goldman Sachs (1994)). Apart from the computation of the CB prices, our method can be extended to the computation of the price sensitivities. Moreover, the extensive CB pricing scheme we suggest here can be easily reduced and specialized in pricing simpler exotic derivatives and, in fact, extend the work of Lord et al. (2008) on the pricing of Bermudan/American vanilla options to the case of stochastic interest rates.

Appendix 8.A—Equivalent martingale measure changes: the t -forward measure

Proposition 24 *For the pure-discount bond price $P_s(t)$, $0 \leq s \leq t$, in (8.4), there exists a martingale measure $\mathbb{P}^*$ defined by its Radon-Nikodým derivative with respect to $\mathbb{P}$*

$$\begin{aligned} \gamma_s^* &: = \frac{d\mathbb{P}^*}{d\mathbb{P}} \Big|_{\mathcal{F}_s} = \exp \left(-\rho^2 \int_0^s \frac{m_u^2(t)}{2} du - (1 - \rho^2) \int_0^s \frac{m_u^2(t)}{2} du \right. \\ &\quad \left. + \rho \int_0^s m_u(t) dW_u + \sqrt{1 - \rho^2} \int_0^s m_u(t) d\tilde{W}_u \right), \quad s \in [0, t], \end{aligned}$$

where W and $\tilde{W}$ are independent $\mathbb{P}$ -standard Brownian motions. We define $\mathbb{P}^*$ as the t -forward measure associated to the numéraire $P_s(t)$.

Proof. The pure-discount bond price $P_s(t)$ satisfies the definition of a numéraire, as in Geman et al. (1995). Then, based on Geman et al. ((1995), Theorem 1), we construct the Radon-Nikodým derivative

$$\gamma_s^* = \frac{d\mathbb{P}^*}{d\mathbb{P}} \Big|_{\mathcal{F}_s} = \frac{P_s(t) e^{-\int_0^s r_u du}}{P_0(t)}.$$

From (8.4), we have that

$$\begin{aligned} \gamma_s^* &= \exp \left(-\int_0^s \frac{m_u^2(t)}{2} du + \int_0^s m_u(t) dW_{r,u} \right) \\ &= \exp \left(-\rho^2 \int_0^s \frac{m_u^2(t)}{2} du - (1 - \rho^2) \int_0^s \frac{m_u^2(t)}{2} du + \rho \int_0^s m_u(t) dW_u \right. \\ &\quad \left. + \sqrt{1 - \rho^2} \int_0^s m_u(t) d\tilde{W}_u \right), \end{aligned}$$

which follows in virtue of the decomposition $W_r = \rho W + \sqrt{1 - \rho^2} \tilde{W}$, for independent Brownian motions W and $\tilde{W}$. ■

Based on Proposition 24, we conclude, in virtue of the Girsanov theorem, that W and W_r retain their semimartingale property and decompose, after the measure change, to

$$\begin{aligned}
W_s &= W_s^* + \int_0^s \rho m_u(t) du, \\
W_{r,s} &= W_{r,s}^* + \int_0^s m_u(t) du,
\end{aligned}$$

where W^* and W_r^* are correlated $\mathbb{P}^*$ -standard Brownian motions with constant correlation $\rho \in (-1, 1)$.

Appendix 8.B—Characterization of the bivariate log-firm value—interest rate process under the forward measure

In the subsequent derivation, for compactness we denote $E(\cdot | \mathcal{F}_t) = E_t$ under a generic probability measure P .

Following Geman et al. ((1995), Corollary 2) and the change to the $\mathbb{P}^*$ measure as described in Proposition 24, we have that

$$\mathbb{E}_s^*(\exp(iv_1 Y_t + iv_2 r_t)) = \frac{\mathbb{E}_s\left(\exp\left(-\int_s^t r_u du + iv_1 Y_t + iv_2 r_t\right)\right)}{P_s(t)}. \quad (8.22)$$

Given the independence of the jump part of the jump diffusion process from the diffusion part and the short rate, we restate the right-hand side of equation (8.22) as

$$\begin{aligned} & \frac{1}{P_s(t)} \mathbb{E}_s \left(\exp \left(-\int_s^t r_u du + iv_1 \left(Y_s + \int_s^t (r_u - \sigma^2/2 - \lambda(\phi_L(-i) - 1)) du \right. \right. \right. \\ & \left. \left. \left. + \sigma \int_s^t dW_u \right) + iv_2 r_t \right) \right) \mathbb{E}_s \left(\exp \left(iv_1 \int_{\mathbb{R}} l(N_t(dl) - N_s(dl)) \right) \right). \end{aligned} \quad (8.23)$$

Then, equation (8.23) is equivalent to

$$\begin{aligned} & \frac{1}{P_s(s+\tau)} \exp \left(Y_s + \frac{\sigma^2}{2} \frac{iv_1}{iv_1 - 1} \tau \right) \mathbb{E}_s \left(\exp \left((iv_1 - 1) (Y_s \right. \right. \\ & \left. \left. + \int_s^{s+\tau} \left(r_u - \frac{1}{2} \left(\sigma \frac{iv_1}{iv_1 - 1} \right)^2 \right) du + \sigma \frac{iv_1}{iv_1 - 1} \int_s^{s+\tau} dW_u \right) + iv_2 r_{s+\tau} \right) \right) \\ & \times \exp(-i\lambda(\phi_L(-i) - 1)v_1\tau) \mathbb{E}_s \left(\exp \left(iv_1 \int_{\mathbb{R}} l(N_{s+\tau}(dl) - N_s(dl)) \right) \right), \end{aligned} \quad (8.24)$$

