On the calibration of Lévy option pricing models

IJH Visagie

20313071

Thesis submitted for the degree *Philosophiae Doctor* in Risk Analysis at the Potchefstroom Campus of the North-West University

Promoter: Prof F Lombard

The financial assistance of the National Research Foundation (NRF) and the DST-NRF Centre of Excellence in Mathematical and Statistical Sciences (CoE-MaSS) is hereby acknowledged. Opinions expressed and conclusions arrived at are those of the author and are not necessarily to be attributed to the NRF or the CoE-MaSS.

October 2015

It all starts here ™

NORTH-WEST UNIVERSITY ®
YUNIBESITI YA BOKONE-BOPHIRIMA
NOORDWES-UNIVERSITEIT

Contents

Acknowledgements

The financial assistance of the National Research Foundation (NRF) and the DST-NRF Centre of Excellence in Mathematical and Statistical Sciences (CoE-MaSS) is hereby acknowledged. Opinions expressed and conclusions arrived at are those of the author and are not necessarily to be attributed to the NRF or the CoE-MaSS.

I would like to thank a number of people for their support and their contribution to my studies:

- My wife, Jeanette, for her love, enthusiasm and unwavering support.
- My parents for their love, encouragement and interest. I owe all of what I am to you.
- My brothers for their support.
- My colleagues for their help with all of my problems, big and small. A special word of thanks to James Allison, Leonard Santana and Charl Pretorius for always being willing to help.
- Fanie Terblance for his help with the optimisation problems.
- Erika Fourie and Mari van Wyk for listening to all of my complaints and for their support.
- Prof Hennie Venter, Prof Jan Swanepoel and Prof Riaan de Jongh for their interest and their help.
- My examiners for helpful comments.
- My promoter, Prof Freek Lombard, for teaching me to think and for inspiring me. Prof, working with you has been a great honour.

Summary

In this thesis we consider the calibration of models based on Lévy processes to option prices observed in some market. This means that we choose the parameters of the option pricing models such that the prices calculated using the models correspond as closely as possible to these option prices. We demonstrate the ability of relatively simple Lévy option pricing models to nearly perfectly replicate option prices observed in financial markets. We specifically consider calibrating option pricing models to barrier option prices and we demonstrate that the option prices obtained under one model can be very accurately replicated using another. Various types of calibration are considered in the thesis.

We calibrate a wide range of Lévy option pricing models to option price data. We consider exponential Lévy models under which the log-return process of the stock is assumed to follow a Lévy process. We also consider linear Lévy models; under these models the stock price itself follows a Lévy process. Further, we consider time changed models. Under these models time does not pass at a constant rate, but follows some non-decreasing Lévy process. We model the passage of time using the lognormal, Pareto and gamma processes. In the context of time changed models we consider linear as well as exponential models.

The normal inverse Gaussian ($N \circ IG$) model plays an important role in the thesis. The numerical problems associated with the $N \circ IG$ distribution are explored and we propose ways of circumventing these problems. Parameter estimation for this distribution is discussed in detail.

Changes of measure play a central role in option pricing. We discuss two well-known changes of measure; the Esscher transform and the mean correcting martingale measure. We also propose a generalisation of the latter and we consider the use of the resulting measure in the calculation of arbitrage free option prices under exponential Lévy models.

Key words: Calibration - option pricing - Lévy processes - normal inverse Gaussian distribution - lognormal distribution - Pareto distribution - generalised mean correcting martingale measure, barrier options.

List of acronyms and abbreviations

AAE	average absolute error
ARE	average relative error
DB	digital barrier
$DMLE$	direct maximum likelihood estimator
DOB	down-and-out barrier
$ECFEA$	empirical characteristic function estimator with an absolute value weight function
$ECFEN$	empirical characteristic function estimator with a normal density weight function
exp	exponential
GMCMM	generalised mean correcting martingale measure
$IMLE$	indirect maximum likelihood estimator
LEMM	locally equivalent martingale measure
lin	linear
MCMM	mean correcting martingale measure
$N \circ IG$	normal inverse Gaussian
$RMSE$	root mean square error
$SN \circ IG$	standard normal inverse Gaussian
Std dev	Standard deviation
TC	Time changed
UOB	up-and-out barrier

Chapter 1

Introduction

This chapter serves as an introduction to the thesis. Section 1.1 is concerned with the motivation for the study as well as a brief overview of the concepts considered. In Section 1.2 we state the objectives of the thesis and discuss the research questions that we attempt to answer. Section 1.3 shows an outline of the chapters to follow and Section 1.4 provides the details of the hardware and software used for numerical calculations throughout the thesis.

1.1 Motivation and overview

A financial derivative is a product that derives its value from some underlying asset, usually a stock or an index comprising multiple stocks. The global derivatives market is vast and expanding. Recent times has seen a lot of research done relating to the pricing of derivatives as well as the quantification of the risk that companies face because of their holdings in these product. The unpredictability of the derivatives market prompts further research into these fields.

The market in financial options make up a large portion of the global derivatives market. Many different models have been proposed for the calculation of option prices, certainly the most famous of these is the Black-Scholes model. Under this model the stock price is assumed to follow an exponential Brownian motion. However, it is well known that many of the assumptions of the Black-Scholes model do not hold in practice. As a result, many generalisations of and alternatives to this model have been proposed.

One of the Black-Scholes assumptions that do not hold in practice is the normality of the

log-returns of the stock price. A generalisation of the Black-Scholes model addressing this problem is obtained by replacing the assumption that the stock price follows an exponential Brownian motion by the assumption that the stock price follows an exponential Lévy process with jumps. These models are called exponential Lévy option pricing models. A specific example of this type of model used extensively throughout the thesis is the exponential $N \circ IG$ model.

We consider various option pricing models based on Lévy processes, including the exponential Lévy models mentioned above as well as models under which the increments in the stock price follow a Lévy process; we refer to these models as linear Lévy models. Another use of Lévy processes in option price models can be found in time change modelling; these models are explained below.

A second assumption of the Black-Scholes model contradicted by empirical analysis of financial data is that the volatility of stock prices are constant over time. Empirical evidence suggests that stock prices go through times of higher activity and times of lower activity. This property can be modelled using option pricing models under which time does not evolve at a constant rate (these models are called time changed models). Rather time is modelled using a non-decreasing stochastic process. In this thesis we model time using three different Lévy processes; the lognormal, Pareto and gamma processes.

One of the Lévy processes often used to model financial returns is the $N \circ IG$ process. In this thesis we highlight some numerical difficulties encountered when using this process in the context of option pricing and we propose ways of solving these numerical problems.

In order to use the $N \circ IG$ process for the calculation of option prices we need to be able to estimate the parameters of the corresponding distribution based on sample data. We consider various estimators for the parameters of this distribution in some detail; we consider the maximum likelihood estimator as well as an empirical characteristic function estimator.

Changes of probability measure form a central part of option pricing theory. In this thesis we consider two well-known changes of measure together with their uses in option

pricing; the Esscher transform and the mean correcting martingale measure. We also propose a generalisation of the mean correcting martingale measure. We consider the use of this change of measure in option pricing and we provide an economic interpretation of this measure change.

The main aim of this thesis is to evaluate calibration as a method for calculating the parameters used in various Lévy option pricing models. Calibrating a model to a set of observed option prices entails choosing the parameters of the model by minimising some distance measure between the option prices calculated using the model and the observed option prices. We discuss various types of calibration as well as the restrictions placed on each of the types of calibration.

Two option types play an important role in the thesis; European options and barrier options. The payoff of a European option on a given stock depends on the terminal value of the stock only, while the payoff of a barrier option is path dependent. The calculation of the prices of European options and the associated calibrations are considered in detail in the first part of the thesis up to the end of Chapter 5. From Chapter 6 the focus shifts to calibrations relating to barrier options.

The main contribution of this thesis lies therein that we demonstrate that it is possible to replicate with one model the barrier option prices calculated using another model with a high degree of accuracy. We shall refer to such models as interchangeable. Furthermore, we show that estimating the parameters of these models from financial time series data also leads to option prices that correspond closely under the various models.

1.2 Objectives

In this thesis we aim to:

- Demonstrate that one model can be used to replicate the barrier option prices obtained using another model with a high degree of accuracy.
- Discuss different types of calibration and consider the implications of the restrictions

placed on these calibration types.

- Suggest ways to work around the numerical problems encountered when using the $N \circ IG$ distribution in option pricing.
- Provide a comparison between different estimators that can be used in order to calculate parameter estimates for the $N \circ IG$ distribution based on sample data.
- Generalise the mean correcting martingale measure and consider the possible role that this generalised measure can play in option pricing.
- Demonstrate the use of time changed option pricing models under which time is assumed to follow a lognormal or Pareto process.

1.3 Thesis outline

Chapter 1: Introduction

The reader is provided with a motivation for the analyses to follow as well as a general overview of the results obtained in the thesis.

Chapter 2: Option pricing and model calibration

We give a short overview of the option pricing literature relevant to the analyses done in the thesis. Special attention is paid to considerations relating to the calibration of option pricing models to observed option prices.

Chapter 3: Lévy processes and infinitely divisible distributions

This chapter gives a brief overview of Lévy processes and infinitely divisible distributions. The $N \circ IG$ distribution is considered and special attention is paid to parameter estimation for this distribution. In this chapter we also discuss a class of infinitely divisible distributions known as generalised Γ-convolutions; this class of distributions include the lognormal and Pareto distributions.

Chapter 4: Martingale measures and exponential Lévy models

This chapter provides a discussion of various changes of measure. Here we consider the requirements placed on changes of measure in order to give rise to arbitrage free option prices. We consider two well-known martingale measures; the Esscher transform and the mean correcting martingale measure. We also propose a generalisation of the mean correcting martingale measure.

Chapter 5: Numerical calibration results for European options

In this chapter numerical results are presented for the calculation of European option prices using exponential Lévy models, including the results obtained using different types of calibration. This chapter includes a discussion of the observed option prices used as well as an explanation of the procedure used to remove outlying option prices from the considered datasets.

Chapter 6: Time changed option pricing models

An overview of time changed option pricing models is provided. It is worth noting that we use time change models where the directing processes are the lognormal, Pareto and gamma processes. The first two of these processes are rarely if ever used in this context. This chapter also includes an overview of the first paper to use time changes in the context of financial modelling.

Chapter 7: Calibrating models to barrier option prices

We discuss the calculation of barrier option prices as well as the calibration of various models to these option prices. In this chapter we show that barrier option pricing models are interchangeable in the following sense; given a set of barrier option prices calculated under some model, it is possible to calibrate another model very accurately to these prices.

Chapter 8: Calibrations and estimations relating to barrier options

In this chapter we use exponential Lévy option pricing models and we show that it is not only possible to calibrate the models considered to option prices with a high degree of accuracy, it is also possible to achieve this result using estimation procedures.

1.4 Hardware and software used

Calibrating a model to observed option prices can be a very computationally expensive task and typically requires large amounts of computer time. As a result, we are interested in comparing different calculation methods in an attempt to speed up the calibration procedures. In our search for numerically efficient calculation methods we report the times taken to calculate certain results obtained in the thesis. In order to objectively rate the speed at which these calculations are performed one requires knowledge of the computer hardware and software used. The results shown in the thesis were obtained using Matlab 2012b on a 64 bit Windows 7 operating system with an Intel Core i7 CPU @ 2.80 GHz (8 CPUs) with 4 GB of RAM.

Chapter 2

Option pricing and model calibration

This chapter introduces some of the concepts relating to the calculation of option prices as well as calibration used throughout the thesis. No new results are presented in this chapter.

This chapter consists of four sections. Section 2.1 contains a brief overview of option pricing. In Section 2.2 various methods used for the calculation of European option prices are discussed and compared. Section 2.3 explains what is meant by the calibration of option pricing models and provides a motivation for the use of calibration. In Section 2.4 we discuss Schoutens *et al.* (2004:66-78); a paper that provides some interesting insights into the calibration of option pricing models to observed option prices.

2.1 An overview of option pricing

This section provides a brief overview of some of the major results in option pricing theory relevant to the thesis. We define the market used throughout together with assets and portfolios available in this market. We define European call and put options; these option types play important roles in the thesis. The analyses in the remainder of the thesis are not limited to these option types; we introduce barrier options in Section 7.1. Below we define an arbitrage opportunity and we consider the conditions for an arbitrage free market. The fundamental theorem of asset pricing is also provided.

The discussion below is based on Björk (2011:61-68) and Cont and Tankov (2004:291-299), the interested reader are referred to these texts as well as Harrison and Pliska (1981:215-260) for a more technical exposition.

2.1.1 Definition of a financial market

We assume a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$, where $\mathcal{F}_t$ is a filtration satisfying the usual conditions; i.e. $\mathcal{F}_t$ is right-continuous and non-decreasing. We assume that there exists a time T_* such that $\mathcal{F}_{T_*} = \mathcal{F}$. On this space we define two processes B_t and S_t. Economically B_t can be interpreted as the value of a risk free bond at time t with $B_0 = 1$. Throughout the thesis we assume a constant interest rate and continuous compounding. Therefore,

$$B_t = \exp(rt), \tag{2.1}$$

where $r \geq 0$ is the known constant risk free interest rate. The price of the bond is used for discounting purposes. S_t is a semi-martingale generating $\mathcal{F}_t$ ($\mathcal{F}_t = \sigma(S_t)$). The process S_t represents a stock price (the value of a single share in a publicly traded company) or an index comprising multiple stocks. In the sequel our use of the word stock can be replaced by the word index. The probability measure P is known as the objective probability measure; this measure determines the distribution of S_t.

Throughout the thesis all of the models used are parameterised daily. This means that S_t denotes the stock price at the end of the t^{th} business day and that r is the continuously compounded daily interest rate. We assume that a year contains 252 business days.

We make a number of standard simplifying assumptions. All assets in this market are perfectly divisible; an investor is not restricted to hold a whole number of assets but can hold fractional assets as well. The market is perfectly liquid; any asset can be immediately bought or sold at the current market price. We assume a frictionless market where no trading costs are incurred when buying or selling any asset. Short selling of assets is allowed, meaning that investors are allowed to hold negative quantities of assets.

A portfolio is a combined holding of assets. Mathematically a portfolio consisting of the bond B and the stock S is a 2 dimensional predictable process,

$$h_t = (h_t^B, h_t^S),$$

where h_t^B and h_t^S respectively represent the number of units of the bond and the stock held by an investor at time t. The value of the portfolio h_t is

$$V_t^h = h_t^B B_t + h_t^S S_t. \tag{2.2}$$

A portfolio strategy is said to be self-financing if no additional funds are entered into or withdrawn out of the portfolio after the initial investment. Therefore a portfolio strategy is self-financing if

$$dV_t^h = h_t^B dB_t + h_t^S dS_t. \tag{2.3}$$

2.1.2 European options

Hull (2009:1) provides the following definition of a financial derivative.

Definition 1 *A derivative can be defined as a financial instrument whose value depends on (or derives from) the values of other, more basic, underlying variables. Very often the variables underlying derivatives are the prices of traded assets.*

The class of financial derivatives considered in this thesis is options. Two option types that play an important role in the thesis are the European call option and the European put option. Cont and Tankov (2004:355) defines a European call option as follows.

Definition 2 *A European call option on an asset S with maturity date T and strike price K is defined as a contingent claim that gives its holder the right (but not the obligation) to buy the asset at a fixed date T for a fixed price K.*

If the price of the underlying asset at time T is more than K, the holder will choose to exercise the option and buy the stock for less than its market value. Since the asset can be immediately sold for its current market value, the holder of the option then realises a profit of $S_T - K$. If the price of the asset at time T is less than K the holder will choose not to exercise the option. As a result, the payoff associated with a European call is given

by

$$(S_T - K)^+ = \begin{cases} S_T - K, & S_T > K \\ \\ 0, & S_T \leq K. \end{cases}$$

A European put option can be defined as follows.

Definition 3 *A European put option on an asset S with maturity date T and strike price K is defined as a contingent claim that gives its holder the right (but not the obligation) to sell the asset at a fixed date T for a fixed price K.*

The payoff function of a European put option is

$$(K - S_T)^+ = \begin{cases} 0, & S_T > K \\ \\ K - S_T, & S_T \leq K. \end{cases}$$

A financial derivative, such as an option, is itself an asset with a price process. We denote the price of a given option at time t by π_t. The introduction of the asset π_t extends the market to contain three assets; B_t, S_t and π_t. The definition of a portfolio can now be extended as follows:

$$h_t = (h_t^B, h_t^S, h_t^\pi),$$

where h_t^π represents the number of units of the option held by an investor at time t. The value of the portfolio given in (2.2) can be similarly extended by adding $h_t^\pi \pi_t$, while the definition of a self-financing portfolio in (2.3) is similarly extended by adding $h_t^\pi d\pi_t$. Multiple options can be available in the market.

Throughout the thesis we are mainly interested in the spot prices of options (the prices of the options at time 0). In the sequel we omit the subscript 0 in the spot price of a given option, meaning that the price at time 0 of a given option is denoted by π.

2.1.3 Arbitrage and the fundamental theorem of asset pricing

Björk (2011:63) provides the following definition of an arbitrage opportunity.

Definition 4 *A portfolio strategy h constitutes an arbitrage opportunity on the time interval $[0,T]$ if:*

1. *h is self-financing.*
2. *The initial value of h is zero; $V_0^h = 0$.*
3. *$P(V_T^h \geq 0) = 1$ and $P(V_T^h > 0) > 0$.*

This definition is general and holds for any market irrespective of the different asset classes available in the market. If there exists a portfolio strategy constituting an arbitrage opportunity then there exists a possibility of realising a profit without any initial capital or exposure to risk. Since this contradicts economic theory a realistic market model should contain no arbitrage opportunities.

Arbitrage is closely linked to the concept of a locally equivalent martingale measure (LEMM), which we define below. Two probability measures, P and $\tilde{P}$, are said to be equivalent (denoted $P \sim \tilde{P}$) if

$$P(A) = 0 \iff \tilde{P}(A) = 0,$$

for all $A \in \mathcal{F}$. Let $P_n = P|\mathcal{F}_n$ and $\tilde{P}_n = \tilde{P}|\mathcal{F}_n$ be the restrictions of the probability measures to $\mathcal{F}_n$. Two probability measures are said to be locally equivalent (denoted $P \overset{loc}{\sim} \tilde{P}$) if the following holds for all $n = 1, 2, ...$

$$P_n(A) = 0 \iff \tilde{P}_n(A) = 0,$$

for all $A \in \mathcal{F}_n$. We can now define a LEMM.

Definition 5 *A probability measure Q is said to be a locally equivalent martingale measure (LEMM) with respect to P if the following two conditions hold:*

1. *Q is locally equivalent to P.*
2. *$S_t/B_t = e^{-rt}S_t$ forms a Q-martingale.*

We now state the first fundamental theorem of asset pricing.