where $\tau = t - s$. Next, we define

$$\begin{aligned} \tilde{\sigma} &= \sigma \frac{v_1}{\tilde{v}_1}, \\ \tilde{v}_1 &= \frac{iv_1 - 1}{i}, \\ \tilde{Y}_t &= Y_s + \int_s^t (r_u - \tilde{\sigma}^2/2) du + \tilde{\sigma} \int_s^t dW_u \end{aligned}$$

and split (8.24) into the product of

$$C_s(\tau) = \frac{1}{P_s(s+\tau)} \exp(Y_s + \sigma \tilde{\sigma} \tau / 2), \quad (8.25)$$

$$\begin{aligned} D_s(v_1, v_2, \tau) &= \mathbb{E}_s \left(\exp \left(i\tilde{v}_1 \left(Y_s + \int_s^{s+\tau} (r_u - \tilde{\sigma}^2/2) du + \tilde{\sigma} \int_s^{s+\tau} dW_u \right) \right. \right. \\ &\quad \left. \left. + i v_2 r_{s+\tau} \right) \right) \\ &= \mathbb{E}_s \left(\exp \left(i\tilde{v}_1 \tilde{Y}_{s+\tau} + i v_2 r_{s+\tau} \right) \right), \end{aligned} \quad (8.26)$$

$$\begin{aligned} F_s(v_1, \tau) &= \exp(-i\lambda(\phi_L(-i) - 1)v_1\tau) \mathbb{E}_s \left(\exp \left(i v_1 \int_{\mathbb{R}} l(N_{s+\tau}(dl) - N_s(dl)) \right) \right) \\ &= \exp(\lambda(-i(\phi_L(-i) - 1)v_1 + \phi_L(v_1) - 1)\tau). \end{aligned} \quad (8.27)$$

As far as (8.26) is concerned, we apply the characterization of regular affine Markov processes of Duffie et al. (2000) and Duffie et al. (2003), as stated in Section 3.3 of this thesis. Based on this, we have that

$$D_s(v_1, v_2, \tau) = \exp(E_0(v_1, v_2, \tau) + E_1(v_1, v_2, \tau)r_s + E_2(v_1, v_2, \tau)Y_s), \quad (8.28)$$

where E_0 , E_1 and E_2 satisfy the subsequent system of generalized Riccati equations:

$$\begin{aligned} \frac{\partial E_0(v_1, v_2, u)}{\partial u} &= \kappa \mu_r E_1(v_1, v_2, u) - \frac{\tilde{\sigma}^2}{2} E_2(v_1, v_2, u) + \frac{1}{2} (\sigma_r^2 E_1^2(v_1, v_2, u) \\ &\quad + 2\tilde{\sigma} \sigma_r \rho E_1(v_1, v_2, u) E_2(v_1, v_2, u) + \tilde{\sigma}^2 E_2^2(v_1, v_2, u)), \\ \frac{\partial E_1(v_1, v_2, u)}{\partial u} &= -\kappa E_1(v_1, v_2, u) + E_2(v_1, v_2, u), \\ \frac{\partial E_2(v_1, v_2, u)}{\partial u} &= 0, \\ E_0(v_1, v_2, s) &= 0, \quad E_1(v_1, v_2, s) = i v_2, \quad E_2(v_1, v_2, s) = i \tilde{v}_1. \end{aligned}$$

Given the fact that the log-firm value model contains no mean-reverting term and the interest rate process is of Ornstein-Uhlenbeck type, we explicitly obtain from Kallsen ((2006), Corollary 3.5) that

$$E_2(v_1, v_2, \tau) = i \tilde{v}_1, \quad (8.29)$$

$$\begin{aligned}
 E_1(v_1, v_2, \tau) &= iv_2 \exp(-\kappa\tau) + \frac{1 - \exp(-\kappa\tau)}{\kappa} i\tilde{v}_1 \\
 &= iv_2 (\cosh(\kappa\tau) - \sinh(\kappa\tau)) + \frac{i\tilde{v}_1}{\kappa} (1 - \cosh(\kappa\tau) + \sinh(\kappa\tau)), \quad (8.30)
 \end{aligned}$$

$$\begin{aligned}
 E_0(v_1, v_2, \tau) &= \int_s^{s+\tau} (\kappa\mu_r E_1(v_1, v_2, u) - i\tilde{\sigma}^2 \tilde{v}_1/2 + (\sigma_r^2 E_1^2(v_1, v_2, u) \\
 &\quad + 2i\tilde{\sigma}\sigma_r \rho E_1(v_1, v_2, u) \tilde{v}_1 - \tilde{\sigma}^2 \tilde{v}_1^2)/2) du \\
 &= \frac{i\tilde{v}_1}{\kappa} \left(\kappa\mu_r + i\tilde{\sigma}\sigma_r \rho \tilde{v}_1 + \frac{i\sigma_r^2 \tilde{v}_1}{2\kappa} \right) \tau + i\tilde{\sigma}^2 \tilde{v}_1 (i\tilde{v}_1 - 1) \tau/2 \\
 &\quad + \frac{i}{\kappa} \left(v_2 - \frac{\tilde{v}_1}{\kappa} \right) (\sinh(\kappa\tau) - \cosh(\kappa\tau) + 1) \left(\kappa\mu_r + i\tilde{\sigma}\sigma_r \rho \tilde{v}_1 + \frac{i\sigma_r^2 \tilde{v}_1}{\kappa} \right) \\
 &\quad - \frac{\sigma_r^2}{4\kappa} \left(v_2 - \frac{\tilde{v}_1}{\kappa} \right)^2 (\sinh(2\kappa\tau) - \cosh(2\kappa\tau) + 1). \quad (8.31)
 \end{aligned}$$

Summarizing, equations (8.25) and (8.27-8.31) lead to

$$\begin{aligned}
 \mathbb{E}_s^*(\exp(iv_1 Y_{s+\tau} + iv_2 r_{s+\tau})) &= C_s(\tau) \exp(E_0(v_1, v_2, \tau) + E_1(v_1, v_2, \tau) r_s \\
 &\quad + E_2(v_1, v_2, \tau) Y_s) F(v_1, \tau).
 \end{aligned}$$

Furthermore, upon considering the transformation

$$(Z, Z_r) = (Y_{s+\tau} - Y_s + \ln P_s(s + \tau), r_{s+\tau} - r_s \exp(-\kappa\tau)),$$

we obtain that

$$\phi^*(v_1, v_2) = \mathbb{E}^*(\exp(iv_1 Z + iv_2 Z_r))$$

is independent of Y_s and r_s .

In fact,

$$\begin{aligned}
\phi^*(v_1, v_2) &= \exp(iv_1(-Y_s + \ln P_s(s + \tau)) - iv_2 r_s \exp(-\kappa\tau)) \\
&\quad \times \mathbb{E}_s^*(\exp(iv_1 Y_{s+\tau} + iv_2 r_{s+\tau})) \\
&= \exp(iv_1(-Y_s + \ln P_s(s + \tau)) - iv_2 r_s \exp(-\kappa\tau) \\
&\quad - \ln P_s(s + \tau) + Y_s + \frac{\sigma \tilde{\sigma}}{2} \tau + E_0(v_1, v_2, \tau) \\
&\quad + E_1(v_1, v_2, \tau) r_s + E_2(v_1, v_2, \tau) Y_s) F(v_1, \tau) \\
&= \exp\left(\frac{\sigma^2}{2} \frac{iv_1}{iv_1 - 1} \tau + (iv_1 - 1) A_s(s + \tau) + E_0(v_1, v_2, \tau)\right) F(v_1, \tau), \quad (8.32)
\end{aligned}$$

where $A_s(s + \tau)$ is given by equation (8.7).

Chapter 9

Concluding remarks

In this thesis, we focus on two financial instruments whose accurate and efficient pricing has been for years a major concern for both the practitioners and academics. The Fourier transform-based numerical integration method we develop in this thesis, deals with convergence issues, and dimensionality and payout structure limitations reported in previous attempts. It offers clear advantages over other numerical techniques in terms of its flexibility to capture real-world specifications for both products, despite their completely different payout structure, and its generality in respect of the modelling assumptions for the underlying state variables. Hence, we believe that this work contributes significantly to an important area of finance.

In the first part of the thesis, the main focus is on the pricing of discretely sampled arithmetic Asian options. Under the assumptions of the Black-Scholes-Merton economy, the proposed method is shown to be numerically competitive to the well-established PDE technique of Večer¹ (2002) and the efficient analytical approximation by Lord (2006a), especially for high asset volatilities. By further expanding to the non-Gaussian Lévy context, our backward price convolution approach clearly distinguishes from the previous forward density convolutions by Carverhill and Clewlow (1990), Benhamou (2002) and Fusai and Meucci (2008) which exhibit non-monotone convergence. By smooth convergence of our scheme, extrapolation in space for a fixed number of monitoring points guarantees highly precise results, whereas extrapolating in the time dimension speeds up the convergence to the price of a continuously sampled option.

¹A detailed numerical comparison against Večer's PDE technique can be found in Černý and Kyriakou (2010).

Furthermore, extension to calculate the price sensitivities is straightforward without affecting the order of computational effort. Comparisons against efficient Monte Carlo schemes, equipped with accurate analytical prices/sensitivities of geometric Asian options as control variates, illustrate the numerical advantage of our convolution technique.

The need to simulate Lévy processes has further motivated our study on the tempered stable process, which has hitherto proved difficult to simulate. The joint Monte Carlo-Fourier transform simulation scheme we propose copes with the limitations of the biased scheme proposed by Madan and Yor (2008), and the approach of Poirrot and Tankov (2006) specifically tailored to path-independent contracts. We are currently exploring a multivariate generalization of the tempered stable process (CGMY subclass) which can accommodate a full range of possible dependence, as opposed to previous contributions in the literature (see Ballotta and Bonfiglioli (2010)). We will then employ our joint Monte Carlo-Fourier transform scheme to evaluate contracts written on more than one underlying, e.g., basket and Asian basket options, which are commonly traded in the credit and energy markets. This is important, since pure numerical integration techniques are too slow to compete with Monte Carlo in pricing this type of contracts due to increasing model dimensionality (e.g., see Lord et al. (2008)).

We also consider an extension of the convolution method from the Lévy to the two-dimensional Heston and Bates stochastic volatility frameworks for pricing discretely sampled Asian options. As a benchmark, we utilize an effective Monte Carlo strategy, which combines the biased sampling mechanism of Andersen (2008) with the geometric Asian price we derive for non-Lévy log-returns as a control variate. Numerical examples illustrate that the convolution procedure has an extra edge for high precision levels, whereas the control variate Monte Carlo is faster when smaller accuracies are sought. The latter also requires that we pre-estimate the inherent bias, which varies with the number of time steps employed; however, to accurately measure the bias, it is necessary to hold in advance the exact option price, which may only be computed via the convolution method. Furthermore, in contrast with the Monte Carlo method, the convolution method provides us with the option prices on a grid of initial variance values.

Overall, our study suggests that, for non-Gaussian Lévy models, the skewness and excess kurtosis are the primary factors driving the option prices, rather than the actual model choice. The absence of skewness and excess kurtosis in the Gaussian distribution also explains the ob-

served gap between the prices for Gaussian and non-Gaussian Lévy log-returns. The impact of the model choice on the option deltas and gammas is even weaker. Introducing, further, stochastic volatility relaxes the assumption of independence of the log-returns, leading to departures from the Lévy-based prices.