Theorem 6 *A market is arbitrage free if and only if there exists a LEMM.*

For a proof in discrete time, see Shiryaev (2003:410-432). Many extensions to this theorem have been proven. For a discussion and proof of the theorem in its most general form the interested reader is referred to Delbaen and Schachermayer (2005:149-190).

A well-known result in option pricing theory is that if Q is a LEMM, then the price of an option with payoff function X and time to maturity T can be calculated as

$$\pi = e^{-rT} E^Q [X] . \tag{2.4}$$

If Q is used in the calculation of option prices we may refer to Q as the pricing measure.

Using (2.4) to calculate option prices entails changing probability measures from P to Q. When changing measure from P to Q the possible price paths of S_t do not change; all events that are impossible under P are impossible under Q, and vice versa. However, changing the probability measure changes the probability of uncertain events. As a result, the processes governed by P and Q have the same possible price paths but different statistical properties. Measure changes play an important role in arbitrage free option pricing theory. We consider measure changes in more detail in Chapter 4.

2.1.4 The Black-Scholes option pricing model

Certainly the most famous option pricing model is the Black-Scholes model. Black and Scholes (1973:637-654) uses differential equations and a replicating portfolio argument to find an option pricing formula. We now briefly discuss this model.

In the framework of the Black-Scholes model the market is assumed to consist of two assets; a bond B_t and a stock S_t. The dynamics of the bond process are given by (2.1). Under the objective probability measure P, the stock price is driven by an exponential Brownian motion,

$$S_t = S_0 \exp\left(\mu t + \sigma W_t\right), \tag{2.5}$$

with $\mu \in \mathbb{R}$, $\sigma > 0$ and W_t a standard Brownian motion. Note that, since a Brownian motion is a Lévy process, the Black-Scholes model is an example of an exponential Lévy option pricing model.

In the Black-Scholes market there exists a unique LEMM. Changing measure from P to the LEMM in this model entails setting the drift of the Brownian motion to

$$\mu = r - \frac{\sigma^2}{2}. \tag{2.6}$$

This means that the arbitrage free price of each option in this market is uniquely determined by σ. The price of a European option is an increasing function of σ. Therefore, it is possible (for a given call or put option) to calculate the unique value of σ that equates the price calculated using this model with the market price of the option. The value of σ that equates these two prices is known as the implied volatility of the option.

In the remainder of the thesis we consider many different option pricing models. Under the majority of these models there exist multiple LEMMs. As a result, there are multiple arbitrage free prices for a single option. Calibration can be used to choose a suitable LEMM, this is explained in Chapter 2.3.

2.2 The calculation of European option prices

Below we consider two methods of calculating European option prices: direct numerical integration and a Fourier method introduced in Attari (2004). Numerical results are included for the prices of European call options under the Black-Scholes model and the exponential $N \circ IG$ model. Under this model the increments in the logarithm of the stock price is

assumed to be independent and identically distributed random $N \circ IG$ variables.

Definition 7 *A random variable $X : \Omega \to \mathbb{R}$ is said to follow a $N \circ IG$ distribution with parameter set $\theta = (\alpha, \beta, \mu, \delta)$ if it has density*

$$f(x;\theta) = \frac{\alpha\delta}{\pi} \exp\left(\delta\sqrt{\alpha^2 - \beta^2} + \beta(x-\mu)\right) \frac{K_1\left(\alpha\sqrt{\delta^2 + (x-\mu)^2}\right)}{\sqrt{\delta^2 + (x-\mu)^2}}, \tag{2.7}$$

where $\alpha > 0$, $|\beta| < \alpha$, $\mu \in \mathbb{R}$, $\delta > 0$ and K_1 denotes the modified Bessel function of the third kind with index 1.

A definition of this Bessel function can be found in Schoutens (2003:148). The characteristic function of the $N \circ IG(\theta)$ distribution is

$$\phi(t;\theta) = \exp\left(i\mu t - \delta\left(\sqrt{\alpha^2 - (\beta + it)^2} - \sqrt{\alpha^2 - \beta^2}\right)\right). \tag{2.8}$$

The $N \circ IG$ distribution is considered in detail in Chapter 3. Both the Black-Scholes and the exponential $N \circ IG$ model play an important role in this thesis.

Implementation of the direct numerical integration method requires knowledge of the density function of the price process, or some functional thereof (usually $\log(S_t)$). Similarly, Attari's method requires knowledge of the characteristic function. If the density and characteristic function of S_t (and $\log(S_t)$) are unknown the methods described below cannot be used. In these cases Monte Carlo simulation can be used to estimate the prices of the options; this method is used in Chapters 7 and 8.

The arbitrage free prices of options can be calculated as expected values taken with respect to some LEMM Q, see (2.4). The arbitrage free price of a European call option (with strike price K and time to maturity T) calculated with respect to Q is given by

$$\pi = e^{-rT} E^Q\left[(S_T - K)^+\right]. \tag{2.9}$$

Below we calculate (2.9) using the two methods mentioned above.

When calibrating models to observed option prices a large number of prices have to be calculated. As a result, we are interested in the computer time necessary to compute the option prices using the two models. The time required to calculate the option prices using the two methods are reported in the numerical examples.

The analysis below is done in order to decide between two competing methods for the calculation of option prices used later in the thesis. The choice is based on the time required for the calculation of option prices using these methods. A secondary aim of this analysis is to serve as a confirmation that the prices of European call options are calculated correctly in the remainder of the thesis.

2.2.1 Direct numerical integration

Let f_t denote the density of $\log(S_t)$ under the probability measure Q. If f_T is known analytically, (2.9) can be evaluated as follows:

$$\pi = e^{-rT} \int_{l}^{\infty} (S_0 e^x - K) f_T(x)\, dx, \tag{2.10}$$

where $l = \log(K/S_0)$. Numerical integration software can be used to compute (2.10).

European put option prices can be calculated similarly.

2.2.2 The use of Fourier transforms

Carr and Madan (1999:61-73) show that fast Fourier transforms can be used to calculate option prices. A numerically efficient simplification to the method developed by Carr and Madan is presented in Attari (2004). The application of this method is demonstrated below.

Let ϕ_t represent the characteristic function of $\log(S_t)$;

$$\phi_t(\omega) = E\left[\exp\left(i\omega \log(S_t)\right)\right].$$

Attari (2004) shows that the price of a European call option can be calculated as follows:

$$\begin{aligned}\pi_0 &= S_0 - e^{-rT}K\left(\frac{1}{2}+\right.\\ &\left.\frac{1}{\pi}\int_0^{\infty}\frac{\left(R_T(\omega)+\frac{I_T(\omega)}{\omega}\right)\cos(\omega l)+\left(I_T(\omega)-\frac{R_T(\omega)}{\omega}\right)\sin(\omega l)}{1+\omega^2}d\omega\right), \quad (2.11)\end{aligned}$$

where $R_T(\omega) = \mathrm{Re}(\phi_T(\omega))$, $I_T(\omega) = \mathrm{Im}(\phi_T(\omega))$ and $l = \log\left(\frac{Ke^{-rT}}{S_0}\right)$. We denote the option price in (2.11) by π_0 in order to distinguish it from the irrational number π in the right hand side of this equation.

The price of a European put option can be calculated using the put-call parity, see Hull (2009:208-211). If we denote the price of a European put option (with strike price K and time to maturity T) by π_0^p then

$$\pi_0^p = \pi_0 - S_0 + e^{-rT}K,$$

where π_0 is calculated using (2.11).

The formula in (2.11) implicitly assumes that the discounted stock price forms a Q-martingale. If this is not the case, then (2.11) cannot be used to calculate the right hand side in (2.9).

2.2.3 Numerical comparison of calculation methods

The two models used in the numerical analyses below are the Black-Scholes model and the exponential $N \circ IG$ model. The exponential $N \circ IG$ model is obtained by replacing the Brownian motion in (2.5) by a $N \circ IG$ process. For the details of the $N \circ IG$ process and the exponential $N \circ IG$ model, see Sections 3.2 and 4.1 respectively. For each of the models considered we arbitrarily choose the parameter set used for option pricing, no special significance is attached to these parameter sets.

Schoutens (2003:155-156) reports the prices of 75 European call options on the Standard

and Poor 500 index that were available on 18 April 2002. The options have varying strike prices and times to maturity. The risk free interest rate on this date is taken to be 0.7% per annum, compounded continuously, and the index price on this day is reported to be \$1124.47.

The prices of the options are calculated using each of the two methods discussed and we report the average relative absolute differences between the prices,

$$R = \frac{1}{75} \sum_{j=1}^{75} \frac{\left| \pi^{N,j} - \pi^{A,j} \right|}{\pi^{N,j}},$$

where $\pi^{N,j}$ and $\pi^{A,j}$ denote the prices of the j^{th} option calculated using direct numerical integration and Attari's method respectively. We report the time required to calculate the option prices using each of the two methods.

The Black-Scholes model

Consider the prices of the options under the Black-Scholes model with $\sigma = 0.01$. Using (2.6) μ is set to -2.22×10^{-5}. Figure 2.1 shows the prices calculated using direct numerical integration as circles and those calculated using Attari's method as crosses.

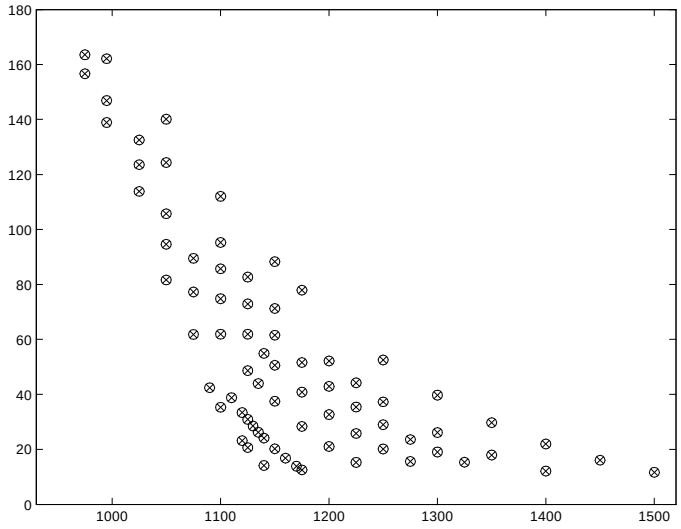

Figure 2.1: Comparison of calculation methods used for European option pricing.

Figure 2.1 shows that the prices calculated using the two methods are visually indistinguishable. The value of R is calculated to be $\$7.02 \times 10^{-12}$ in this example. Therefore, for practical purposes the prices obtained using the two methods are equal. Using direct numerical integration the 75 option prices above were calculated in 0.234 seconds, while calculating the option prices using Attari's method required 0.047 seconds.

The exponential $N \circ IG$ model

The exponential $N \circ IG$ model with parameter set $(\alpha, \beta, \mu, \delta) = (20, -15, 0.002, 0.002)$ is considered next. As before the option prices calculated using the two methods are visually indistinguishable. Therefore, we do not include a graph of the calculated prices here. The value of R is calculated to be $\$6.075 \times 10^{-9}$ in this example. As was the case in the Black-Scholes model, we conclude that the difference between the prices are not of practical importance.

Using direct numerical integration to calculate the 75 option prices required 0.14 seconds, while calculating the option prices using Attari's method required only 0.047 seconds.

Conclusions

We conclude that there is no practically significant difference between the option prices calculated using the two methods. It was noted earlier that Attari's method implicitly assumes that $e^{-rt}S_t$ forms a martingale under the pricing measure Q. For reasons explained below this assumption does not always hold. In these cases we use direct numerical integration for the calculation of option prices. However, if $e^{-rt}S_t$ forms a martingale under the pricing measure, then we use Attari's method since it is the faster of the methods. The difference between the time required for the implementation of direct numerical integration and Attari's method (in the case of the exponential $N \circ IG$ model) is explained by the difference in the complexity of the integrands used in the two methods. Direct numerical integration requires the repeated evaluation of the density function given in (2.7). This density function contains a Bessel function which is numerically complicated to evaluate.

Attari's method relies on the evaluation of the, much simpler, characteristic function given in (2.8).

2.3 The calibration of option pricing models

In Section 2.2 we consider the calculation of option prices (under some model) given a certain parameter set. Below we consider the reverse situation; we are given a set of option prices and we need to obtain parameters corresponding to these option prices. This process is known as calibration.

Calibrating a given option pricing model to a set of observed option prices entails choosing the parameters of the model so that the option prices calculated using the model correspond as closely as possible to the observed option prices. This explanation is made more precise below.

Below we consider the need for calibration. We explain that in some option pricing models multiple LEMMs exist. In these cases the arbitrage free price of an option is not unique. Calibration can be used in order to choose a pricing measure used for the calculation of option prices. We also define some of the distance measures routinely used in financial modelling. This section concludes with a discussion of various types of calibration as well as the restrictions imposed on the resulting LEMM under these types of calibration.

2.3.1 The need for calibration

In financial modelling we assume that the asset price is governed by some probability measure P. P is not known, but can be estimated from historical price data on S_t. Following this step arbitrage free prices of the options available in the market can be calculated as the discounted expected value of the payoff taken with respect to a LEMM. When implementing this approach practically we are faced with a problem; the LEMM may not be unique. Below we discuss how this problem can be solved using calibration.

In Section 2.1.4 we discuss the Black-Scholes model. In this model there exists a single LEMM that can be used for arbitrage free option pricing. This measure is obtained by

changing the drift parameter of the Brownian motion model, see (2.6). Since there exists a single LEMM in this market the arbitrage free prices of options in this market are uniquely defined. Markets exhibiting this property are called complete. The reason for the existence of unique prices in this market is that the payoff of any option can be replicated using a portfolio not containing the option itself. The arbitrage free price of the option must then necessarily correspond to the value of this portfolio; any difference between the option price and the value of the portfolio will immediately lead to an arbitrage opportunity.

In complete markets, where the payoff of the option can be replicated by a portfolio of assets, the price of the option is completely determined by P. As a result, the prices of options in these markets do not provide any information not contained in P. In fact, options are redundant assets in these markets; the options can be removed from these markets without diminishing the number of portfolios with unique value functions that can be constructed.

The situation described above, where every option possesses a unique arbitrage free price, is the exception and not the rule in financial modelling. Consider the exponential $N \circ IG$ model. Under this model the LEMM (and therefore the price of a given option) is not uniquely defined. This is because it is not possible to replicate the payoff of an option in this market using a portfolio that does not contain the option itself. Markets in which multiple LEMMs exist are called incomplete. In the case of incomplete markets options are not redundant assets and option prices provide information that should be taken into consideration when choosing the pricing measure.

Consider an incomplete market containing n observed option prices. When calibrating a model to these option prices we aim to find a pricing measure Q such that

$$\pi^{O,j} = e^{-rT_j} E^Q \left[X_j \right], \tag{2.12}$$

for all $j = 1, 2, ..., n$, where T_j and X_j are the time to maturity and the payoff function of the j^{th} option respectively and $\pi^{O,j}$ is the observed price of the j^{th} option. In a market

consisting of multiple options it is usually not possible to find a Q such that (2.12) holds for all $j = 1, 2, ..., n$. The solution to the calibration problem is found by minimising some distance measure between the observed option prices and the option prices obtained using Q. We discuss various distance measures below.

Usually calibration is done using European call options because of the high level of liquidity normally associated with these options. However, there is no theoretical reason that other option types cannot be used for calibration purposes.

2.3.2 Distance measures

When calibrating a given option pricing model to observed option prices we endeavour to minimise some distance measure between the observed prices and the prices calculated under the model. Various distance measures are used in model calibration; we define three of the most popular below.

Consider a market with n options and denote by $\pi^{O,j}$ and $\pi^{M,j}$ the observed price and the price under the model respectively of the j^{th} option. The average absolute error (AAE) is defined as

$$AAE = \frac{1}{n} \sum_{j=1}^{n} \left| \pi^{O,j} - \pi^{M,j} \right|. \tag{2.13}$$

The root mean square error ($RMSE$) is given by

$$RMSE = \sqrt{\frac{1}{n} \sum_{j=1}^{n} \left(\pi^{O,j} - \pi^{M,j} \right)^2}.$$

The average relative error (ARE) is defined as

$$ARE = \frac{1}{n} \sum_{j=1}^{n} \frac{\left| \pi^{O,j} - \pi^{M,j} \right|}{\pi^{O,j}}.$$

The three distance measures defined above are routinely used in financial modelling. When calibrating a model the pricing measure obtained is a function of the distance mea-

sure used. Choosing between the three distance measures defined above comes down to personal preference. In the remainder of the thesis we perform calibrations with respect to the AAE. Our reason for choosing this distance measure is its simple interpretation; the AAE is the average amount (in the relevant currency) that the model used misprices the options considered. Details of the numerical implementation of the calibration algorithms used are given in Section 5.2.

2.3.3 Different types of calibration

As was explained previously, calibrating a model to a set of option prices entails minimising some distance measure between the observed option prices and those calculated using the model. The distance measure that we choose to minimise is the AAE, see (2.13). Mathematically this can be expressed as follows:

$$\begin{aligned} Q^{\star} &= \arg\min_{Q\in\mathcal{Q}} AAE\left(Q\right) \\ &= \arg\min_{Q\in\mathcal{Q}} \frac{1}{n}\sum_{j=1}^{n}\left|\pi^{O,j} - e^{-rT_j}E^{Q}\left[X_j\right]\right|, \end{aligned} \tag{2.14}$$

where $\mathcal{Q}$ is the set of possible pricing measures. Below we consider the restrictions placed on the set $\mathcal{Q}$.

Denote an individual element of $\mathcal{Q}$ by Q. The requirements placed on Q play an important role in the calibration process. Consider the two possible restrictions imposed on Q and their implications for arbitrage free option pricing; the martingaleness of $e^{-rt}S_t$ under Q and the local equivalence of Q to the objective probability measure P. The exposition below considers the necessity of these two restrictions while avoiding technical details. The discussion below is intended as a guide to aid an intuitive understanding of the restrictions placed on calibration.

In order to see the importance of the requirement that $e^{-rt}S_t$ forms a Q-martingale,

assume that Q does not satisfy this requirement. Assume that

$$e^{-rT}E^{Q}\left[S_T\right] > S_0. \tag{2.15}$$

If π is the price of a European call option with time to maturity T and strike price 0, then the price of the option with respect to Q is

$$\pi = e^{-rT}E^{Q}\left[S_T\right] > S_0.$$

At $t = 0$ the price of the option is greater than the price of the stock, but at $t = T$ the two assets will have the same price; S_T. An investor can realise an arbitrage profit by selling the option short and buying the stock. Reversing the inequality in (2.15) results in an arbitrage opportunity for an investor buying the option while selling the stock short. As a result, the only way to avoid arbitrage is by enforcing the requirement that the discounted value of the stock price forms a Q-martingale.

Next, consider the importance of the requirement that Q be locally equivalent to P. Assume that this condition is not met. Let A be an event such that the probability of A under P is strictly positive and the probability of A under Q is 0. Consider an option with payoff function given by $\mathbb{I}(A)$. Since the probability of A is 0 under Q, the price of the option is 0. However, there exists a positive probability of the event A occurring in the market. Therefore, buying the option at the market price of 0 leads to an arbitrage opportunity. In the reverse situation A is possible under Q and impossible under P. Therefore the option price is positive and the probability that A will occur in the market is 0. In this case an arbitrage profit can be realised by selling the option short.

From the arguments above we conclude that, in order to guarantee the absence of arbitrage in a given market, Q must be a LEMM. Therefore, in order to guarantee that the market is arbitrage free $\mathcal{Q}$ should be defined as the class of locally equivalent martingale measures. However, Fouque *et al.* (2001:34-37) points out that when a calibration procedure is employed in practice, the history of S_t is often completely ignored (or in some cases

the analyst might not be able to obtain historical data on S_t). This means that financial practitioners often do not estimate P and that the requirement that the pricing measure be locally equivalent to the objective measure is discarded. A possible reason for this is that in some instances the calibration procedure is significantly simplified when the local equivalence condition is ignored. This methodology is not restricted to practitioners; often when a new model is proposed in the literature the model is calibrated to option prices while ignoring the price history of the stock.

When using a calibration procedure under which the local equivalence requirement is ignored $\mathcal{Q}$ is the class of measures such that $e^{-rt}S_t$ forms a Q-martingale for all $Q \in \mathcal{Q}$. In the sequel we refer to a calibration procedure satisfying the martingaleness requirement (but not the requirement of local equivalence) as martingale restricted calibration.

A second method of calibration entails finding $Q^\star$ in (2.14) by minimising the AAE, ignoring both the local equivalence requirement and the martingaleness requirement. We shall refer to this type of calibration as full calibration.

Finding $Q^\star$ in (2.14) is generally a non-convex optimisation problem containing multiple local minima. As a result, the solution that is obtained is sensitive to the starting values used by the optimisation algorithm. In an attempt to ensure that the optimisation algorithm converges to an acceptable solution we consider multiple possible starting values for each calibration. This is discussed in more detail in Section 3.3.3.

2.4 An overview of Schoutens *et al.* (2004:66-78)

An interesting paper relating to the calibration of option pricing models by Schoutens, Simons and Tistaert was published in 2004. The title of this paper is: A perfect calibration! Now what? The main results of the thesis are presented in Chapters 7 and 8. The results obtained in these chapters are obtained using techniques similar to those used in Schoutens *et al.* (2004:66-78). Below we briefly discuss this paper.

The paper demonstrates, by way of a numerical example, that various option pricing models can be calibrated very accurately to observed European call option prices. However,

when the calibrated models are used to calculate the prices of path dependent options (known as exotic options) the resulting prices differ substantially.

2.4.1 Calibration method used and results obtained

Various models are considered in the paper; two of which, the Heston stochastic volatility model and the time changed exponential $N \circ IG$ model, are used later in the thesis, see Section 7.2. Each of the models are calibrated to a set of observed European call option prices in turn. For each of the models the resulting option prices correspond very closely to the observed option prices. In the paper the authors use martingale restricted calibration with a mean correcting argument. This means that the calibration procedure allows each of the parameters in the model to vary freely, with the exception of the location parameter. This parameter is set to the value required for the discounted stock price to form a martingale. The distance measure minimised in the calibration procedure is the $RMSE$.

The word "perfect" in the title of the article might be a slight exaggeration since there are small differences between the observed option prices and the prices of the options obtained using the various models. However, from the results obtained in the paper it is clear that the models considered are able to mimic the behaviour of the option prices nearly perfectly. The results also indicate that various models can be used to arrive at similar prices for European call options.

2.4.2 A numerical calibration example

The dataset used in the paper consists of 144 European call option prices available on the Eurostoxx 50 index on 7 October 2003. The Eurostoxx 50 is an index comprised of the stocks of 50 large companies from 12 countries in the Eurozone. On 7 October 2003 the index closed at €2476.61. We take the continuously compounded risk free interest rate to be 3% per annum in accordance with the assumptions made by the authors. Below we use the exponential $N \circ IG$ model to calculate the prices of the options available in this market. We use an estimation procedure as well as a calibration procedure.

In order to estimate P, the probability measure governing the distribution of S_t, we proceed as follows. We calculate the realised log-returns of S_t for a period of one year, and we fit a $N \circ IG$ distribution to these log-returns. Since Schoutens *et al.* (2004:66-78) uses a mean correcting argument in order to perform the calibration procedure we proceed in a similar fashion here. We use a maximum likelihood parameter estimation procedure in which α, β and δ are allowed to vary freely, but μ is chosen so as to ensure that the discounted price process forms a martingale. Parameter estimation for the $N \circ IG$ distribution is discussed in more detail in Section 3.3. The parameter estimates obtained using this method are

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (74.8181, 9.8668, -0.004, 0.0297)\,. \tag{2.16}$$

Figure 2.2 shows a density estimate of the calculated log-returns (as a solid line) with the density function of the $N \circ IG$ distribution with the estimated parameters superimposed (as a dashed line). The density estimate below (and all other density estimates in the remainder of the thesis) are calculated using Matlab's ksdensity.m. For a detailed discussion of density estimation, see Van der Vaart (1998:341-349).

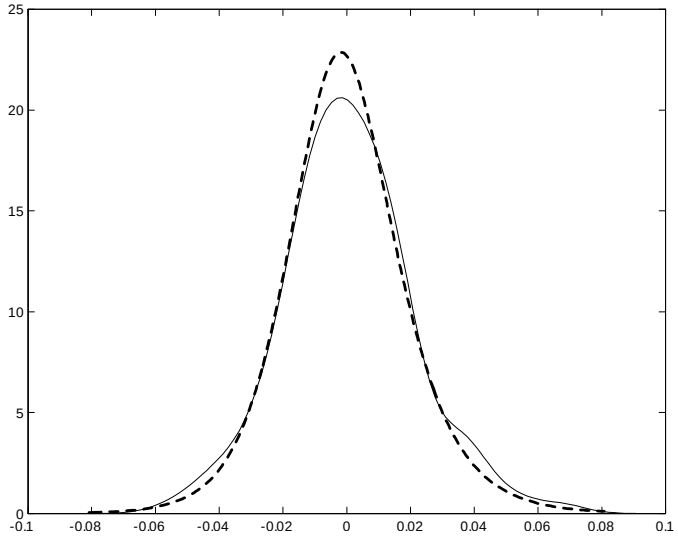

Figure 2.2: Density estimate of the log-returns (solid line) with the estimated $N \circ IG$ density (dashed line) superimposed.

We calculate the prices of the options using the parameters reported in (2.16) and we compare these prices to the observed option prices. Figure 2.3 provides a comparison of the observed option prices (circles) together with the corresponding option prices calculated under the exponential $N \circ IG$ model with the estimated parameters (crosses).

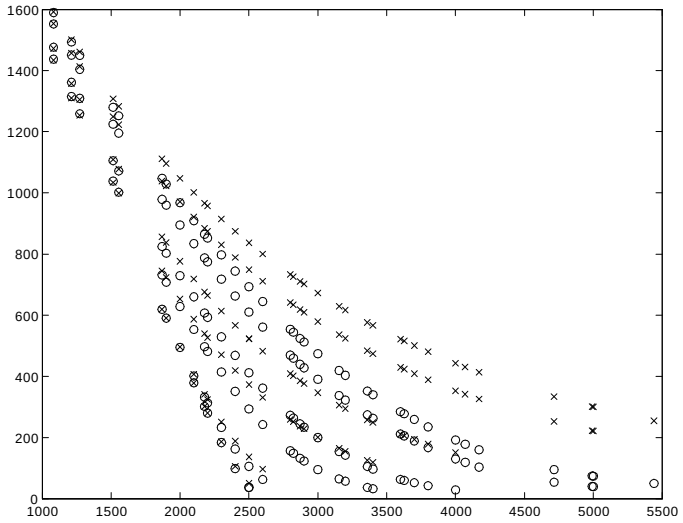

Figure 2.3: Observed option prices (circles) and option prices calculated using the exponential $N \circ IG$ model with estimated parameters (crosses).

Clearly the prices of the options calculated using the exponential $N \circ IG$ model with the estimated parameters do not fit the observed option prices very well. We calculate the following distance measures between the two sets of option prices:

$$AAE = 106.2328, \quad ARE = 0.676, \quad RMSE = 132.6805. \tag{2.17}$$

Next we consider a calibration algorithm. Similarly to the approach used in the paper, we use martingale restricted calibration with a mean correcting argument. This means that α, β and δ are allowed to vary freely while μ is chosen so as to ensure that the discounted stock price forms a martingale under the pricing measure. The calibration algorithm used minimises the AAE and results in the following parameter estimates:

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (50.1131, -48.0697, 0.0012, 0.0004)\,.$$

The calibration of the exponential $N \circ IG$ model is discussed in more detail in Chapter 5.

The observed option prices are indicated as circles in Figure 2.4, while the prices calculated using the calibrated exponential $N \circ IG$ model are shown as crosses.

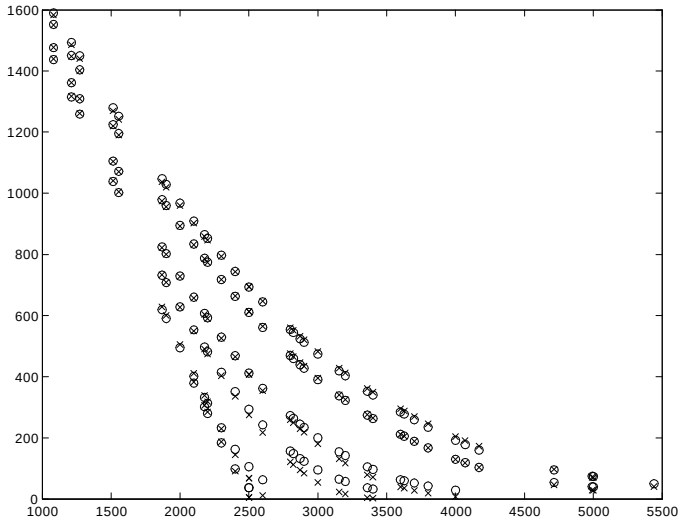

Figure 2.4: Observed option prices (circles) and option prices calculated using the exponential $N \circ IG$ model with calibrated parameters (crosses).

Figure 2.4 shows that the option prices calculated using the calibrated model closely resemble the observed option prices. The same is not true for the model obtained by estimating the parameters from the historical price process, see Figure 2.2.

We calculate the following distance measures between the observed option prices and the option prices calculated under the calibrated model:

$$AAE = 10.2444, \quad ARE = 0.0996, \quad RMSE = 15.0713. \tag{2.18}$$

Comparing the distance measures in (2.17) and (2.18) we see that the calibration procedure results in substantially lower distance measures than does the estimation procedure in this example.

2.4.3 Exotic options

After the authors calibrate various models to the observed European call option prices they turn their attention to different types of exotic options. The payoff of a European option depends on the terminal value of the stock price, but not on the history of the stock price before the maturity of the option. The payoff of an exotic option is a function of the terminal price of the stock as well as the path of the stock price. Schoutens *et al.* (2004:66-78) considers various types of exotic options including barrier options. In Chapters 7 and 8 we consider calibration and estimation techniques relating to three types of barrier options; digital barrier calls, down-and-out barrier calls and up-and-out barrier calls. Each of these option types are defined in Section 7.1.

The European call option prices calculated using the calibrated models in Schoutens *et al.* (2004:66-78) are almost identical. However, when these models are used to calculate the prices of exotic options (including the three types of barrier options mentioned above) the resulting option prices differ markedly. Monte Carlo simulation is used for the calculation of the exotic option prices.

The authors do not calibrate the models to exotic option prices. One of the possible reasons for this is that the prices of these options are not as readily available as the prices of European options.

2.4.4 Conclusions

The calibrated models are able to replicate the prices of the European call option nearly perfectly. Since the prices of these options are functions of the distribution of S_t at various times the authors conclude that the option pricing models used are able to reflect the distributional properties of S_t implied by the market. However, the exotic option prices calculated using the calibrated models vary markedly. This means that the path properties of S_t differ substantially from one calibrated model to the next.

The authors explain that the path properties of the stock price under the various models need to be considered in order to more accurately model exotic options. It remains unclear

whether or not it is possible to calibrate option pricing models to exotic option prices with the same level of precision as can evidently be achieved when using European call options. We discuss this question in detail in Chapters 7 and 8 for barrier options.

Chapter 3

Lévy processes and infinitely divisible distributions

Lévy processes are stochastic processes with stationary and independent increments. These processes are named after the French mathematician, Paul Lévy, who pioneered the field. Lévy processes are closely related to infinitely divisible distributions; for each Lévy process there exists a corresponding infinitely divisible distribution and vice versa. Lévy processes and infinitely divisible distributions are often used in financial modelling.

Throughout the thesis we use Lévy processes and infinitely divisible distributions to model the returns associated with a stock price. The definition of these processes and their connection to infinitely divisible distributions are discussed in Section 3.1. Here we also include the Lévy-Khintchine representation of Lévy processes. This representation is used in Section 4.4 in order to generalise a martingale measure often used in option pricing.

The $N \circ IG$ distribution is a popular choice of model for financial returns. Section 3.2 provides the definition of this distribution and discusses its properties. Parameter estimation for the $N \circ IG$ distribution is discussed in some detail in Section 3.3. Another popular choice of model is Brownian motion. However, the properties of Brownian motion can be found in any standard text on stochastic processes and are not discussed here.

In Section 3.4 we consider a subclass of Lévy processes called generalised gamma convolutions. The lognormal and Pareto processes are discussed as specific examples of processes belonging to this class. These processes are used in the context of time changed models in Chapters 6 and 7.

3.1 The definitions of Lévy processes and infinitely divisible distributions

Below we denote by X a random variable defined on a probability space $(\Omega, \mathcal{F}, P)$ and by L_t a Lévy process defined on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ where the filtration $\mathcal{F}_t$ is increasing in t and generated by the process L_t ($\mathcal{F}_t = \sigma(L_t)$).

Sato (2005:3) provides the following definition for a Lévy process.

Definition 8 *A process $L = (L_t; t \geq 0)$ defined on $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ is a one-dimensional Lévy process if the following conditions hold:*

1. For any $n \geq 1$ and $0 \leq t_0 \leq t_1 \leq \cdots \leq t_n$, the random variables L_{t_0}, $L_{t_1} - L_{t_0}$,...,$L_{t_n} - L_{t_{n-1}}$ are independent.

2. $L_0 = 0$ almost surely.

3. $L_{s+t} - L_s$ is independent of s for all $s \geq 0$ and $t > 0$.

4. L is stochastically continuous.

5. There is $\Omega_0 \in \mathcal{F}$ with $P(\Omega_0) = 1$ such that, for every $\omega \in \Omega_0$, $L_t(\omega)$ is right-continuous in $t \geq 0$ and has left limits.

In Chapter 6 we use a subclass of Lévy processes called subordinators in connection with time changed option pricing models.

Definition 9 *A Lévy process is called a subordinator if the paths of this process are almost surely non-decreasing.*

There exists a close connection between Lévy processes and infinitely divisible distributions. Sato (2005:31) provides the following definition of an infinitely divisible distribution. Let μ^n denote the n-fold convolution of a probability measure μ by itself.

Definition 10 *A probability measure μ on $\mathbb{R}$ is infinitely divisible if, for any positive integer n, there exists a probability measure μ_n on $\mathbb{R}$ such that $\mu = \mu_n^n$.*

The connection between infinitely divisible distributions and Lévy processes is expressed by the following theorem.

Theorem 11 *1. If L_t is a Lévy process on $\mathbb{R}$, then the distribution of L_t is infinitely divisible for all $t \geq 0$.*

2. Conversely, if μ is an infinitely divisible distribution on $\mathbb{R}$, then there is a Lévy process L_t, such that μ is the law of L_1.

3. If L_t and L_t' are Lévy processes on $\mathbb{R}$ such that the distribution of L_1 and L_1' are equal, then $L_t = L_t'$.

See Sato (2005:35-37) for a proof of this theorem.

3.1.1 The Lévy-Khintchine formula and the triplet of Lévy characteristics

A useful characterisation of Lévy processes is that the characteristic exponent of all Lévy processes satisfies the Lévy-Khintchine formula;

$$\log\left(E^P\left[\exp\left(iuL_t\right)\right]\right) = t\left\{iu\gamma - \frac{1}{2}\sigma^2u^2 + \int_{-\infty}^{\infty}\left(e^{iux} - 1 - iuh\left(x\right)\right)\nu\left(dx\right)\right\},$$

where $\gamma \in \mathbb{R}$, $\sigma^2 \geq 0$ and $h\left(x\right)$ is some truncation function. Throughout the thesis we use $h\left(x\right) = \mathbb{I}\left(|x| < 1\right)$. This is a popular choice of truncation function which has become standard in Lévy processes literature; see, for example, Sato (2005:37-38) and Shiryaev (2003:194-196). ν is a measure on $\mathbb{R}\backslash\left\{0\right\}$ such that

$$\int_{-\infty}^{\infty}\inf\left\{1, x^2\right\}\nu\left(dx\right) < \infty.$$

A Lévy process is uniquely determined by $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ for a given truncation function. The triplet $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ is called the triplet of Lévy characteristics. Any Lévy process can be decomposed into three independent parts; a straight line, a Brownian motion and a pure jump process. For a Lévy process with triplet of Lévy characteristics $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ the slope of the straight line component is γ, the variance of the Brownian motion component is σ^2, and the jumps made by the process are governed by $\nu\left(dx\right)$. Consider a set of possible

jump sizes A, where A is bounded away from 0. The number of jumps with sizes within the set A per unit time follows a Poisson process with intensity parameter $\int_A \nu\,(dx)$, see Schoutens (2003:45).

3.2 The normal inverse Gaussian ($N \circ IG$) distribution

The $N \circ IG$ distribution was introduced by Barndorff-Nielsen, see Barndorff-Nielsen (1997:1-13) and Barndorff-Nielsen (1998:41-68). This versatile distribution is often used in financial modelling. The discussion below is partially based on Schoutens (2003:59-60).