Interesting extensions for future research include the pricing of forward-start and in-progress Asian options, with time to maturity greater or smaller than the averaging period respectively. From the implementation side, the next step in our research would be to implement a computational algorithm which involves sequential evaluation in the Fourier space, skipping successive transformations between the state and Fourier spaces; a novel method towards this direction has been proposed recently by Feng and Linetsky (2008) based on the fast Hilbert transform, which involves a discrete approximation with exponentially decaying errors. Of similar performance is also the technique suggested by Fang and Oosterlee (2008b) based on Fourier-cosine series expansions.

In the second part of the thesis, we propose a new valuation framework for CBs, which comprises a jump diffusion model for the firm value process and correlated stochastic interest rate movements. Adopting such an approach allows default to be reached following a number of consecutive shocks in the value of the firm. This complements previous contribution by Bermúdez and Webber (2004) who, by limitation of their PDI scheme, assume that default occurs only once, and thereafter the firm value evolves as a pure diffusion, eliminating the possibility of future exogenous default events to occur. Our numerical analysis has shown that the inclusion of market shocks in the firm value process increases the probability of default, resulting in lower CB values when compared to the pure diffusion paradigm. Furthermore, the numerical scheme we propose is able to accommodate flexibly a number of contract-design features, such as discrete cash flows, conversion forced by a call from the issuer, or voluntary conversion at the option of the CB holder prior to a dividend payment.

Although firm-value approaches provide a natural link between debt and equity, which is ideal in the CB valuation, the fact that the firm value is an unobserved state variable, leaves open the problem of model calibration. We have discussed a few possible routes towards solving this issue, the implementation and testing of which is left for future research. Alternatively, the presented valuation scheme can be tailored to an equity-based setting with jumps and stochastic

CONCLUDING REMARKS

interest rates, subject to suitably modified payoff functions. It can be further specialized in pricing of other Bermudan/American options and barrier options under stochastic interest rates.

References

- Abate, J. and W. Whitt (1992). The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems* 10, 5–88.
- Abramowitz, M. and I. A. Stegun (1968). *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables*. New York: Dover.
- Aït-Sahalia, Y. (1996). Nonparametric pricing of interest rate derivative securities. *Econometrica* 64, 527–560.
- Albrecher, H. and M. Predota (2002). Bounds and approximations for discrete Asian options in a Variance-Gamma model. *Grazer Mathematische Berichte* 345, 35–57.
- Albrecher, H. and M. Predota (2004). On Asian option pricing for NIG Lévy processes. *Journal of Computational and Applied Mathematics* 172, 153–168.
- Albrecher, H. and W. Schoutens (2005). *Static hedging of Asian options under stochastic volatility models using Fast Fourier Transform*. In Exotic Option Pricing and Advanced Lévy Models, A. Kyprianou, W. Schoutens and P. Wilmott, eds. New York: Wiley.
- Ammann, M., A. Kind, and C. Wilde (2008). Simulation-based pricing of convertible bonds. *Journal of Empirical Finance* 15, 310–331.
- Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance* 11, 1–42.
- Andersen, L. and D. Buffum (2003). Calibration and implementation of convertible bond models. *Journal of Computational Finance* 7, 1–34.

- Andreasen, J. (1998). The pricing of discretely sampled Asian and lookback options: a change of numeraire approach. *Journal of Computational Finance* 2, 5–30.
- Andricopoulos, A. D., M. Widdicks, D. P. Newton, and P. W. Duck (2007). Extending quadrature methods to value multi-asset and complex path dependent options. *Journal of Financial Economics* 83, 471–499.
- Asquith, P. (1995). Convertible bonds are not called late. *Journal of Finance* 50, 1275–1289.
- Asquith, P. and D. W. Mullins (1991). Convertible debt: corporate call policy and voluntary conversion. *Journal of Finance* 46, 1273–1289.
- Ayache, E., P. A. Forsyth, and K. R. Vetzal (2003). Valuation of convertible bonds with credit risk. *Journal of Derivatives* 11, 9–29.
- Bailey, D. H. and P. N. Swarztrauber (1991). The fractional Fourier transform and applications. *SIAM Review* 33, 389–404.
- Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics* 55, 205–238.
- Ballotta, L. (2010). Efficient pricing of ratchet equity index annuities in a VG economy. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1135696. To appear in North American Actuarial Journal.
- Ballotta, L. and E. Bonfiglioli (2010). Multivariate asset models using Lévy processes and applications. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1695527. Working paper, Cass Business School.
- Barndorff-Nielsen, O. E. (1998). Processes of normal inverse Gaussian type, *Finance and Stochastics* 2, 41–68.
- Barndorff-Nielsen, O. E. and N. Shephard (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society, Series B* 63, 167–241.

- Barone-Adesi, G., A. Bermúdez, and J. Hatgioannides (2003). Two-factor convertible bonds valuation using the method of characteristics/finite elements. *Journal of Economic Dynamics and Control* 27, 1801–1831.
- Bates, D. S. (1996). Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. *Review of Financial Studies* 9, 69–107.
- Bayraktar, E. and H. Xing (2011). Pricing Asian options for jump diffusion. *Mathematical Finance* 21, 117–143.
- Benhamou, E. (2000). An application of Malliavin calculus to continuous time Asian options. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=265284. Working paper, London School of Economics.
- Benhamou, E. (2002). Fast Fourier transform for discrete Asian options. *Journal of Computational Finance* 6, 49–68.
- Benhamou, E. (2003). Optimal Malliavin weighting function for the computation of the Greeks. *Mathematical Finance* 13, 37–53.
- Bermúdez, A. and N. Webber (2004). An asset based model of defaultable convertible bonds with endogenized recovery. Available from http://www.cls.dk/caf/pp_04.134.pdf. Working paper, City University.
- Bernis, G., E. Gobet, and A. Kohatsu-Higa (2003). Monte Carlo evaluation of Greeks for multidimensional barrier and lookback options. *Mathematical Finance* 13, 99–113.
- Bielecki, T. R. and M. Rutkowski (2002). *Credit risk: Modeling, Valuation and Hedging*. New York: Springer.
- Black, F. and J. C. Cox (1976). Valuing corporate securities: some effects of bond indenture provisions. *Journal of Finance* 31, 351–367.
- Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.

- Bluestein, L. I. (1968). A linear filtering approach to the computation of the discrete Fourier transform. *IEEE Northeast Electronics Research and Engineering Meeting* 10, 218–219.
- Bochner, S. and K. Chandrasekharan (1949). *Fourier transforms*. Annals of Mathematics Studies, No. 19. Princeton: Princeton University Press.
- Boyarchenko, S. I. and S. Z. Levendorskiĭ (2002). *Non-Gaussian Merton-Black-Scholes Theory*. River Edge: World Scientific.
- Boyle, P. and F. Boyle (2001). *Derivatives: The Tools that Changed Finance*. London: Risk Books.
- Boyle, P., M. Broadie, and P. Glasserman (1997). Monte Carlo method for security pricing. *Journal of Economic Dynamics and Control* 21, 1267–1321.
- Boyle, P. and A. Potapchik (2008). Prices and sensitivities of Asian options: a survey. *Insurance: Mathematics and Economics* 42, 189–211.
- Brennan, M. J. and E. S. Schwartz (1977). Convertible bonds: valuation and optimal strategies for call and conversion. *Journal of Finance* 32, 1699–1715.
- Brennan, M. J. and E. S. Schwartz (1980). Analyzing convertible bonds. *Journal of Financial and Quantitative Analysis* 15, 907–929.
- Broadie, M. and J. Detemple (2004). Option pricing: valuation models and applications. *Management Science* 50, 1145–1177.
- Broadie, M. and P. Glasserman (1996). Estimating security price derivatives using simulation. *Management Science* 42, 269–285.
- Broadie, M. and O. Kaya (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* 54, 217–231.
- Broadie, M. and Y. Yamamoto (2003). Application of the fast Gauss transform to option pricing. *Management Science* 49, 1071–1008.
- Broadie, M. and Y. Yamamoto (2005). A double-exponential fast Gauss transform algorithm for pricing discrete path-dependent options. *Operations Research* 53, 764–779.

- Cai, N. and S. G. Kou (2010). Pricing Asian options under a general jump diffusion model. Available from http://www.ieor.columbia.edu/pdf-files/IEOR_Oxford_Kou.pdf. Working paper, Columbia University.
- Carayannopoulos, P. (1996). Valuing convertible bonds under the assumption of stochastic interest rates: an empirical investigation. *Quarterly Journal of Business and Economics* 35, 17–31.
- Carayannopoulos, P. and M. Kalimipalli (2003). Convertible bond prices and inherent biases. *Journal of Fixed Income* 13, 64–73.
- Carr, P., H. Geman, D. B. Madan, and M. Yor (2002). The fine structure of asset returns: an empirical investigation. *The Journal of Business* 75, 305–332.
- Carr, P., H. Geman, D. B. Madan, and M. Yor (2003). Stochastic volatility for Lévy processes. *Mathematical Finance* 13, 345–382.
- Carr, P. and D. B. Madan (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance* 2, 61–73.
- Carr, P. and L. Wu (2004). Time-changed Lévy process and option pricing. *Journal of Financial Economics* 71, 113–141.
- Carverhill, A. and L. Clewlow (1990). Flexible convolution. *Risk* 3, 25–29.
- Černý, A. (2004). Introduction to FFT in finance. *Journal of Derivatives* 12, 73–88.
- Černý, A. and I. Kyriakou (2010). An improved convolution algorithm for discretely sampled Asian options. Available from <http://ssrn.com/abstract=1323252>. To appear in Quantitative Finance.
- Chambers, J., C. Mallows, and B. Stuck (1976). A method of simulating stable random variables. *Journal of American Statistical Association* 71, 340–344.
- Chen, N. and S. Kou (2009). Credit spreads, optimal capital structure, and implied volatility with endogenous default and jump risk. *Mathematical Finance* 19, 343–378.

- Cheung, W. and I. Nelken (1994). Costing the converts. *Risk* 7, 47–49.
- Chourdakis, K. (2004). Option pricing using the fractional FFT. *Journal of Computational Finance* 8, 1–18.
- Cont, R. and P. Tankov (2004a). *Financial modelling with jump processes*. Boca Raton: Chapman & Hall/CRC Press.
- Cont, R. and P. Tankov (2004b). Non-parametric calibration of jump-diffusion option pricing models. *Journal of Computational Finance* 7, 1–49.
- Conze, A. and Viswanathan (1991). European path dependent options: the case of geometric averages. *Finance* 91, 7–22.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica* 53, 385–407.
- CreditGrades (2002). CreditGrades technical document. Available from <http://www.riskmetrics.com/publications/techdocs/cgtdovv.html>. Risk Technical Documents.
- Curran, M. (1994a). Valuing Asian and portfolio options by conditioning on the geometric mean price. *Management Science* 40, 1705–1711.
- Curran, M. (1994b). Strata Gems. *Risk*, 70–71.
- Curran, M. (1998). *Greeks in Monte Carlo*. In Monte Carlo Methodologies and Applications for Pricing and Risk Management, B. Dupire, ed. London: Risk Books.
- Dao, B. and M. Jeanblanc (2006). Double exponential jump diffusion process: a structural model of endogenous default barrier with roll-over debt structure. Available from http://www.maths.univ-evry.fr/pages_perso/jeanblanc/publication.html. Working paper, Université Paris Dauphine and Université d’Évry.
- Davies, B. (2002). *Integral Transforms and Their Applications* (3rd ed.). New York: Springer.
- Davis, M. and F. R. Lischka (2002). *Convertible bonds with market risk and credit risk*. In Studies in Advanced Mathematics, Series 26, R. H. Chan, Y. K. Kwok, D. Yao and Q. Zhang, eds. American Mathematical Society and International Press.