Below we consider the definition of the $N \circ IG$ distribution together with some of its properties and we show that the normal distribution can be obtained as a limit case of this distribution. We include an alternative parameterisation of the $N \circ IG$ distribution in this section. Below we also demonstrate some of the numerical difficulties that are encountered when working with this distribution.

3.2.1 Definition and properties

The density and characteristic functions of a $N \circ IG$ random variable are provided in Section 2.2. However, for ease of reference we include these functions below.

Definition 12 *A random variable $X : \Omega \to \mathbb{R}$ is said to follow a $N \circ IG$ distribution with parameter set $\theta = (\alpha, \beta, \mu, \delta)$ if it has density*

$$f(x;\theta) = \frac{\alpha\delta}{\pi} \exp\left(\delta\sqrt{\alpha^2-\beta^2} + \beta(x-\mu)\right) \frac{K_1\left(\alpha\sqrt{\delta^2+(x-\mu)^2}\right)}{\sqrt{\delta^2+(x-\mu)^2}}, \tag{3.1}$$

where $\alpha > 0$, $|\beta| < \alpha$, $\mu \in \mathbb{R}$, $\delta > 0$ and K_1 denotes the modified Bessel function of the third kind with index 1.

For the definition of this Bessel function, see Schoutens (2003:148). We denote a random variable following this distribution by $X \sim N \circ IG(\theta)$. The characteristic function of the

$N \circ IG\,(\theta)$ distribution is

$$\phi(t;\theta) = \exp\left(i\mu t - \delta\left(\sqrt{\alpha^2 - (\beta + it)^2} - \sqrt{\alpha^2 - \beta^2}\right)\right). \tag{3.2}$$

The first four standardised central moments of the distributions are used extensively in Section 3.3. We denote by $\nu_1\,(\theta)$, $\nu_2\,(\theta)$, $\nu_3\,(\theta)$ and $\nu_4\,(\theta)$ the mean, variance, skewness and kurtosis of the $N \circ IG\,(\theta)$ distribution. In the sequel we refer to $\nu_1\,(\theta)$ as the mean implied by the parameter set θ, the same convention is used for the implied variance, skewness and kurtosis. We also refer to the first four standardised central moments simply as the moments of the distribution in the remainder of the thesis. If $X \sim N \circ IG(\theta)$, then the expected value, variance, skewness and kurtosis of X are

$$\begin{aligned}
\nu_1\,(\theta) &= \mu + \frac{\delta\beta}{\sqrt{\alpha^2 - \beta^2}}, \\
\nu_2\,(\theta) &= \frac{\alpha^2\delta}{\sqrt{(\alpha^2 - \beta^2)^3}}, \\
\nu_3\,(\theta) &= \frac{3\beta}{\alpha\sqrt{\delta(\alpha^2 - \beta^2)^{1/2}}}, \\
\nu_4\,(\theta) &= 3\left(1 + \frac{\alpha^2 + 4\beta^2}{\delta\alpha^2\sqrt{\alpha^2 - \beta^2}}\right).
\end{aligned} \tag{3.3}$$

Consider convolutions of $N \circ IG$ random variables. If $Y_j,\ j = 1, 2, ..., n$, denotes n independent and identically distributed $N \circ IG\,(\alpha, \beta, \mu, \delta)$ random variables, then

$$\sum_{j=1}^{n} Y_j \sim N \circ IG(\alpha, \beta, n\mu, n\delta).$$

If $X \sim N \circ IG(\alpha, \beta, \mu, \delta)$ and $\gamma > 0$, then

$$\gamma X \sim N \circ IG\left(\frac{\alpha}{\gamma}, \frac{\beta}{\gamma}, \gamma\mu, \gamma\delta\right). \tag{3.4}$$

This result is used to generalise a martingale measure often used in option pricing in Section 4.4.

The $N \circ IG(\alpha, \beta, \mu, \delta)$ distribution has the normal distribution as a limit case.

Theorem 13 *If $\alpha \to \infty$, $\delta \to \infty$ and $\frac{\delta}{\alpha} \to \sigma^2$ while β is held constant, then the $N \circ IG$ distribution converges to the $N\left(\mu + \beta\sigma^2, \sigma^2\right)$ distribution.*

Proof. Consider the characteristic function of the $N \circ IG$ distribution with parameter set $\theta = (\alpha, \beta, \mu, \delta)$ where $\alpha \to \infty$:

$$\begin{aligned}
\phi(t;\theta) &= \exp\left(it\mu - \delta\left(\sqrt{\alpha^2 - (\beta + it)^2} - \sqrt{\alpha^2 - \beta^2}\right)\right) \\
&= \exp\left(it\mu - \alpha\delta\left(\sqrt{1 - \frac{(\beta + it)^2}{\alpha^2}} - \sqrt{1 - \frac{\beta^2}{\alpha^2}}\right)\right) \\
&= \exp\left(it\mu - \alpha\delta\left(-\frac{(\beta + it)^2}{2\alpha^2} + \frac{\beta^2}{2\alpha^2} + O\left(\alpha^{-4}\right)\right)\right) \\
&= \exp\left(it\left(\mu + \beta\frac{\delta}{\alpha}\right) - \frac{t^2}{2}\frac{\delta}{\alpha} + O\left(\alpha^{-3}\right)\right).
\end{aligned}$$

If $\delta = \alpha\sigma^2 + O\left(\alpha^{\lambda}\right)$ where $\lambda < 1$, then

$$\begin{aligned}
\phi(t;\theta) &= \exp\left(it\left(\mu + \beta\left(\sigma^2 + O\left(\alpha^{\lambda-1}\right)\right)\right) - \frac{t^2}{2}\left(\sigma^2 + O\left(\alpha^{\lambda-1}\right)\right) + O\left(\alpha^{-3}\right)\right) \\
&\to \exp\left(it\left(\mu + \beta\sigma^2\right) - \frac{t^2}{2}\sigma^2\right),
\end{aligned}$$

as $\alpha \to \infty$. This is the characteristic function of the $N\left(\mu + \beta\sigma^2, \sigma^2\right)$ distribution. ■

The fact that the $N \circ IG$ distribution has the normal distribution as a limit case is a source of difficulty in parameter estimation. This is discussed in Section 3.3.

A stochastic process L_t is a $N \circ IG$ process with parameter set $(\alpha, \beta, \mu, \delta)$ if L_t is a Lévy process such that the random variable L_1 follows a $N \circ IG(\alpha, \beta, \mu, \delta)$ distribution.

The increments of a $N \circ IG(\alpha, \beta, \mu, \delta)$ process have the following distribution:

$$L_{t+s} - L_t \sim N \circ IG(\alpha, \beta, \mu s, \delta s),$$

for all $s > 0$. The triplet of Lévy characteristics of a $N \circ IG$ process with parameter set $\theta = (\alpha, \beta, \mu, \delta)$ is

$$\left(\mu + \frac{2\delta\alpha}{\pi} \int_0^1 \sinh(\beta x) K_1(\alpha x)\, dx, 0, \frac{\delta\alpha}{\pi} \frac{\exp(\beta x) K_1(\alpha |x|)}{|x|} dx \right).$$

Venter *et al.* (2005:79-101) introduces an alternative parameterisation for the $N \circ IG$ distribution known as the standard normal inverse Gaussian ($SN \circ IG$) distribution.

Definition 14 *A random variable $X : \Omega \to \mathbb{R}$ is said to follow a $SN \circ IG$ distribution with parameter set $\theta = (\lambda, \psi, \mu, \sigma)$ if it has density function*

$$\begin{aligned} f(x;\theta) &= \frac{\psi}{\pi\sigma} \left(\frac{\psi^2 + \lambda^2}{\psi^2 + \left(\frac{x-\mu}{\sigma}\right)^2} \right)^{1/2} \exp\left(\psi^2 + \lambda \left(\frac{x-\mu}{\sigma} \right) \right) \times \\ &\quad K_1 \left(\sqrt{\psi^2 + \lambda^2} \sqrt{\psi^2 + \left(\frac{x-\mu}{\sigma} \right)^2} \right), \end{aligned}$$

with $\lambda \in \mathbb{R}$, $\psi > 0$, $\mu \in \mathbb{R}$, $\sigma > 0$ and K_1 as before.

When using the $SN \circ IG$ parameterisation we encounter numerical problems similar to those associated with the $N \circ IG$ parameterisation. In the remainder of the thesis we use the $N \circ IG$ parameterisation defined in (3.1).

Cont (2001:223-236) provides a detailed analysis of the statistical properties of asset returns. The distributional properties of the log-returns of stock prices listed in the paper include heavier tails than that of the normal distribution as well as asymmetry. The $N \circ IG$ distribution provides the modeller with both of these properties.

Cont (2001:223-236) advocates for the use of distributions with at least four parameters when modelling log-returns. Four parameters are necessary to realistically model the

location, scale, asymmetry and tail behaviour of the log-returns.

3.2.2 Numerical difficulties

When evaluating the density function of the $N \circ IG$ distribution we encounter numerical difficulties for certain parameter sets because of the limited computation power provided by the software used for numerical calculation. This problem can be remedied using characteristic function inversion.

The Bessel function $K_1(z) \downarrow 0$ as $z \to \infty$. Since most computer packages used for numerical work round all sufficiently small positive numbers to 0, the Bessel function is rounded to 0 for sufficiently large arguments. If the $\alpha \to \infty$ and $\delta \to \infty$, then the exponential function and the Bessel function in (3.1) both have large arguments. If these arguments are large enough the computer package used sets the values of the Bessel and the exponential functions to 0 and ∞ respectively and the calculation of the density function breaks down.

For parameter sets leading to these numerical problems, the density function can be obtained numerically using the well-known Fourier inversion formula. If we denote by $f(x,\theta)$ and $\phi(t,\theta)$ respectively the density and characteristic functions of the $N \circ IG$ distribution with parameter set θ, then the Fourier inversion formula can be expressed as

$$f(x;\theta) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-itx} \phi(t;\theta)\, dt, \tag{3.5}$$

see Feller (1971:509-510). Since ϕ is smooth and bounded the right hand side in (3.5) can be approximated very accurately by the Riemann sum

$$\frac{1}{2\pi} \sum_{t \in I} e^{-itx} \phi(t)\, \Delta t, \tag{3.6}$$

where $I = \left\{ \frac{i}{m}; |i| \leq ma \right\}$ for suitably large values of m and a. The values of m and a are

chosen such that

$$\phi\left(ma\right) \approx 0,$$

and the intervals $\frac{1}{m}$ are small enough to accurately approximate ϕ. We provide a numerical example below in which we report the values chosen for m and a.

A numerical example

Consider the use of (3.6) to calculate the density function of a $N \circ IG\left(10, -5, 0, 0.1\right)$ random variable. Figure 3.1 shows the characteristic function of this random variable on the interval $[-5; 5]$. The solid and dashed lines represent the real and imaginary parts of the characteristic function respectively.

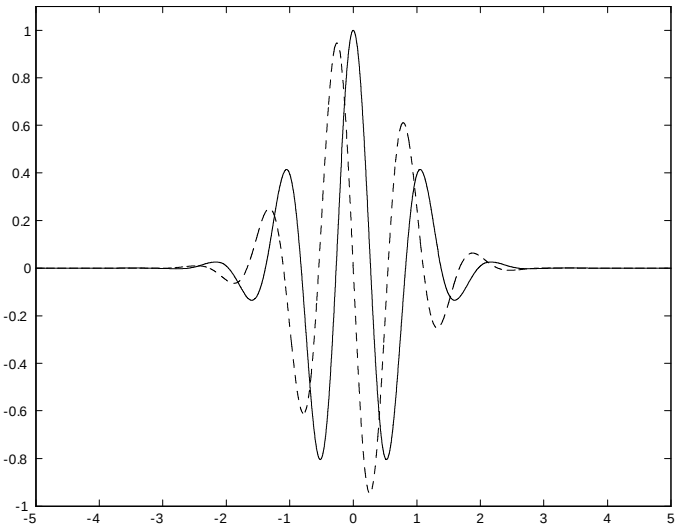

Figure 3.1: Real (solid line) and imaginary (dashed line) parts of the characteristic function of a $N \circ IG(10, -5, 0, 0.1)$ random variable.

The density function calculated using (3.6) is compared to the density function calculated directly using (3.1) in Figure 3.2 below. The dashed line represents the density function calculated using formula (3.6), while the solid line represents the density function calculated using (3.1).

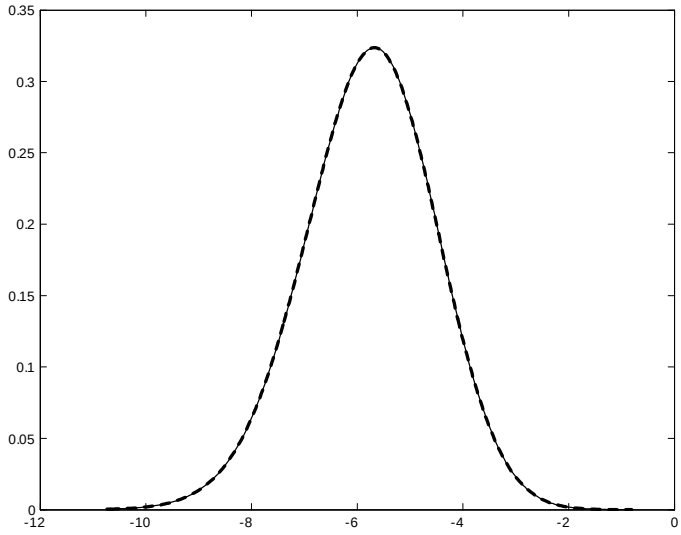

Figure 3.2: The $N \circ IG(10, -5, 0, 1)$ density calculated directly (solid line) and using characteristic function inversion (dashed line).

Figure 3.2 is obtained using the approximation in (3.6) where I consists of a set of 50001 equally spaced points on the interval $[-7; 7]$. This corresponds to $I = \left\{\frac{i}{m}; |i| \leq ma\right\}$ with

$$m = \frac{50000}{14} \approx 3571 \quad \text{and} \quad a = \frac{7}{m} = 0.00196.$$

Figure 3.2 indicates that the density function calculated using characteristic function inversion is practically indistinguishable from the density function obtained by direct calculation.

When using (3.6), the density function is obtained by integrating a complex integrand. To demonstrate that this method works satisfactorily it remains to show that the imaginary part of the density function obtained is close to 0 (the only reason for the density function to contain a complex part is numerical rounding errors). In the example used the complex part of the density function (on the same interval as above) is calculated to be between -4×10^{-17} and 5×10^{-17}. We conclude that the small complex part of the density function is a product of numerical rounding errors.

To illustrate why the direct calculation of the density function in (3.1) fails computationally, consider the density function of the $N \circ IG\,(10, 50, -50, 100)$ distribution at

$x = 0$. In this case the arguments of the exponential and Bessel functions are 1116.0254 and 1118.034 respectively. As noted previously, Matlab is limited in terms of the largest and smallest numbers that it can calculate. As a result, the values of the exponential and Bessel functions are set to ∞ and 0 respectively and (3.1) cannot be evaluated directly. However, the Fourier method described above is successful in evaluating the density at $x = 0$.

3.3 Parameter estimation for the $N \circ IG$ distribution

In this section we discuss parameter estimation for the $N \circ IG$ distribution. We consider two estimators; the maximum likelihood estimator (MLE) and the empirical characteristic function estimator ($ECFE$). Below we illustrate the behaviour of these estimators through some numerical examples.

When using the MLE we encounter numerical difficulties. We propose a solution to the problem using Fourier inversion. Two variations of the MLE are considered. The first suffers from the numerical problems mentioned, while the second circumvents these difficulties using Fourier inversion. We also consider two variations of the $ECFE$, the difference between the two estimators being the weight function used.

3.3.1 The maximum likelihood estimator

Consider the MLE and its application to the $N \circ IG$ distribution. Since the density function of this distribution is known explicitly it is, in principle, possible to find MLEs. However, there are no analytical expressions for these estimators and the estimates have to be found numerically using an optimisation algorithm.

As explained in Section 3.2.2, the Bessel function goes to 0 if its argument becomes large. If, because of limited computational power, the Bessel function is rounded to 0, then the log-likelihood function assumes the value $-\infty$. When calculating MLEs for a given dataset we use two different procedures. The first procedure allows the log-likelihood function to assume the value $-\infty$ in cases where the argument of the Bessel function is

sufficiently large. Since the optimisation algorithm seeks a parameter set that maximises the log-likelihood function, this procedure simply does not converge to parameter values that lead to excessively large arguments of the Bessel function in (3.1). In effect, this procedure maximises the log-likelihood function subject to the restriction that the Bessel function can be calculated without being rounded to 0. Below we denote this estimator $DMLE$ (for direct MLE). The second procedure for the calculation of MLEs uses Fourier inversion to calculate the log-likelihood function in order to avoid the situation where the value of this function is set to $-\infty$. Below we denote this estimator $IMLE$ (for indirect MLE).

Note that the two variations of MLE differ only in terms of practical implementation; both of these methods calculate MLEs. The $DMLE$ is restricted in terms of the estimates that can be obtained, but it converges much faster than the $IMLE$ because no (numerically intensive) Fourier inversion is necessary.

3.3.2 The empirical characteristic function estimator

Another estimator for the parameters is based on the empirical characteristic function

$$\phi_n(u) = \frac{1}{n}\sum_{j=1}^{n} \exp\left(iux_j\right),$$

where $x_1, x_2, ..., x_n$ is some dataset. The ECFE of θ is defined as the minimiser over θ of the expression

$$D_n\left(\theta\right) = \int_{-\infty}^{\infty} \left|\phi_n(u) - \phi(u;\theta)\right|^2 w(u)du, \tag{3.7}$$

with ϕ defined in (3.2) and w some positive weight function satisfying

$$\int_{-\infty}^{\infty} w\left(u\right) du < \infty.$$

We use two variations of the weight function,

$$w_1(u) = \exp\left(-|u|\right), \tag{3.8}$$

and

$$w_2(u) = \exp\left(-u^2\right). \tag{3.9}$$

Below we denote by $ECFEA$ and $ECFEN$ respectively the estimators obtained using (3.8) and (3.9).

Heathcote (1977:255-264) explains that the use of estimators based on the empirical characteristic function may be preferable to MLEs in cases where the likelihood function is much more complicated than the characteristic function. The mentioned paper derives a central limit theorem for the $ECFE$ and presents results relating to the consistency of this estimator. For a more general exposition of the $ECFE$ and related Fourier methods, see Feuerverger and McDunnough (1981:379-387). Yu (2004:93-123) provides a more practical discussion of this estimator including an example using financial data.