- Deelstra, G. and F. Delbaen (1998). Convergence of discretized stochastic (interest rate) processes with stochastic drift term. *Applied Stochastic Models and Data Analysis* 14, 77–84.
- Delbaen, F. and W. Schachermayer (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen* 300, 463–520.
- Dhaene, J., M. Denuit, M. J. Goovaerts, R. Kaas, and D. Vyncke (2002). The concept of comonotonicity in actuarial science and finance: theory. *Insurance: Mathematics and Economics* 31, 3–33.
- Duffie, D., D. Filipović, and W. Schachermayer (2003). Affine processes and applications in finance. *The Annals of Applied Probability* 13, 984–1053.
- Duffie, D. and P. Glynn (1995). Efficient Monte Carlo simulation of security prices. *Annals of Applied Probability* 5, 897–905.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Duffie, D. and K. J. Singleton (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies* 12, 687–720.
- Dufresne, D. (2000). The Laguerre series for Asian and other options. *Mathematical Finance* 10, 407–428.
- Eydeland, A. (1994). A fast algorithm for computing integrals in function spaces: financial applications. *Computational Economics* 7, 277–285.
- Fang, F. and C. W. Oosterlee (2008a). A novel pricing method for European options based on Fourier-Cosine series expansions. *SIAM Journal on Scientific Computing* 31, 826–848.
- Fang, F. and C. W. Oosterlee (2008b). Pricing early-exercise and discrete barrier options by Fourier-Cosine series expansions. Available from <http://ta.twi.tudelft.nl/mf/users/oosterle/oosterlee/bermCOS.pdf>. Working paper, Delft University of Technology.

- Fang, F. and C. W. Oosterlee (2009). The COS method in pricing Bermudan and barrier options under Heston’s model. Available from <http://cermics.enpc.fr/cnf/Fang.pdf>. Presentation, 3rd Conference on Numerical Methods in Finance, April 2009.
- Feng, L. and X. Lin (2009). Hilbert transform approach for pricing Bermudan options in Lévy models. Presentation, Spectral and Cubature Methods in Finance and Econometrics Workshop, June 2009.
- 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, 337–384.
- Filipović, D. and E. Mayerhofer (2009). Affine diffusion processes: theory and applications. *Radon Series on Computational and Applied Mathematics* 8, 1–40.
- Fishman, G. S. (1996). *Monte Carlo: Concepts, Algorithms, and Applications*. New York: Springer.
- Fons, J. S. (1994). Using default rates to model the term structure of credit risk. *Financial Analysts Journal* 50, 25–32.
- Fouque, J-P. and C-H. Han (2003). Pricing Asian options with stochastic volatility. *Quantitative Finance* 3, 353–362.
- Fouque, J-P. and C-H. Han (2004a). Asian options under multiscale stochastic volatility. *Proceedings of the AMS-IMS-SIAM Summer Conference on Mathematics of Finance*, 125–138.
- Fouque, J-P. and C-H. Han (2004b). Variance reduction for Monte Carlo methods to evaluate option prices under multi-factor stochastic volatility models. *Quantitative Finance* 4, 597–606.
- Fouque, J-P., G. Papanicolaou, and K. R. Sircar (2000). *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge: Cambridge University Press.
- Fournié, E., J-M. Lasry, J. Lebuchoux, P-L. Lions, and N. Touzi (1999). Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance and Stochastics* 3, 391–412.

- Fournié, E., J-M. Lasry, J. Lebuchoux, P-L. Lions, and N. Touzi (2001). Applications of Malliavin calculus to Monte Carlo methods in finance. II. *Finance and Stochastics* 5, 201–236.
- Fu, M. C., D. B. Madan, and T. Wang (1999). Pricing continuous Asian options: a comparison of Monte Carlo and Laplace transform inversion methods. *Journal of Computational Finance* 2, 49–74.
- Fusai, G. (2004). Pricing Asian options via Fourier and Laplace transforms. *Journal of Computational Finance* 7, 87–106.
- Fusai, G. and A. Meucci (2008). Pricing discretely monitored Asian options under Lévy processes. *Journal of Banking and Finance* 32, 2076–2088.
- García, D. (2003). Convergence and biases of Monte Carlo estimates of American option prices using a parametric exercise rule. *Journal of Economic Dynamics and Control* 27, 1855–1879.
- Geman, H., N. El Karoui, and J. Rochet (1995). Changes of numéraire, changes of probability measure and option pricing. *Journal of Applied Probability* 32, 443–458.
- Geman, H. and M. Yor (1993). Bessel processes, Asian options and perpetuities. *Mathematical Finance* 3, 349–375.
- Geske, R. and K. Shastri (1985). Valuation by approximation: a comparison of alternative option valuation techniques. *Journal of Financial and Quantitative Analysis* 20, 45–71.
- Gil-Pelaez, J. (1951). Note on the inversion theorem. *Biometrika* 37, 481–482.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. New York: Springer.
- Glasserman, P. and K. Kim (2008a). Gamma expansion of the Heston stochastic volatility model. Available from <http://www2.gsb.columbia.edu/faculty/pglasserman/Other/>. To appear in *Finance and Stochastics*.
- Glasserman, P. and K. Kim (2008b). Beta approximation for bridge sampling. *Proceedings of the Winter Simulation Conference*, 569–577.
- Glynn, P. W. and W. Whitt (1992). The asymptotic efficiency of simulation estimators. *Operations Research* 40, 505–520.

- Gobet, E. and A. Kohatsu-Higa (2003). Computation of Greeks for barrier and lookback options using Malliavin calculus. *Electronic Communications in Probability* 8, 51–62.
- Goldberg, R. R. (1961). *Fourier transforms*. Cambridge Tracts in Mathematics and Mathematical Physics, No. 52. New York: Cambridge University Press.
- Goldman Sachs (1994). Valuing convertible bonds as derivatives. Available from http://www.ederman.com/new/docs/gs-valuing_convertibles.pdf. Quantitative Strategies Research Notes.
- Gurland, J. (1948). Inversion formulae for the distribution of ratios. *Annals of Mathematical Statistics* 19, 228–237.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Higham, D. J. and X. Mao (2005). Convergence of Monte Carlo simulations involving the mean-reverting square root process. *Journal of Computational Finance* 8, 35–62.
- Hilberink, B. and L. C. G. Rogers (2002). Optimal capital structure and endogenous default. *Finance and Stochastics* 6, 237–263.
- Ho, T. S. Y. and D. M. Pfeffer (1996). Convertible bonds: model, value, attribution and analytics. *Financial Analysts Journal* 52, 35–44.
- Hughett, P. (1998). Error bounds for numerical inversion of a probability characteristic function. *SIAM Journal on Numerical Analysis* 35, 1368–1392.
- Hull, J. C. (2003). *Options, futures and other derivatives* (5th ed.). Upper Saddle River: Prentice Hall.
- Ingersoll, J. E. (1977a). A contingent-claims valuation of convertible securities. *Journal of Financial Economics* 4, 289–321.
- Ingersoll, J. E. (1977b). An examination of corporate call policies on convertible securities. *Journal of Finance* 32, 463–478.

- Ingersoll, J. E. (1987). *Theory of Financial Decision Making: Studies in Financial Economics*. Savage: Rowman & Littlefield.
- Ju, N. (2002). Pricing Asian and basket options via Taylor expansion. *Journal of Computational Finance* 5, 79–103.
- Kaas, R., J. Dhaene, and M. J. Goovaerts (2000). Upper and lower bounds for sums of random variables. *Insurance: Mathematics and Economics* 27, 151–168.
- Kallsen, J. (2006). *A didactic note on affine stochastic volatility models*. In From Stochastic Calculus to Mathematical Finance: The Shiryaev Festschrift, Y. Kabanov, R. Liptser and J. Stoyanov, eds. New York: Springer.
- Kallsen, J., J. Muhle-Karbe, N. Shenkman, and R. Vierthauer (2009). *Discrete-time variance-optimal hedging in affine stochastic volatility models*. In Alternative Investments and Strategies, R. Kiesel, M. Scherer and R. Zagst, eds. Singapore: World Scientific.
- Kallsen, J. and A. Pauwels (2009). Variance-optimal hedging for time-changed Lévy processes. Available from <http://www.numerik.uni-kiel.de/~jk/personen/kallsen.html>. To appear in Applied Mathematical Finance.
- Kemna, A. and T. Vorst (1990). A pricing method for options based on average asset values. *Journal of Banking and Finance* 14, 113–129.
- King, R. (1986). Convertible bond valuation: an empirical test. *Journal of Financial Research* 9, 53–69.
- Kohatsu-Higa, A. and K. Yasuda (2009). *A review of some recent results of Malliavin calculus and its applications*. In Advanced Financial Modelling, H. Albrecher, W. Runggaldier and W. Schachermayer, eds. Berlin: de Gruyter.
- Koponen, I. (1995). Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process. *Physical Review E* 52, 1197–1199.
- Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science* 48, 1086–1101.

- Küchler, U. and M. Sørensen (1997). *Exponential families of stochastic processes*. New York: Springer.
- Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research* 2, 99–120.
- L’Ecuyer, P. and G. Perron (1994). On the convergence rates of IPA and FDC derivative estimators. *Operations Research* 42, 643–656.
- Leland, H. E. (1994). Corporate debt value, bond covenants, and optimal capital structure. *Journal of Finance* 49, 1213–1252.
- Leland, H. E. and K. B. Toft (1996). Optimal capital structure, endogenous bankruptcy and the term structure of credit spreads. *Journal of Finance* 51, 987–1019.
- Levy, E. (1992). Pricing European average rate currency options. *Journal of International Money and Finance* 11, 474–491.
- Lewis, A. L. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. Available from <http://www.optioncity.net/>. Working paper, OptionCity.net.
- Linetsky, V. (2004). Spectral expansions for Asian (average price) options. *Operations Research* 52, 856–867.
- Longstaff, F. A. and E. S. Schwartz (1995). A simple approach to valuing risky fixed and floating rate debt. *Journal of Finance* 50, 789–819.
- Longstaff, F. A. and E. S. Schwartz (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* 14, 113–147.
- Lord, R. (2006a). Partially exact and bounded approximations for arithmetic Asian options. *Journal of Computational Finance* 10, 1–52.
- Lord, R. (2006b). Pricing of baskets, Asians and swaptions in general models. Available from http://www.rogerlord.com/condition_and_conquer_lunteren.pdf. Presentation, 5th Winter school on Financial Mathematics, January 2006.

- 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, 1678–1705.
- Lord, R. and C. Kahl (2008). Complex logarithms in Heston-like models. Available from <http://ssrn.com/abstract=1105998>. To appear in *Mathematical Finance*.
- Lord, R., R. Koekkoek, D. van Dijk (2010). A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance* 10, 177–194.
- Lvov, D., A. Yigitbasioglu, and N. El Bachir (2004). Pricing convertible bonds by simulation. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=950213. ISMA Centre Discussion Papers in Finance, University of Reading.
- Madan, D. B. (2000). Pricing the risks of default. Working paper, University of Maryland.
- Madan, D. B., P. P. Carr, and E. C. Chang (1998). The variance gamma process and option pricing. *European Finance Review* 2, 79–105.
- Madan, D. B. and F. Milne (1991). Option pricing with V.G. martingale components. *Mathematical Finance* 1, 39–55.
- Madan, D. B. and E. Seneta (1990). The variance gamma (V.G.) model for share market returns. *Journal of Business* 63, 511–524.
- Madan, D. B. and M. Yor (2008). Representing the CGMY and Meixner Lévy processes as time changed Brownian motions. *Journal of Computational Finance* 12, 22–47.
- McConnell, J. J. and E. S. Schwartz (1986). LYON Taming. *Journal of Finance* 41, 561–576.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- Merton, R. C. (1974). On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* 29, 449–470.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.