3.3.3 Starting values

Neither the MLEs nor the $ECFE$s are known analytically, meaning that optimisation algorithms are used to calculate these estimates numerically. These optimisation algorithms require starting values. The resulting parameter estimates calculated using the various estimators are sensitive to the starting values used. We use the following procedure to obtain the starting values for the MLE. A grid of possible starting values are specified for each parameter. The possible starting values for the parameter estimates are as follows; for α we use $(1, 5, 20, 50, 100, 200, 300)$, for β we use $(-95, -20, -10, 0, 10, 20, 95)$, for μ we use $(-20, -10, -5, 0, 5, 10, 20)$ and for δ we use $(0.001, 0.1, 1, 10, 20, 100, 250)$. The log-likelihood function is calculated for each combination that can be constructed from the grids of possible starting values. The parameter set associated with the largest log-likelihood value is chosen as starting values for the optimisation algorithm.

For the $ECFE$ the same approach is used to find starting values with the exception that the log-likelihood function is not calculated at each point in the grid, rather we calculate the distance defined in (3.7). The parameter set resulting in the smallest distance is used as the starting values for the optimisation algorithm.

In the case of the $IMLE$ the time required to calculate the log-likelihood function for all combinations of the possible starting values is prohibitive in some cases. In order to alleviate this computation burden we do not calculate the log-likelihood function for each possible combination of starting values. Instead we calculate the mean and the standard deviation of the dataset under consideration as well as the mean implied by each of the parameter sets. All of the parameter sets implying a mean that is not within three standard deviations of the mean of the dataset are discarded. For each of the remaining datasets the log-likelihood function is calculated and the parameter set associated with the largest log-likelihood is used as starting values. Similar techniques to reduce the time required to calculate the remaining estimates were deemed unnecessary.

3.3.4 Empirical results

We now consider some empirical results obtained as follows. We simulate 250 realisations (this roughly corresponds to the number of business days in a year) from a $N \circ IG$ distribution with known parameter set and we fit a $N \circ IG$ distribution to the simulated data; we calculate the $DMLE$s, $IMLE$s, $ECFEA$s and $ECFEN$s. This process is repeated 1000 times for each of the two examples considered below.

Recall that the normal distribution is a limit case of the $N \circ IG$ distribution. In the second example we simulate from a $N \circ IG$ distribution resembling the normal distribution in that it exhibits little skewness and excess kurtosis. In the first example this is not the case. For both examples we include boxplots of the estimates obtained for α. Boxplots of the remaining parameters are omitted for the sake of brevity. We omit a few outliers in several of the boxplots so as not to obscure the distribution of the estimates.

For each example we report the mean times required in order to calculate the parameter

estimates using the different estimators.

Example 1: The $N \circ IG\,(10, 5, 5, 0.1)$ distribution

Figure 3.3 shows boxplots of the estimates of α calculated using the $N \circ IG\,(10, 5, 5, 0.1)$ distribution.

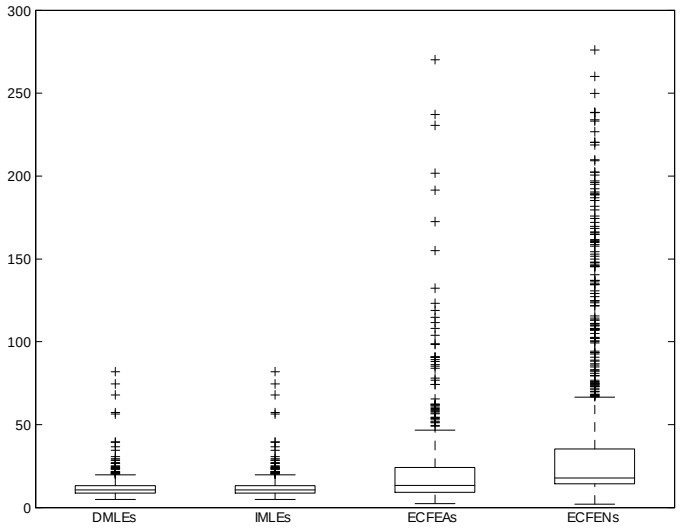

Figure 3.3: Boxplots of the estimates of α in the $N \circ IG\,(10, 5, 5, 0.1)$ example.

The results of the $DMLE$ and the $IMLE$ coincide for this example. In the second example it will become apparent that this is not generally the case. Note that all of the estimators overestimate the true value of α substantially. The distributions of the estimates are also positively skewed. The maximum estimate for α obtained using the $ECFEN$ is 1353374. This value (and a few other large values) are omitted in the boxplots above. Similar configurations of values of the estimates of β, μ and δ are observed.

It is clear that the results obtained using the different estimators differ. However, the density functions associated with the different parameter estimates are closer than one might expect when looking at the values of the estimates. Denote by $\hat{\theta}_{DMLE}$ and $\hat{\theta}_{ECFEA}$ the estimates calculated using the $DMLE$ and the $ECFEA$ respectively (similarly, denote the estimates calculated using the $IMLE$ and the $ECFEN$ by $\hat{\theta}_{IMLE}$ and $\hat{\theta}_{ECFEN}$

respectively). For one of the datasets we obtained the following estimates:

$$\hat{\theta}_{DMLE} = (20.9882, 11.967, 4.9467, 0.1488)\,,$$

$$\hat{\theta}_{ECFEA} = (15.4296, 6.7976, 4.9804, 0.1488)\,.$$

Figure 3.4 shows the density functions that are calculated when $\hat{\theta}_{DMLE}$ (solid line) and $\hat{\theta}_{ECFEA}$ (dashed line) are substituted into (3.1).

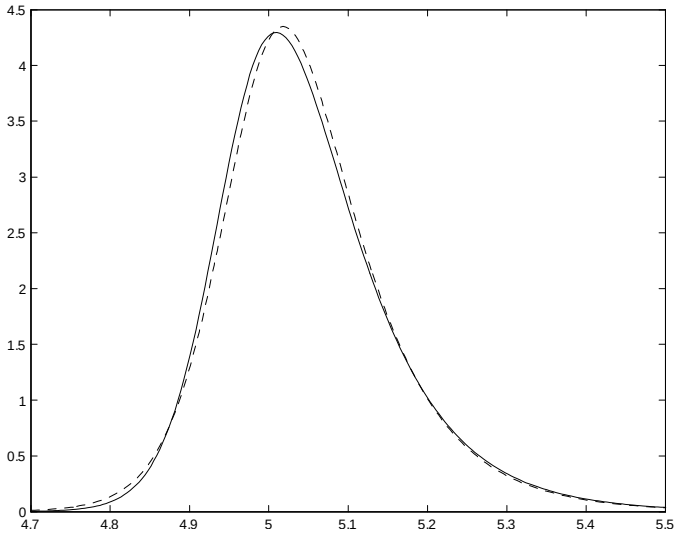

Figure 3.4: Comparison of estimated density functions.

The density functions obtained are similar in spite of substantial differences in the estimates of α and β.

For each realised $\hat{\theta}_{DMLE}$, $\hat{\theta}_{IMLE}$, $\hat{\theta}_{ECFEA}$ and $\hat{\theta}_{ECFEN}$ we calculate the implied moments. Table 3.1 shows the average and standard deviation (in brackets) of the implied moments associated with each estimator. The table also shows the average times required in seconds to calculate the estimates. The $N \circ IG\,(10, 5, 5, 0.1)$ distribution has mean, variance, skewness and kurtosis equal to 5.058, 0.015, 1.612 and 9.928.

	$DMLE$	$IMLE$	$ECFEA$	$ECFEN$
$\nu_1\left(\hat{\theta}\right)$	5.058 (0.008)	5.058 (0.008)	5.058 (0.008)	5.057 (0.008)
$\nu_2\left(\hat{\theta}\right)$	0.016 (0.003)	0.016 (0.003)	0.015 (0.003)	0.015 (0.003)
$\nu_3\left(\hat{\theta}\right)$	1.584 (0.46)	1.584 (0.46)	1.528 (0.651)	1.088 (0.927)
$\nu_4\left(\hat{\theta}\right)$	10.097 (3.37)	10.097 (3.37)	9.776 (6.017)	7.923 (14.632)
Time	0.468	18.316	11.834	6.867

Table 3.1: The moments implied by the various estimators for the $N \circ IG\,(10, 5, 5, 0.1)$ example.

The moments implied by the estimates calculated for each estimator are close to the moments of the true distribution, possibly with the exception of the skewness and kurtosis implied by the $ECFEN$s. Note that the calculation of the $DMLE$s required the least amount of time, while the $IMLE$s required the most.

Example 2: The $N \circ IG\,(230, -25, 0, 90)$ distribution

As a second example we consider the $N \circ IG\,(230, -25, 0, 90)$ distribution. As was pointed out above, this distribution resembles the normal distribution in that it exhibits a skewness of nearly 0 and a kurtosis close to 3. Figure 3.5 shows a boxplot of the $DMLE$s of α for this distribution.

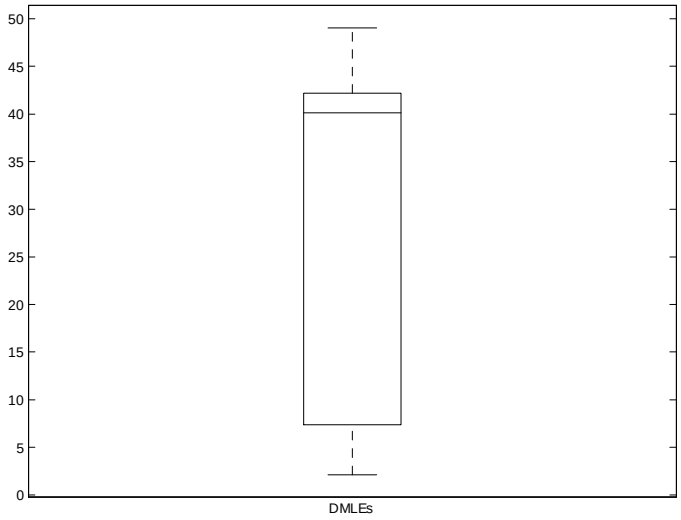

Figure 3.5: Boxplot of the DMLEs of α in the $N \circ IG\,(230, -25, 0, 90)$ example.

Note that the $DMLE$s underestimate the true value of α. The boxplots of the remaining estimators for α are shown in Figure 3.6.

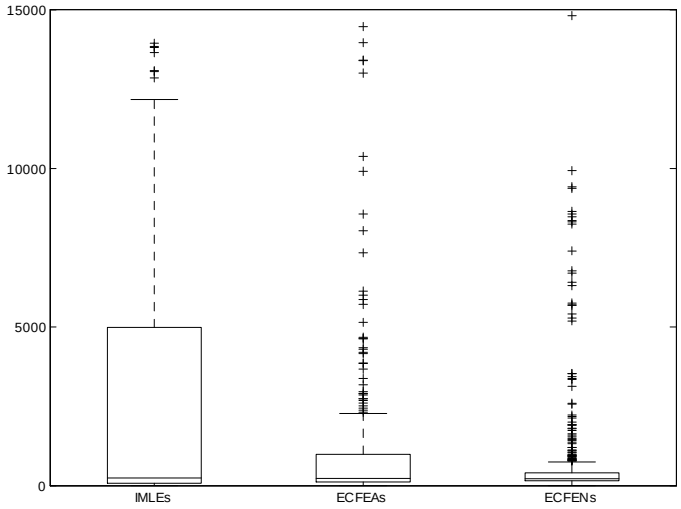

Figure 3.6: Boxplots of the estimates of α in the $N \circ IG\,(230, -25, 0, 90)$ example.

In contrast to the $DMLE$, the estimators shown in Figure 3.6 clearly overestimate the true value of α. As was the case in the previous example, the distributions of the $IMLE$, the $ECFEA$ and the $ECFEC$ shown in Figure 3.6 are positively skewed.

The mean and standard deviation of the implied moments associated with each of the estimators, as well as the mean times required in seconds to calculate the estimates, are shown

in Table 3.2. The mean, variance, skewness and kurtosis of the $N \circ IG\,(230, -25, 0, 90)$ distribution are -9.841, 0.398, -0.002 and 3.0002 respectively.

	$DMLE$	$IMLE$	$ECFEA$	$ECFEN$
$\nu_1\left(\hat{\theta}\right)$	$-9.839\,(0.041)$	$-9.84\,(0.041)$	$-9.838\,(0.043)$	$-9.839\,(0.042)$
$\nu_2\left(\hat{\theta}\right)$	$0.397\,(0.035)$	$0.397\,(0.035)$	$0.399\,(0.039)$	$0.397\,(0.036)$
$\nu_3\left(\hat{\theta}\right)$	$-0.005\,(0.1)$	$0.015\,(0.059)$	$0.02\,(0.062)$	$0.011\,(0.046)$
$\nu_4\left(\hat{\theta}\right)$	$3.141\,(0.242)$	$3.039\,(0.14)$	$3.041\,(0.17)$	$3.01\,(0.077)$
Time	0.474	477.207	13.858	7.549

Table 3.2: The moments implied by the various estimators for the $N \circ IG\,(230, -25, 0, 90)$ distribution.

The moments implied by the parameter estimates closely correspond to the moments of the true distribution, perhaps with the exception of the kurtosis implied by the $DMLE$s. Numerical difficulties in the calculation of the log-likelihood function occur for parameter sets that imply a kurtosis close to that of the normal distribution. As a result, the $DMLE$ does not converge to these values. As was the case in the previous example, calculating the $DMLE$s required the least amount of time, while calculating the $IMLE$s required the most time. In this example the calculation of the $IMLE$s took roughly 1000 times as long as the calculation of the $DMLE$s.

Conclusions

We consider four different estimators for the parameters of the $N \circ IG$ distribution. The estimates obtained using these methods differ wildly, but the moments implied by the various sets of estimates correspond closely to each other and to that of the distribution from which the data are realised.

The differences between the calculated estimates are especially large if the distribution underlying the data exhibits little skewness and excess kurtosis. In these situations the

implied kurtosis associated with the $DMLE$ tends to overestimate that of the underlying distribution, which can be a problem if the tail behaviour of the estimated distribution is of particular interest. This problem can be remedied by using characteristic function inversion to calculate the $IMLE$. However, the time required to calculate the $IMLE$ can be prohibitive if a large number of parameter estimates are required.

3.4 Generalised Γ-convolutions

The generalised Γ-convolutions are a subclass of infinitely divisible distributions introduced in Thorin (1977a:31-40). Thorin (1977a:31-40) and Thorin (1977b:121-148) respectively show that the Pareto and the lognormal distributions belong to this class.

Below we consider the definition of generalised Γ-convolutions and we include the motivation for the chosen name. We also derive the triplet of Lévy characteristics of a generalised Γ-convolution Lévy process. Special attention is paid to the Pareto and lognormal distributions; the Lévy processes corresponding to these distributions are used as subordinators in Chapters 6 and 7 in the context of time changed models.

3.4.1 Definition of a generalised Γ-convolution

Thorin (1977a:31-40) defines a generalised Γ-convolution as a distribution for which the moment generating function can be written as

$$M(s) = \exp\left(as + \int_0^\infty \log\left(\frac{1}{1-\frac{s}{y}}\right) dU(y)\right), \tag{3.10}$$

where $\mathrm{Re}(s) \leq 0$, $a \geq 0$ and U is non-decreasing and satisfies

$$\begin{aligned} U(0) &= 0, \\ \int_0^1 |\log(y)| \, dU(y) &< \infty, \\ \int_1^\infty \frac{dU(y)}{y} &< \infty. \end{aligned} \tag{3.11}$$

A generalised Γ-convolution is characterised by a and U in (3.10). In Sections 3.4.4 and 3.4.5 we report the values of a and U that characterise the Pareto and lognormal distributions respectively.

In order to see why generalised Γ-convolution is a sensible name for this class of distributions choose U as an increasing step function with m steps occurring at $y_1 < y_2 < ... < y_m$ (where $y_1 > 0$). Then U can be written as

$$U(y) = \sum_{j=1}^{m} u_j \mathbb{I}(y \geq y_j), \tag{3.12}$$

where u_j denotes the size of the j^{th} step. Substituting (3.12) into (3.10) we obtain

$$\begin{aligned} M(s) &= \exp\left(as + \int_0^\infty \log\left(\frac{1}{1 - \frac{s}{y}} \right) dU(y) \right) \\ &= \exp\left(as + \sum_{j=1}^{m} u_j \log\left(\left(1 - \frac{s}{y} \right)^{-1} \right) \right) \\ &= e^{as} \times \exp\left(\sum_{j=1}^{m} \log\left(\left(1 - \frac{s}{y} \right)^{-u_j} \right) \right) \\ &= e^{as} \prod_{j=1}^{m} \left(\left(1 - \frac{s}{y} \right)^{-u_j} \right), \end{aligned}$$

which corresponds to the moment generating function of an m-fold gamma convolution translated by a, see Schoutens (2003:52).

3.4.2 Triplet of Lévy characteristics

Theorem 15 *The triplet of Lévy characteristics, with respect to the truncation function* $\mathbb{I}(|x| \leq 1)$*, of a generalised* Γ*-convolution Lévy process with moment generating function given by (3.10) is*

$$\left(a + \int_0^\infty \frac{1 - e^{-y}}{y} dU(y), 0, \mathbb{I}(x \geq 0) \int_0^\infty e^{-yx} dU(y) \frac{dx}{x} \right).$$

Proof. Assume, without loss of generality, that $a = 0$.

Since a generalised Γ-convolution is infinitely divisible with support $[0, \infty)$, its moment generating function can be written as

$$M(s) = \exp\left(\int_0^\infty \frac{e^{sx} - 1}{x} dQ(x) \right), \tag{3.13}$$

where

$$\int_1^\infty \frac{dQ(x)}{x} < \infty,$$

see Feller (1971:450).

In order to prove the theorem we shall show that Q in (3.13) can be expressed as

$$Q(x) = \mathbb{I}(x \geq 0) \int_0^\infty \frac{1 - e^{-yx}}{y} dU(y). \tag{3.14}$$

If Q is defined by (3.14) then its derivative is

$$Q'(x) = \mathbb{I}(x \geq 0) \int_0^\infty e^{-yx} dU(y). \tag{3.15}$$

Substituting (3.15) into (3.13) the moment generating function of a generalised Γ-convolution reduces to

$$\begin{aligned} M(s) &= \exp\left(\int_0^\infty \frac{e^{sx}-1}{x}\int_0^\infty e^{-yx}dU(y)\,dx\right) \\ &= \exp\left(\int_0^\infty\int_0^\infty \frac{e^{-(y-s)x}-e^{-yx}}{x}dxdU(y)\right) \\ &= \exp\left(\int_0^\infty \log\left(\frac{y-s}{y}\right)dU(y)\right) \\ &= \exp\left(\int_0^\infty \log\left(\frac{1}{1-\frac{s}{y}}\right)dU(y)\right), \end{aligned} \tag{3.16}$$

where the third equality above is obtained using Frullani's integral:

$$\int_0^\infty \frac{f(ax)-f(bx)}{x}dx = \log\left(\frac{b}{a}\right)(f(0)-f(\infty)),$$

if $f'(x)$ is continuous and

$$-\infty < \int_1^\infty \frac{f(x)-f(\infty)}{x}dx < \infty,$$

see Spiegel and Liu (1999:100). Since (3.16) equals (3.10) with $a = 0$ we have shown that $Q(x)$ and $Q'(x)$ are given by (3.14) and (3.15) respectively.