- Milevsky, M. and S. Posner (1998). Asian options, the sum of lognormals, and the reciprocal gamma distribution. *Journal of Financial and Quantitative Analysis* 33, 409–422.
- Nielsen, J. A. and K. Sandmann (2003). Pricing bounds on Asian options. *Journal of Financial and Quantitative Analysis* 38, 449–473.
- Norberg, R. (2006). Dynamic greeks. *Insurance: Mathematics and Economics* 39, 123–133.
- Poirot, J. and P. Tankov (2006). Monte Carlo option pricing for tempered stable (CGMY) processes. *Asia-Pacific Financial Markets* 13, 327–344.
- Rabiner, L. R., R. W. Schafer, and C. M. Rader (1969). The chirp z-transform algorithm and its application. *Bell Systems Technical Journal* 48, 1249–1292.
- Raible, S. (2000). Lévy processes in finance: theory, numerics and empirical facts. PhD thesis, Institut für Mathematische Stochastik, Albert-Ludwigs-Universität, Freiburg.
- Rogers, L. C. G. and Z. Shi (1995). The value of an Asian option. *Journal of Applied Probability* 32, 1077–1088.
- Rosiński, J. (2001). *Series representations of Lévy processes from the perspective of point processes*. In Lévy processes - Theory and Application, O. E. Barndorff-Nielsen, T. Mikosch and S. I. Resnick, eds. Boston: Birkhauser.
- Schöbel, R. and J. Zhu (1999). Stochastic volatility with an Ornstein-Uhlenbeck process: an extension. *European Finance Review* 3, 23–46.
- Schoutens, W. (2003). *Lévy Processes in Finance: Pricing Financial Derivatives*. New York: Wiley.
- Shreve, S. and J. Večer (2000). Options on a traded account: vacation calls, vacation puts and passport options. *Finance and Stochastics* 4, 255–274.
- Stein, E. and J. Stein (1991). Stock-price distributions with stochastic volatility - an analytic approach. *Review of Financial Studies* 4, 727–752.
- Takahashi, A., T. Kobayashi, and N. Nakagawa (2001). Pricing convertible bonds with default risk: a Duffie-Singleton approach. *Journal of Fixed Income* 11, 20–29.

- Talvila, E. (2001). Necessary and sufficient conditions for differentiating under the integral sign. *The American Mathematical Monthly* 108, 544–548.
- Tsiveriotis, K. and C. Fernandes (1998). Valuing convertible bonds with credit risk. *Journal of Fixed Income* 8, 95–102.
- Turnbull, S. M. and L. M. Wakeman (1991). A quick algorithm for pricing European average options. *Journal of Financial and Quantitative Analysis* 26, 377–389.
- van Haastrecht, A., R. Lord, A. Pelsser, and D. Schrager (2009a). Pricing long-maturity equity and FX derivatives with stochastic interest rates and stochastic volatility. *Insurance: Mathematics and Economics* 45, 436–448.
- van Haastrecht, A., R. Lord, and A. Pelsser (2009b). Monte Carlo pricing in the Schöbel-Zhu model and its extensions. Available from <http://www.rogerlord.com/szmc.pdf>. Working paper, Delta Lloyd, University of Amsterdam, Cardano and Maastricht University.
- Vanmaele, M., G. Deelstra, J. Liinev, J. Dhaene, and M. J. Goovaerts (2006). Bounds for the price of discretely sampled arithmetic Asian options. *Journal of Computational and Applied Mathematics* 185, 51–90.
- Vašíček, O. A. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics* 5, 177–188.
- Večeř, J. (2001). A new PDE approach for pricing arithmetic average Asian Options. *Journal of Computational Finance* 4, 105–113.
- Večeř, J. (2002). Unified Asian pricing. *Risk* 15, 113–116.
- Večeř, J. and M. Xu (2004). Pricing Asian options in a semimartingale model. *Quantitative Finance* 4, 170–175.
- Weron, R. (1996). On the Chambers-Mallows-Stuck method for simulating skewed stable random variables. *Statistics & Probability Letters* 28, 165–171.
- Wong, H. Y. and Y. L. Cheung (2004). Geometric Asian options: valuation and calibration with stochastic volatility. *Quantitative Finance* 4, 301–314.

- Zabolotnyuk, Y., R. A. Jones, and C. Veld (2009). An empirical comparison of convertible bond valuation models. Available from <http://ssrn.com/abstract=994805>. Working paper, Carleton University, Simon Fraser University and University of Stirling.
- Zhang, J. E. (2001). A semi-analytical method for pricing and hedging continuously sampled arithmetic average rate options. *Journal of Computational Finance* 5, 59–79.
- Zhou, C. (1997). A jump-diffusion approach to modeling credit risk and valuing defaultable securities. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=39800. Working paper, Peking University.
- Zvan, R., P. A. Forsyth, and K. R. Vetzal (1998). A general finite element approach for PDE option pricing models. Available from <http://citeseerx.ist.psu.edu/viewdoc>. Working paper, University of Waterloo.
- Zvan, R., P. A. Forsyth, and K. R. Vetzal (2001). A finite volume approach for contingent claims valuation. *IMA Journal of Numerical Analysis* 21, 703–731.