Using (3.13) the logarithm of the moment generating function can be expressed as follows:

$$\begin{aligned}
\log\left(M\left(s\right)\right) &= \int_0^\infty \frac{e^{sx}-1}{x} dQ\left(x\right) \\
&= \int_0^\infty \left(e^{sx}-1-sx\mathbb{I}\left(x\leq 1\right)+sx\mathbb{I}\left(x\leq 1\right)\right)\frac{Q\left(dx\right)}{x} \\
&= s\left(Q\left(1\right)-Q\left(0\right)\right) \\
&\quad + \int_0^\infty \left(e^{sx}-1-sx\mathbb{I}\left(x\leq 1\right)\right)\frac{Q\left(dx\right)}{x}. \qquad (3.17)
\end{aligned}$$

In the last step above we split the integral under the assumption that $Q\left(1\right)-Q\left(0\right)$ is finite. We show that this is indeed the case below, see (3.18). From (3.14), $Q\left(0\right)=0$. From (3.17) we identify the Lévy triplet of a generalised Γ-convolution Lévy process with $a=0$. The drift is given by $Q\left(1\right)$, no Brownian motion component is present, and the Lévy measure is

$$\frac{Q\left(dx\right)}{x} = \mathbb{I}\left(x\geq 0\right)\int_0^\infty e^{-yx} dU\left(y\right)\frac{dx}{x}.$$

Introducing $a>0$ corresponds to a translation. This affects neither the Brownian motion component of the Lévy triplet nor the Lévy measure. However, a translation of a serves to increase the value of the drift by a units. Therefore, the drift of a generalised Γ-convolution Lévy process is

$$a+Q\left(1\right)=a+\int_0^\infty \frac{1-e^{-y}}{y} dU\left(y\right).$$

■

3.4.3 Path properties

Below we consider the path properties of generalised Γ-convolution Lévy processes. Specifically we provide proofs that these processes are infinite activity processes and have bounded variation. Below we denote the Lévy measure of such a process by ν.

Theorem 16 *The paths of generalised* Γ*-convolution Lévy processes have bounded variation.*

Proof. To show that the paths of these processes have bounded variation it suffices to show that

$$\int_{|x|\leq 1} |x| \, \nu(dx) < \infty, \tag{3.18}$$

see Sato (2005:140-141). Choose $\delta \in (0,1)$ such that

$$\frac{1-e^{-x}}{x} \leq |\log(x)|,$$

for all $x \in (0,\delta]$. In order to show that (3.18) holds, consider

$$\begin{aligned}
\int_{|x|\leq 1} |x| \, \nu(dx) &= \int_0^1 x\nu(dx) \\
&= Q(1) \\
&= \int_0^\infty \frac{1-e^{-x}}{x} dU(x) \\
&= \int_0^\delta \frac{1-e^{-x}}{x} dU(x) + \int_\delta^1 \frac{1-e^{-x}}{x} dU(x) + \int_1^\infty \frac{1-e^{-x}}{x} dU(x). \qquad (3.19)
\end{aligned}$$

In order to complete the proof we need to show that the three integrals in (3.19) are finite.

Consider the first integral in (3.19). From the definition of δ

$$\int_0^\delta \frac{1-e^{-x}}{x} dU(x) \leq \int_0^\delta |\log(x)| \, dU(x) < \infty, \tag{3.20}$$

where we have used the properties of U reported in (3.11).

Consider the second integral in (3.19);

$$\begin{aligned}
\int_{\delta}^{1} \frac{1-e^{-x}}{x} dU(x) &\leq e^{\frac{1}{2}} \int_{\delta}^{1} \frac{\exp\left(-\frac{x}{2}\right) - \exp\left(-\frac{3x}{2}\right)}{x} dU(x) \\
&\leq e^{\frac{1}{2}} \int_{0}^{\infty} \frac{\exp\left(-\frac{x}{2}\right) - \exp\left(-\frac{3x}{2}\right)}{x} dU(x) \\
&\leq e^{\frac{1}{2}} \left\{ Q\left(\frac{3}{2}\right) - Q\left(\frac{1}{2}\right) \right\} \\
&= e^{\frac{1}{2}} \int_{\frac{1}{2}}^{\frac{3}{2}} x\nu(dx) \\
&< \infty, \qquad (3.21)
\end{aligned}$$

from the properties of a Lévy measure, see Schoutens (2003:44-45).

It remains to show that the third integral in (3.19) is finite;

$$\int_{1}^{\infty} \frac{1-e^{-x}}{x} dU(x) \leq \int_{1}^{\infty} \frac{dU(x)}{x} < \infty, \qquad (3.22)$$

from the properties of U shown in (3.11).

The proof is completed by substituting (3.20), (3.21) and (3.22) into (3.19). ■

Theorem 17 *Generalised* Γ*-convolution Lévy processes are infinite activity processes.*

Proof. In order to show that generalised Γ-convolution Lévy processes are infinite activity processes it suffices to show that

$$\nu(0, \infty) = \infty, \qquad (3.23)$$

see Sato(2005:136). To show that (3.23) holds for generalised Γ-convolutions, consider

$$\begin{aligned}
\nu(0,\infty) &= \int_0^\infty \frac{Q(dx)}{x} \\
&= \int_0^\infty \frac{1}{x} \int_0^\infty e^{-yx} dU(y)\, dx \\
&= \int_0^\infty \int_0^\infty \frac{1}{x} e^{-yx} dx dU(y) \\
&= \infty.
\end{aligned}$$

This completes the proof. ■

3.4.4 The Pareto distribution

Definition 18 *A random variable X is said to follow a Pareto distribution with parameters $g > 0$ and $h > 0$ if it has density*

$$f(x) = \frac{hg^h}{(g+x)^{h+1}},$$

for $x \geq 0$.

We denote a random variable X following a Pareto distribution with parameters g and h by $X \sim Pareto(g,h)$. Thorin (1977a:31-40) shows that the $Pareto(g,h)$ distribution is a generalised Γ-convolution with $a = 0$ and

$$\begin{aligned}
U(x) &= \frac{1}{\pi} \operatorname{arccot} \left\{ \frac{\Gamma(h)}{\pi} e^{gx} \times \right. \\
&\quad \left. \left((gx)^{-h} + \sum_{j=1}^{\infty} \frac{(gx)^{j-h}}{(h-1)...(h-j)} \right) - \cot(\pi h) \right\}.
\end{aligned}$$

3.4.5 The lognormal distribution

Definition 19 *A random variable X is said to be lognormally distributed with parameters $\mu \in \mathbb{R}$ and $\sigma > 0$ if the density of X is given by*

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\log(x) - \mu)^2}{2\sigma^2}\right), \tag{3.24}$$

for $x > 0$.

A lognormal random variable with parameters μ and σ is denoted by $X \sim \log N\left(\mu, \sigma^2\right)$. Thorin (1977b:121-148) uses an alternative parameterisation of the lognormal distribution. However, the parameterisation shown in (3.24) is used throughout the thesis.

Thorin (1977b:121-148) shows that the lognormal distribution belongs to the class of generalised Γ-convolutions. For this distribution $a = 0$ and

$$U(x) = \frac{1}{\pi} \arctan\left(\frac{\mathrm{Im}(\lambda(x))}{\mathrm{Re}(\lambda(x))}\right), \tag{3.25}$$

where

$$\lambda(x) = \exp\left(\frac{\pi^2}{2\sigma^2}\right) \int_{-\infty}^{\infty} \exp\left(-xe^{\mu-\sigma t} + \frac{i\pi t}{\sigma}\right) \phi(t)\, dt, \tag{3.26}$$

with ϕ the density function of a standard normal random variable. Substituting (3.26) into (3.25) and simplifying, U can be written as

$$U(x) = \frac{1}{\pi} \arctan\left(\frac{\int_{-\infty}^{\infty} \exp\left(-xe^{\mu-\sigma t} - \frac{1}{2}t^2\right) \sin\left(\frac{\pi t}{\sigma}\right) dt}{\int_{-\infty}^{\infty} \exp\left(-xe^{\mu-\sigma t} - \frac{1}{2}t^2\right) \cos\left(\frac{\pi t}{\sigma}\right) dt}\right). \tag{3.27}$$

However, arctan is a multi-valued function and U must be non-decreasing. Thorin (1977b:121-148) states that the function in (3.27) should be defined such that $U(0) = 0$ and behave continuously. We consider this in more detail below.

Denote by $x_1 < x_2 < ...$ the values of x for which $\mathrm{Re}(\lambda(x)) = 0$. For $j = 1, 2, ...$ U is strictly increasing and continuous on each interval (x_j, x_{j+1}). However, each value of x_j is

a discontinuity point of U;

$$\lim_{x \uparrow x_j} U(x) = \frac{1}{\pi} \arctan(+\infty) = \frac{1}{2},$$

and

$$\lim_{x \downarrow x_j} U(x) = \frac{1}{\pi} \arctan(-\infty) = -\frac{1}{2}.$$

To remove the discontinuities define

$$U^*(x) = U(x) + \sum_{j \geq 0} j \mathbb{I}(x_j \leq x \leq x_{j+1}),$$

where U is defined in (3.27) and $x_0 = 0$.

To show that replacing U with U^* does not change the Lévy measure of the lognormal distribution, define

$$\begin{aligned}
Q^*(x) &= \mathbb{I}(x \geq 0) \int_0^\infty \frac{1 - e^{-yx}}{y} dU^*(y) \\
&= \mathbb{I}(x \geq 0) \sum_{j \geq 0} \int_{x_j}^{x_{j+1}} \frac{1 - e^{-yx}}{y} dU^*(y) \\
&= \mathbb{I}(x \geq 0) \sum_{j \geq 0} \int_{x_j}^{x_{j+1}} \frac{1 - e^{-yx}}{y} dU(y) \\
&= \mathbb{I}(x \geq 0) \int_0^\infty \frac{1 - e^{-yx}}{y} dU(y) \\
&= Q(x).
\end{aligned}$$

In conclusion a lognormal distribution is a generalised Γ-convolution with $a = 0$ and

$$\begin{aligned} U(x) &= \frac{1}{\pi} \arctan\left(\frac{\int_{-\infty}^{\infty} \exp\left(-xe^{\mu-\sigma t} - \frac{1}{2}t^2\right) \sin\left(\frac{\pi t}{\sigma}\right) dt}{\int_{-\infty}^{\infty} \exp\left(-xe^{\mu-\sigma t} - \frac{1}{2}t^2\right) \cos\left(\frac{\pi t}{\sigma}\right) dt} \right) \\ &\quad + \sum_{j \geq 0} j \mathbb{I}\left(x_j \leq x \leq x_{j+1}\right), \end{aligned}$$

where $x_1 < x_2 < ...$ are the values of x for which $\mathrm{Re}\left(\lambda(x)\right) = 0$, with $\lambda(x)$ defined in (3.26).

Chapter 4

Martingale measures and exponential Lévy models

In this chapter we discuss exponential Lévy models as well as various martingale measures used together with these models. The Black-Scholes and exponential $N \circ IG$ models are considered as specific examples within the class of exponential Lévy models.

In Section 4.1 we define an exponential Lévy model. In Sections 4.2 and 4.3 we consider two methods of constructing martingale measures often used in financial modelling; the Esscher transform and the mean correcting martingale measure (MCMM). The discussions relating to the Esscher transform and the MCMM are based, in part, on Schoutens (2003:77-80).

In Section 4.4 we propose a generalisation of the MCMM which we refer to as the generalised mean correcting martingale measure (GMCMM). We consider the characteristics of the GMCMM as well as the use of this measure in option pricing. We also include an economic interpretation of the GMCMM. In Section 4.5 we discuss the calculation of multiple option prices using the MCMM and the GMCMM.

4.1 Definition of an exponential Lévy model

The market considered throughout the thesis is defined in Section 2.1.1. We use this market below, with the additional specification that under the probability measure P

$$S_t = S_0 \exp(L_t), \tag{4.1}$$

where S_0 is the initial stock price and L_t is some Lévy process. We consider two exponential Lévy models; the Black-Scholes model (under which L_t follows a Brownian motion) and the exponential $N \circ IG$ model (under which L_t follows a $N \circ IG$ process).

In the remainder of this chapter we consider various changes of measure that can be used to obtain pricing measures for the calculation of arbitrage free option prices.

4.2 The Esscher transform

Below we consider the Esscher transform and its application to the Black-Scholes and the exponential $N \circ IG$ model. Let f_t be the density of L_t under the objective probability measure P. For a real number λ such that

$$\int_{-\infty}^{\infty} e^{\lambda x} f_t(x)\, dx < \infty,$$

define a new density as follows:

$$f_t^{\lambda}(x) = \frac{e^{\lambda x} f_t(x)}{\int_{-\infty}^{\infty} e^{\lambda x} f_t(x) dx}. \tag{4.2}$$

Denote by P^{λ} the probability measure under which the density of L_t is f_t^{λ}.

A family of probability measures are obtained by varying the value of λ in (4.2). The application of the Esscher transform entails choosing the value of λ such that $e^{-rt}S_t$ forms a P^{λ}-martingale. Let ϕ_t denote the characteristic function of L_t under P. The discounted price process $e^{-rt}S_t$ forms a P^{λ}-martingale if and only if

$$e^r \phi_1(-i\lambda) = \phi_1(-i(\lambda + 1)), \tag{4.3}$$

see Schoutens (2003:78). Typically, (4.3) has a unique solution $\lambda = \lambda^*$.

The probability measure obtained using the Esscher transform is a martingale measure by definition. In order to see that the obtained measure is equivalent to P, note that the

Radon-Nikodym derivative of P^λ to P is

$$\frac{dP^\lambda}{dP} = \frac{e^{\lambda x}}{\int_{-\infty}^{\infty} e^{\lambda x} f_t(x)\, dx}, \tag{4.4}$$

see (4.2). Since (4.4) is strictly positive for all $x \in \mathbb{R}$, P^λ is equivalent to P. As a result, the measure obtained by the application of the Esscher transform is a $LEMM$ and can be used for the calculation of arbitrage free option prices.

In both the Black-Scholes model and the exponential $N \circ IG$ model L_t follows the same type of Lévy process under P^λ as it does under P. However, this change of measure changes the parameters of L_t. Formulae for the changes in parameters brought about by the application of the Esscher transform are provided below.

For a detailed discussion on the Esscher transform the interested reader is referred to Gerber and Shiu (1994:99-191).

4.2.1 The Black-Scholes model

Consider the case where L_t is a Brownian motion with parameters $\mu \in \mathbb{R}$ and $\sigma^2 > 0$; i.e. L_t in equation (4.1) is given by

$$L_t = \mu t + \sigma W_t,$$

where W_t denotes a standard Brownian motion. The application of the Esscher transform results in a change of the location parameter μ, while the value of σ^2 remains unchanged. If we denote by μ^* the value of the location parameter after the application of the Esscher transform, then

$$\mu^* = r - \frac{\sigma^2}{2}.$$

When discussing the Black-Scholes model in Section 2.1.4 we explain that there exists a unique LEMM under this model. The application of the Esscher transform results in this LEMM, compare (2.6).

4.2.2 The exponential $N \circ IG$ model

Let L_t be a $N \circ IG$ process with parameter set $(\alpha, \beta, \mu, \delta)$. The application of the Esscher transform leads to a change in the parameters of L_t so that L_t follows a $N \circ IG$ process with parameters $(\alpha, \beta + \gamma^*, \mu, \delta)$ where $\gamma = \gamma^*$ is the solution to

$$r = \mu + \delta \left(\sqrt{\alpha^2 - (\beta + \gamma)^2} - \sqrt{\alpha^2 - (\beta + \gamma + 1)^2} \right),$$

see Schoutens (2003:79).

4.3 The mean correcting martingale measure

We now turn our attention to the mean correcting martingale measure (MCMM) and its role in the calculation of arbitrage free option prices. The MCMM changes the location parameter of L_t so as to ensure that $e^{-rt} S_t$ forms a martingale. The values of the remaining parameters of L_t are not affected by the application of the MCMM. Below we denote by Q the probability measure obtained when applying the MCMM to P. Define Y as

$$Y_t = L_t + rt - \log \left(E^P \left[\exp \left(L_t \right) \right] \right).$$

Let Q be the probability measure such that

$$\begin{aligned} P\left(L_t \leq a \right) &= Q\left(Y_t \leq a \right) \\ &= Q\left(L_t \leq a - rt + \log \left(E^P \left[\exp \left(L_t \right) \right] \right) \right), \end{aligned}$$

for all $a \in \mathbb{R}$. Changing measure from P to Q increases the drift of L by $r - \log \left(E^P \left[\exp \left(L_1 \right) \right] \right)$. Let μ^* denote the location parameter of L under Q in this section. Changing measure from P to Q entails setting

$$\mu^* = r - \log \left(E^P \left[\exp \left(L_1 \right) \right] \right). \tag{4.5}$$

Q is a martingale measure. To see this consider

$$\begin{aligned} e^{-rt}S_t &= S_0 \exp\left(-rt + L_t\right) \\ &= S_0 \exp\left(L_t - \log\left(E^P\left[\exp\left(L_t\right)\right]\right)\right) \\ &= \frac{S_0 \exp\left(L_t\right)}{E^P\left[\exp\left(L_t\right)\right]}, \end{aligned} \tag{4.6}$$

which shows that the discounted price process $e^{-rt}S_t$ forms a Q-martingale.

Next we consider the equivalence of Q to P. Let the Lévy triplet of L be given by $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ under P. Yao *et al.* (2011:593-597) proves the following theorem.

Theorem 20 *Q is equivalent to P if and only if $\sigma^2 > 0$.*

Since $\sigma^2 = 0$ for many Lévy processes (including the $N \circ IG$ process) the application of the MCMM does not, in general, result in an equivalent martingale measure. However, Yao *et al.* (2011:593-597) shows that, under certain conditions, the measure obtained using the MCMM can be used for arbitrage free option pricing when $\sigma^2 = 0$. This is discussed below.

4.3.1 The use of the MCMM in option pricing

Consider a European option with payoff function g and time to maturity T. Denote by $\pi\left(Q\right)$ the option price calculated using the pricing measure Q;

$$\pi\left(Q\right) = e^{-rT}E\left[g\left(S_T\right)\right].$$

Theorem 21 *Let L_t be a Lévy process with Lévy measure $\nu\left(dx\right)$ containing no Brownian motion component under probability measure P. If $\nu\left(dx\right)$ satisfies the following two properties,*

$$\nu\left(\left(-\infty, a\right]\right) > 0, \tag{4.7}$$

for all $a \in \mathbb{R}$ *and*

$$\nu \text{ has no atom and } \int_{[-1,0)} |x| \, \nu(dx) = \int_{(0,1]} x\nu(dx) = \infty, \tag{4.8}$$

then there exists an equivalent martingale measure Q^* *such that*

$$\pi(Q) = \pi(Q^*).$$

For a proof, see Yao *et al.* (2011:593-597). This theorem can be interpreted as follows. In cases where L_t contains no Brownian motion component (and hence Q is not a LEMM) Q can still be used to calculate an arbitrage free price of a European option if the conditions specified in (4.7) and (4.8) are satisfied. The reason for this is that there exists some LEMM that produces the same option price as is obtained when using Q.

Below we discuss the use of the MCMM in the Black-Scholes and exponential $N \circ IG$ models.

4.3.2 The Black-Scholes model

Consider the case where L_t is a Brownian motion with parameters $\mu \in \mathbb{R}$ and $\sigma^2 > 0$ under P. The application of the MCMM changes the value of the location parameter so that

$$\mu^* = r - \frac{\sigma^2}{2}. \tag{4.9}$$

Note that the changes in the parameter sets affected by the Esscher transform and the MCMM coincide if L_t is a Brownian motion.

Since L_t contains a diffusion part ($\sigma^2 > 0$) the probability measure obtained is equivalent to P, meaning that Q is a LEMM. As a result, Q can be used for the calculation of arbitrage free option prices.

4.3.3 The exponential $N \circ IG$ model

Let L_t be a $N \circ IG$ process with parameter set $(\alpha, \beta, \mu, \delta)$. Applying the MCMM to the $N \circ IG$ distribution entails setting the location parameter to

$$\mu^* = r + \delta \left(\sqrt{(\alpha^2 - (\beta+1)^2)} - \sqrt{\alpha^2 - \beta^2} \right), \tag{4.10}$$

see (3.2) and (4.5).

Since L contains no Brownian motion part under this model Q is not equivalent to P. However, conditions (4.7) and (4.8) above hold for the $N \circ IG$ process, see Appendix A for a proof. As a result, Q can be used to calculate an arbitrage free price of a European option.

4.4 The generalised mean correcting martingale measure

Below we propose a generalisation of the MCMM. We refer to this change of measure as the generalised mean correcting martingale measure (GMCMM). We consider the use of this measure in option pricing and provide an economic interpretation of the GMCMM. We also consider the application of this change of measure to the Black-Scholes model and the exponential $N \circ IG$ model.

Let L_t be a Lévy process with triplet of Lévy characteristics $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ under the objective probability measure P. Define Y as follows:

$$Y_t = \theta L_t + rt - \log\left(E^P\left[\exp\left(\theta L_t\right)\right]\right), \tag{4.11}$$

with $\theta > 0$. Let P^θ be the probability measure such that

$$\begin{aligned} P\left(L_t \leq a\right) &= P^\theta\left(Y_t \leq a\right) \\ &= P^\theta\left(L_t \leq \frac{a - rt + \log\left(E^P\left[\exp\left(L_t\right)\right]\right)}{\theta}\right), \end{aligned}$$

for all $a \in \mathbb{R}$. The triplet of Lévy characteristics of L_t under the measure P^θ is

$$\left(r - \frac{\theta^2\sigma^2}{2} - \int_{-\infty}^{\infty} \left(e^x - 1 - x\mathbb{I}\left(|x| < 1\right)\right) \nu\left(\frac{dx}{\theta}\right), \theta^2\sigma^2, \nu\left(\frac{dx}{\theta}\right) \right); \tag{4.12}$$

see Appendix B for a derivation.

The probability measure P^θ is a martingale measure. To see this consider the discounted stock price under P^θ;

$$\begin{aligned} e^{-rt}S_t &= S_0 \exp\left(-rt + Y_t\right) \\ &= S_0 \exp\left(\theta L_t - \log\left(E^P\left[\exp\left(\theta L_t\right)\right]\right)\right) \\ &= \frac{S_0 \exp\left(\theta L_t\right)}{E^P\left[\exp\left(\theta L_t\right)\right]}, \end{aligned}$$

which forms a martingale for all $\theta > 0$. Note that, in the case where $\theta = 1$, the GMCMM reduces to the MCMM. Indeed, the measure obtained by applying the GMCMM for general $\theta > 0$ corresponds to the measure obtained by applying the MCMM to the process θL_t.

We have shown above that $e^{-rt}S_t$ forms a P^θ-martingale, we now consider the local equivalence of P^θ to P, specifically we consider the case where $\theta \neq 1$. If $\theta = 1$ the GMCMM reduces to the MCMM; the equivalence of the MCMM to P is discussed in Section 4.3.1. The following theorem provides a test for local equivalence.

Theorem 22 *Consider two probability measures Q and Q' of Lévy processes. Let the Lévy process under consideration have triplet of Lévy characteristics $(c, b, F(dx))$ under Q and $\left(\tilde{c}, \tilde{b}, \tilde{F}(dx)\right)$ under Q'. The measures Q and Q' are locally equivalent if and only if*

(α) $\tilde{F}$ is equivalent to F and $\int_{-\infty}^{\infty}\left(1 - \left(d\tilde{F}/dF\right)^{\frac{1}{2}}\right) dF < \infty$,

(β) $\tilde{c} = c$,

(γ) either $c > 0$ or $c = 0$ and $\tilde{b} = b + \int_{-\infty}^{\infty} x\mathbb{I}\left(|x| < 1\right)\left(\tilde{F} - F\right)(dx)$.

For a proof, see Barndorff-Nielsen and Shiryaev (2010:143).

Under probability measure P the process L has triplet of Lévy characteristics given by $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$, while the triplet of Lévy characteristics of L under P^θ is given in (4.12). Consider first the case where L contains a Brownian motion component under P (the case where $\sigma^2 > 0$). In this case, for P^θ to be equivalent to P we require, by condition (β), that

$$\theta^2\sigma^2 = \sigma^2,$$

which holds only when $\theta = 1$. Therefore, if $\sigma^2 > 0$ and $\theta \neq 1$, then P^θ is not locally equivalent to P.

We now consider the case where L contains no Brownian motion part under P (the case where $\sigma^2 = 0$). For P^θ to be locally equivalent to P we require, by condition (γ), that

$$\gamma = r - \int_{-\infty}^{\infty} \left(e^x - 1 - x\mathbb{I}\left(|x| < 1\right)\right)\nu\left(\frac{dx}{\theta}\right) - \int_{-1}^{1} x\left[\nu\left(\frac{dx}{\theta}\right) - \nu\left(dx\right)\right]. \tag{4.13}$$

Since r is the risk free interest rate specified by the market, (4.13) does not hold in general and P^θ is not locally equivalent to P. Therefore, P^θ is generally not a LEMM, the exception being when $\sigma^2 > 0$ and $\theta = 1$.

4.4.1 The use of the GMCMM in option pricing

As was the case for the MCMM, the GMCMM can, in some cases, be used in arbitrage free option pricing in spite of the fact that this change of measure does not give rise to a LEMM. This is discussed below. The arguments used are similar to the arguments used in Yao *et al.* (2011:593-597) to show the same result for the MCMM.

Let L_t be a Lévy process containing no Brownian motion component under P. Also, let $\mathcal{M}$ be the class of measures locally equivalent to P under which $e^{-rt}S_t$ forms a martingale, and let $\mathcal{M}^{'}$ be the subclass of $\mathcal{M}$ such that L_t is a Lévy process under Q for all $Q \in \mathcal{M}^{'}$. Consider a European call option with strike price K and time to maturity T. Denote

by $\pi^c(Q)$ the arbitrage free price of the option calculated with respect to the probability measure Q;

$$\pi^c(Q) = e^{-rT} E^Q\left[(S_T - K)^+\right].$$

Let I_c and $I_c^{'}$ respectively denote the range of option prices attainable under probability measures in $\mathcal{M}$ and $\mathcal{M}^{'}$;

$$I_c = \left\{\pi^c(Q) : Q \in \mathcal{M}\right\},$$

and

$$I_c^{'} = \left\{\pi^c(Q) : Q \in \mathcal{M}^{'}\right\}.$$

Theorem 23 *If the Lévy measure $\nu(dx)$ of L_t under probability measure P satisfies the properties specified in (4.7) and (4.8), then $\mathcal{M}$ is not empty, I_c is the full interval*

$$\left(\left(S_0 - e^{-rT}K\right)^+, S_0\right), \tag{4.14}$$

and $I_c^{'}$ is dense in this interval.

For a proof of this theorem, see Eberlein and Jacod (1997:131-140). This theorem shows that any price in the interval given in (4.14) is an arbitrage free price for a European call option (if L_t has no Brownian motion component and satisfies conditions (4.7) and (4.8)).

Theorem 24 *If S_t follows an exponential Lévy process containing no Brownian motion component and satisfying conditions (4.7) and (4.8), and P^θ is the GMCMM for some $\theta > 0$, then $\pi^c\left(P^\theta\right)$ is contained in the interval given in (4.14).*

For a proof of this theorem, see Appendix C.

Since $\pi^c\left(P^\theta\right)$ is contained in the interval given in (4.14) there exists a LEMM Q^* such that the price of a European call option calculated with respect to Q^* equals the price calculated with respect to P^θ;

$$\pi^c(P^\theta) = \pi^c(Q^*).$$

As a result, P^{θ} can be used to calculate an arbitrage free price of a European call option.

We now turn our attention to the calculation of European put option prices. Let $\mathcal{M}$ and $\mathcal{M}^{'}$ be defined as above, and let I_p and $I_p^{'}$ respectively denote the range of put option prices attainable under the probability measures contained in $\mathcal{M}$ and $\mathcal{M}^{'}$;

$$I_p = \left\{ \pi^p \left(Q \right) : Q \in \mathcal{M} \right\},$$

and

$$I_p^{'} = \left\{ \pi^p \left(Q \right) : Q \in \mathcal{M}^{'} \right\}.$$

Theorem 25 *If the Lévy measure $\nu\left(dx\right)$ of L_t under probability measure P satisfies the properties specified in (4.7) and (4.8), then $\mathcal{M}$ is not empty, I_p is the full interval*

$$\left(\left(e^{-rT} K - S_0 \right)^+, e^{-rT} K \right), \tag{4.15}$$

and $I_p^{'}$ is dense in this interval.

The proof of this theorem follows immediately from Theorem 23 and the put-call parity, see Hull (2009:208-211).

Theorem 26 *If S_t follows an exponential Lévy process containing no Brownian motion component and satisfying conditions (4.7) and (4.8), and P^{θ} is the GMCMM for some $\theta > 0$, then $\pi^p \left(P^{\theta} \right)$ is contained in the interval given in (4.15).*

The proof of this theorem is similar to that of Theorem 24 and is omitted.

Theorems 23 and 25 show that, if L_t has no Brownian motion component and satisfies (4.7) and (4.8), then the GMCMM can be used to calculate an arbitrage free price for a European option.

The Black-Scholes model

If L_t is a Brownian motion with parameters μ and σ^2 under P, then the application of the GMCMM entails changing the parameter set of L_t to $\left(r - \frac{\theta^2\sigma^2}{2}, \theta^2\sigma^2\right)$.

As is explained above, P^θ can be used for the calculation of arbitrage free option prices under the Black-Scholes model only if $\theta = 1$. When the GMCMM is used to calculate a single option price θ can be chosen so as to exactly equate the calculated option price with the observed option price. In this case $\theta^2\sigma^2$ is the implied volatility of the option.

The exponential $N \circ IG$ model

If L_t is a $N \circ IG$ process with parameter set $(\alpha, \beta, \mu, \delta)$ and $\theta > 0$, then θL_1 follows a $N \circ IG(\alpha/\theta, \beta/\theta, \theta\mu, \theta\delta)$ distribution, see (3.4). Using (3.2)

$$E^P\left[\exp\left(\theta L_1\right)\right] = \exp\left(-\delta\left(\sqrt{\alpha^2 - (\beta+\theta)^2} - \sqrt{\alpha^2 - \beta^2}\right) + \theta\mu\right).$$

As a result, under P^θ, L_t is a $N \circ IG$ process with parameter set $(\alpha/\theta, \beta/\theta, \mu_\theta, \theta\delta)$, where

$$\mu_\theta = r + \delta\left(\sqrt{\alpha^2 - (\beta+\theta)^2} - \sqrt{\alpha^2 - \beta^2}\right).$$

P^θ is not a LEMM under this model. However, since (4.7) and (4.8) hold for the $N \circ IG$ distribution and the $N \circ IG$ process contains no Brownian motion component, P^θ can be used to calculate an arbitrage free price of a European option.

4.4.2 An economic interpretation of the GMCMM

A European option can be interpreted as an insurance policy against a specified movement in the market price of the underlying stock. The price of a European call option can be considered to be the price of an insurance policy against the stock price rising above the strike price of the option. Similarly, a European put option can be interpreted as an insurance policy against the stock price falling below the strike price. The price of

both these insurance policies are increasing functions of the variance of L_t. The larger the variance of L_t, the greater the uncertainty regarding the future levels of the stock price and the more valuable the insurance policies.

Consider the effect of the GMCMM on the moments of L_t. Denote by m_2 the variance of L_t under probability measure P. The variance of L_t under the measure obtained using the GMCMM is $\theta^2 m_2$, while the skewness and kurtosis of L_t are not affected by this change of measure. The change in the mean of L_t is necessary to ensure that the discounted value of the stock price forms a martingale.

Consider some financial market containing European option prices. An analyst wishing to calculate the prices of the options in this market (under a given model) might estimate the objective probability measure P using historical log-returns. The analyst might then employ some change of measure to obtain the pricing measure Q and calculate the option prices under the model by taking expectations with respect to Q. If the majority of the participants in the market calculate the prices of the options under Q (or some other pricing measure leading to similar option prices) and accepts this as the true values of the options, then the option prices calculated using Q will be close to the observed option prices. However, if the majority of the market participants believe, for example, that the variance of the log-returns of the stock price will be significantly higher in the future than in the past, then the option prices calculated using the model will substantially underestimate the observed option prices.

In contrast to the Esscher transform and the MCMM, the GMCMM contains a free parameter θ. If, for example, the market expects that future log-returns will be more volatile than past log-returns, a value for θ larger than 1 can be specified in order to increase the variance of L_t under the model. Indeed the model can match the variance expected by the market by choosing θ such that $Var(\theta L_1)$ equals the expected future variance of the log-returns of the stock price. Note that in this case θ^2 plays a role similar to that of the implied volatility in the Black-Scholes model.

4.5 A note on the use of the MCMM and the GMCMM in the pricing of multiple options

In the discussion above we explain that the application of the MCMM and the GMCMM do not, in general, lead to LEMMs. However, each of the resulting measures can be used to calculate an arbitrage free price of a European option if certain conditions are satisfied. This is demonstrated (in the case of the MCMM as well as in the case of the GMCMM) by showing that there exists a LEMM Q^* such that the price of the option under Q^* equals the price of the option under the measure obtained using the MCMM or GMCMM.

This argument holds for markets containing a single option. If we introduce a second option into the market and we calculate the price of this option using the MCMM or the GMCMM, then again there exists a LEMM that produces the same option price. However, the LEMMs that equate the prices of the two options to those calculated under the MCMM or the GMCMM are not necessarily equal. It is unclear whether or not the market remains arbitrage free if we introduce multiple options and we calculate the prices of these options using the MCMM or the GMCMM. In Chapter 5 we use the MCMM as well as the GMCMM to calculate option prices in markets containing multiple options.

Chapter 5

Numerical calibration results for European options

In this chapter we use the techniques discussed and developed in the previous chapters to calibrate exponential Lévy models to observed European option prices; we use the exponential $N \circ IG$ model and the Black-Scholes model.

In Section 5.1 we discuss the market data used to obtain the empirical results shown in Section 5.2.

5.1 Market data

Section 5.2 presents numerical results obtained using various types of calibration; we discuss the observed financial data used for the calibrations below. We consider European call and put options. For every option the price, time to maturity, strike price and open interest associated with the option are available. The open interest of an option is the number of open positions currently held by investors in that option; this is a measure of how actively the option is traded. The market data used also contains historical time series data on the index and stock prices used.

Below we describe various problems encountered when analysing the market data as well as the steps taken in an effort to ensure that the empirical results are not negatively affected by these problems.

5.1.1 Option price data

The options considered are obtained from the American financial market, which is the largest and most liquid financial market in the world. As a result, it is hard for a single

market participant to manipulate option prices.

The data consist of European call and put option prices on the S&P 500 index, the PowerShares index and Google shares. The market data discussed below is as at close of business on 11 May 2012. The six datasets considered were obtained from http://finance.yahoo.com. Since all of these datasets are obtained at the same time and in the same market, the same risk free interest rate will be used for each dataset. Kemp (2009:105) provides the following definition of a risk free interest rate.

Definition 27 *The 'risk-free rate' might be deemed to be the interest rate available by investing in debt that is considered to be truly risk-free for the currency in question (typically, this involves debt issued by the government that controls the issuance of the relevant currency).*

Keeping this definition in mind, we use the discount rate of Treasury bills in the secondary market as the risk free interest rate. For each of the datasets below the average of the times to maturity is roughly six months, so we use the discount rate on the six month treasury bill as the risk free interest rate. On 11 May 2012 this rate was 0.15% per annum.

Two of the three assets underlying the option prices discussed above are indexes. These indexes do not pay dividends. Google is one of the few traded companies that do not pay out any dividends. As a result, our models do not need to be adjusted to include a dividend yield.

On 11 May 2012 the S&P 500 index closed at \$1353.39. We obtained the prices of 576 call options and 779 put options on this index. The PowerShares index is an index weighted by market capatilisation that consists of the largest non-financial companies the world over. On this day the PowerShares index closed at \$64.18. The recorded options on this index include 413 call options and 480 put options. On the same date the share price of Google Inc. closed at \$605.23. We obtained the prices of 545 call options and 532 put options on this stock.

Some of the option prices in the captured data seem to be either higher or lower than can

reasonably be expected. We believe that this could be the result of human error (including typing mistakes) or some other shortcoming of the method used to capture the data. We believe that, if these unexpected option prices are allowed to remain in the datasets, the empirical results shown in Section 5.2 will be negatively impacted. Below we describe three steps used to remove the relevant options from the datasets. The first step is to remove option prices that lead to immediate arbitrage opportunities. The second is to identify and remove outlying option prices and the final step is to remove duplicated options.

Option prices leading to arbitrage opportunities

Hull (2009:205-208) provides upper and lower bounds for arbitrage free European call and put option prices. These bounds are not model specific and the arbitrage free intervals provided below are the largest intervals possible under any option pricing model. The bounds for a European call option price π^c are

$$\left(S_0 - Ke^{-rT}\right)^+ \leq \pi^c \leq S_0, \tag{5.1}$$

and the bounds for a European put option price π^p are

$$\left(Ke^{-rT} - S_0\right)^+ \leq \pi^p \leq Ke^{-rT}. \tag{5.2}$$

If a European call option price falls outside the interval given in (5.1) then this option constitutes an immediate arbitrage opportunity. Similarly a European put option price not contained in the interval (5.2) provides an immediate arbitrage opportunity. Since the options considered are actively traded, it is reasonable to assume that these arbitrage opportunities would be exploited immediately. If, for example, an option price falls below the lower bound specified above traders would buy vast numbers of these options in order to realise risk free profits. As a result, the demand for this option would exceed the supply and normal market forces would increase the price of the option until the arbitrage opportunity is removed. Similarly, if an option price was to exceed the upper bound investors would

sell this option in large quantities and the supply would exceed the demand, this would lower the price and remove the arbitrage opportunity.

In order to ensure that the data used represent realistic option prices we remove all call options from the data that do not satisfy (5.1). Similarly we remove all put options that fail to satisfy (5.2).

Outlying option prices

After removing the option prices not contained in realistic intervals we turn our attention to possible outliers remaining in the data. For each dataset the following method is used to identify outliers and to remove these option prices from the datasets used. For each dataset the options are split into several smaller datasets according to their time to maturity; i.e. we group all options with equal times to maturity together. We fit a second degree polynomial (using linear regression) to each of these groups of option prices; the strike price of the option is used as the predictor variable and the option price is used as the response. A second degree polynomial is used to compensate for the volatility smile, see Fouque *et al.* (2001:33-37). Figure 5.1 shows the prices of the call options on the Google shares with a time to maturity of 160 days with the fitted polynomial superimposed.

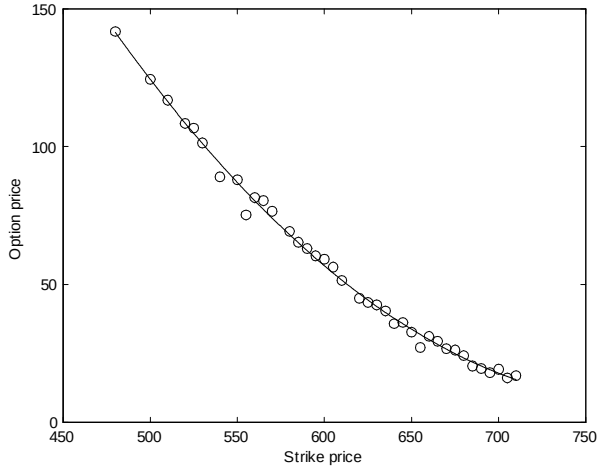

Figure 5.1: Google call option prices with the fitted second degree polynomial.

We calculate the residuals associated with the regression. Outliers in the option prices are identified by identifying outliers in the residuals using a confidence interval, this interval is derived below.

Consider the residuals from a given group of options with a common time to maturity. Denote the number of residuals in this group by n and the j^{th} residual by r_j; $j = 1, 2, ..., n$. Assume that the residuals are independently and identically distributed normal random variables with a mean of 0 and standard deviation of σ. Let $\alpha \in (0, 1)$ be the probability that there is at least one residual in the group not contained in some interval $[-c, c]$ with $c > 0$. Let A be the event that the group considered contains no such residuals;

$$P(A) = 1 - \alpha. \tag{5.3}$$

From independence

$$\begin{aligned} P(A) &= \prod_{j=1}^{n} P(-c \leq r_j \leq c) \\ &= [P(-c \leq r_1 \leq c)]^n \\ &= \left[2\Phi\left(\frac{c}{\sigma}\right) - 1\right]^n, \end{aligned} \tag{5.4}$$

where Φ denotes the standard normal distribution function. Equating (5.3) and (5.4) we obtain

$$\alpha = 1 - \left[2\Phi\left(\frac{c}{\sigma}\right) - 1\right]^n,$$

which can be solved for c:

$$c = \sigma\Phi^{-1}\left(\frac{1}{2} + \frac{1}{2}(1 - \alpha)^{1/n}\right). \tag{5.5}$$

In the argument used above we assume that σ is known. However, this is not the case. In order to implement this procedure we estimate σ using the standard deviation of the realised residuals. This value is used in equation (5.5) in the place of σ to obtain an estimate for c (denoted $\hat{c}$). Any residual not contained in the interval $[-\hat{c}, \hat{c}]$ is deemed an outlier and the corresponding option is removed from the dataset. We implement the method with $\alpha = 0.01$. Figure 5.2 shows the residuals associated with the call options on the Google shares with a time to maturity of 160 days. The horizontal lines represent $\hat{c}$ and $-\hat{c}$ respectively.

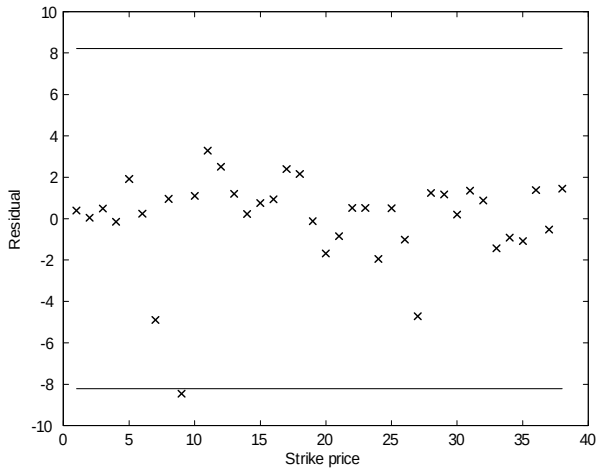

Figure 5.2: Residuals together with an estimated 99% confidence interval.

In this example a single residual lies outside the interval $[-\hat{c}, \hat{c}]$. This residual is deemed an outlier and the corresponding option price is removed from the dataset.

Duplicated options

The prices of some of the recorded options are not unique, i.e. there are multiple options with equal strike prices and times to maturity but different prices. It is not possible for an option to have more than one price, since this would lead to an instant arbitrage opportunity. If an option had two different prices traders would buy the option at the cheaper price and sell it at the more expensive to realise an immediate risk free profit. As a result, the demand for the cheaper option and the supply of the more expensive option

would increase. This in turn would force the price of the cheaper option to rise, while the price of the more expensive option would decrease. This process would continue until the option prices were equal.

A European call (put) option, with given time to maturity, is a decreasing (increasing) function of the strike price. As before, if these conditions do not hold, it will lead to an arbitrage opportunity. It is therefore reasonable to enforce this condition in the data used. All options with multiple prices are identified. The prices of these options are compared to the prices of the options with strike prices immediately smaller and larger than their own strike price (where all options considered have the same time to maturity). In the case of a call (put) option with multiple prices, the option is removed from the dataset if its price is larger (smaller) than that of the option with strike price immediately larger than its own, or if its price is smaller (larger) than that of the option with strike price immediately smaller than its own.

The procedure described above does not remove all of the duplicates in the data. To remove the remaining duplicates we consider the recorded open interest of these options; the open interest of an option is a measure of how actively the option is traded. We reason that the more actively an option is traded, the less likely it is that the recorded price of the option will be incorrect. As a result, for each unique strike price and time to maturity, all of the duplicated options are removed save the one with the largest open interest.

The resulting data

After applying the methods described above the resulting datasets are as follows. The S&P 500 dataset consists of 430 call options and 605 put options, while the PowerShares dataset contains 293 call options and 281 put options. The Google Inc. dataset consists of 519 call options and 445 put options.

5.1.2 Historical stock and index prices

The historical prices of the stock and the indexes underlying the options were obtained from http://finance.yahoo.com.

When deciding on the duration of the historical period in which to consider the stock or index prices a balance must be struck between the desire for a large dataset and the relevance of the observed prices. The historical period considered is used to gain information on the (constantly changing) current market conditions. If this period extends too far into the past it is possible that the market conditions have changed considerably within this historical period. In the empirical results below we use historical stock and index prices for a period of one year prior to the date on which the option prices were recorded.

5.2 Empirical option pricing results

In this section we calibrate the exponential $N \circ IG$ model and the Black-Scholes model to the option prices discussed in Section 5.1. In the discussion below we consider the exponential $N \circ IG$ model in detail. The results pertaining to the Black-Scholes model are shown at the end of the chapter.

We distinguish between two methods of fitting a given model to a dataset; estimation and calibration. Using an estimation procedure entails estimating the objective probability measure P, and affecting a change of measure (such as the Esscher transform or the MCMM) to obtain a pricing measure Q. When using this approach no free parameters are available that can be manipulated to reduce the AAE.

We use an estimation procedure to obtain parameter estimates for the exponential $N \circ IG$ model by fitting a $N \circ IG$ distribution to the log-returns of the underlying stock or index. In the results obtained below we use the indirect maximum likelihood estimator in order to estimate the parameters of the $N \circ IG$ distribution, for an explanation of this estimator see Section 3.3.1. In the case of the Black-Scholes model maximum likelihood is also used in order to estimate the parameters of the normal distribution. After the

parameters of the underlying distribution have been estimated we use the Esscher transform as well as the MCMM (both discussed in Chapter 4) in order to obtain the pricing measure used and we calculate the corresponding option prices.

We also use various calibration procedures. This entails choosing the parameters in the model so that the option prices calculated under the model correspond as closely as possible to the market prices of the options. The distance measure that we use for calibration is the AAE, see (2.13). Below we distinguish between calibration procedures based on the number of free parameters available. We consider three calibrations with a single free parameter. Thereafter we consider two calibrations where all of the parameters of the $N \circ IG$ distribution are allowed to vary freely save one. This parameter is used to ensure the martingaleness of the discounted price process using either the Esscher transform or the MCMM. We refer to this type of calibration as martingale restricted calibration. Finally we consider what we refer to as full calibration. Here all of the parameters are varied in an attempt to find the minimum possible AAE (disregarding the martingaleness requirement).

For the sake of brevity, we discuss in detail only the results relating to the call options available on the S&P 500 index on 11 May 2012 below. The results obtained for the remaining datasets can be found in Appendix D.

Figure 5.3 shows the market prices of the options under consideration as a function of the strike price.

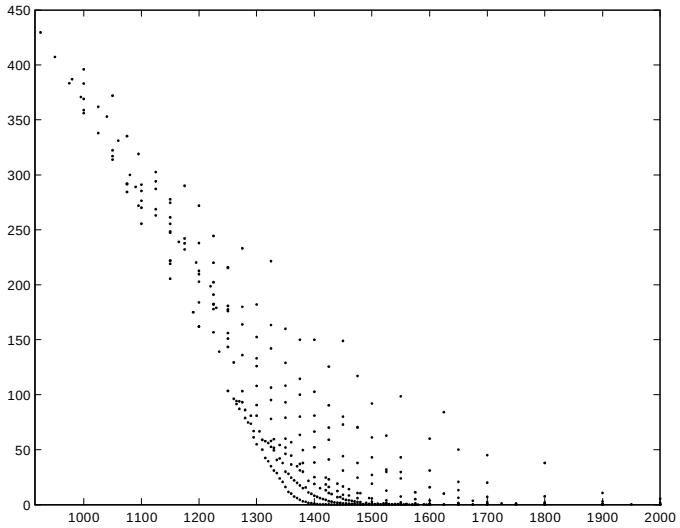

Figure 5.3: Market option prices as a function of strike price.

Because of the large number of options considered we choose to forgo the use of visual tests to aid our assessment of the goodness of fit of a given model in this chapter, instead we rely solely on the calculated distance measures.

Below we consider the estimation of the parameters of the distribution underlying the log-returns of the index before considering the model estimation and calibration results.

5.2.1 Parameter estimation

Parameter estimation for the $N \circ IG$ distribution is discussed in Section 3.3. We fit a $N \circ IG$ distribution to one year's daily log-returns calculated from the index price using the indirect maximum likelihood estimator, see Section 3.3.1. The possible starting values considered for each of the parameter estimates are the same as those used in Section 3.3.3.

The parameter estimates calculated using this method are

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (49.673, -6.5019, 0.0014, 0.0107)\,. \tag{5.6}$$

Figure 5.4 shows a density estimate of the log-returns as solid line; the estimated $N \circ IG$ density is superimposed as a dashed line.

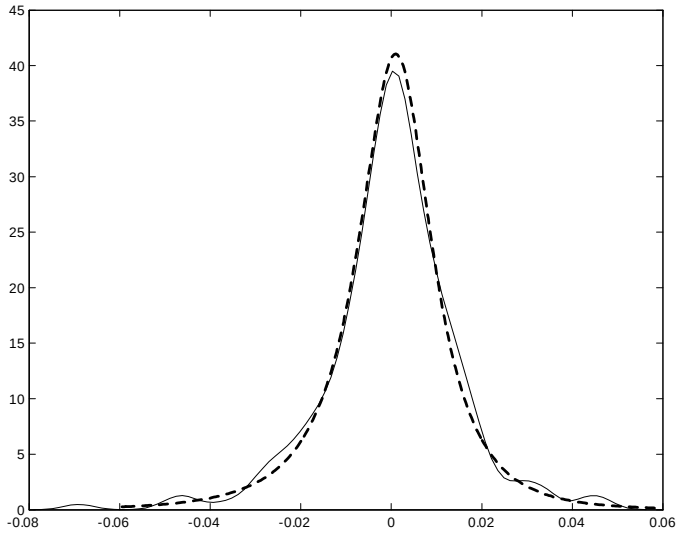

Figure 5.4: Density estimate of the log-returns (solid line) with the estimated $N \circ IG$ density (dashed line) superimposed.

Cont (2001:223-236) advocates the use of distributions with at least four parameters to realistically model the location, scale, asymmetry and tail behaviour of the log-returns. Table 5.1 shows the average, variance, skewness and kurtosis of the log-returns as well as those of the fitted $N \circ IG$ distribution.

	Log-returns	Fitted $N \circ IG$ distribution
Average	1.3923×10^{-5}	1.3922×10^{-5}
Variance	2.2211×10^{-4}	2.1311×10^{-4}
Skewness	-0.541	-0.4985
Kurtosis	9.084	6.0154

Table 5.1: The moments of the log-returns and those of the fitted $N \circ IG$ distribution.

5.2.2 Estimation techniques

Application of the estimation techniques requires that the parameters of the model be estimated from the log-return data. These parameters are changed so as to ensure the

martingaleness of the discounted index price using the Esscher transform and the MCMM.

The Esscher transform

Applying the Esscher transform to the parameters reported in (5.6) results in setting the value of β to -7.0377 while the remaining parameters are left unchanged. We use the resulting parameter set to calculate the prices of each of the options considered under the exponential $N \circ IG$ model. The values of the various distance measures between the market prices and the prices calculated using the model are as follows:

$$AAE = 16.9058, \quad ARE = 3.5015, \quad RMSE = 23.1673. \tag{5.7}$$

The application of the Esscher transform took 0.359 seconds.

The mean correcting martingale measure

Next we consider the MCMM. When applying this change of measure to the parameters in (5.6) the location parameter μ is set to 0.0013. The remaining parameters are not affected by this change of measure. The values of the distance measures between the option prices calculated using these parameters and the market option prices are as follows:

$$AAE = 16.8282, \quad ARE = 3.4935, \quad RMSE = 23.0692. \tag{5.8}$$

Applying the MCMM took 0.28 seconds.

Remarks

When comparing the distance measures reported in (5.7) to those in (5.8) we see that the distance measures obtained using the Esscher transform and the MCMM are similar. The application of each of these measure changes took less than a second.

5.2.3 Calibrating a single parameter

The application of the Esscher transform and the MCMM brings about a change in one of the parameters using a predetermined formula. In the case of the Esscher transform the value of β is changed, while the application of the MCMM adjusts the value of μ. Below we endeavour to determine whether or not the AAE can be substantially reduced by disregarding the martingaleness requirement and varying the value of β or that of μ reported in (5.6) freely.

In this section we also consider the GMCMM. Under this measure θ is a free parameter that can be calibrated to the observed option prices.

Calibrating the value of β

In order to assess whether or not the AAE can be substantially reduced by adjusting the value of β we proceed as follows. We specify a grid of possible values for β and we calculate the AAE corresponding to each of the β values (together with the values of α, μ and δ reported in (5.6)) in turn. Figure 5.5 shows the effect of varying the value of β on the AAE.

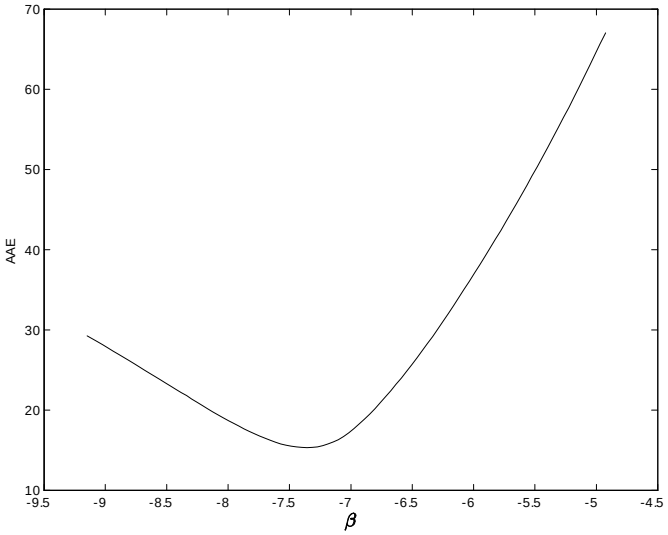

Figure 5.5: AAE as a function of β.

The smallest AAE is achieved by setting $\beta = -7.3576$. The distance measures between the option prices calculated using this value and the market prices are as follows:

$$AAE = 15.3201, \quad ARE = 3.0706, \quad RMSE = 21.0684. \tag{5.9}$$

All three of the distance measures reported in (5.9) are lower than those reported in (5.7). However, the decreases in the distance measures are not large. We conclude that, in this example, the AAE cannot be materially decreased by calibrating the value of β.

The application of this procedure took 30.514 seconds.

Calibrating the value of μ

The MCMM changes the value of μ reported in (5.6) in order to ensure the martingaleness of the discounted price process under the model. Similar to the procedure used in the previous calibration we calculate the value of the AAE for a grid of μ values. Figure 5.6 shows the AAE as a function of μ.

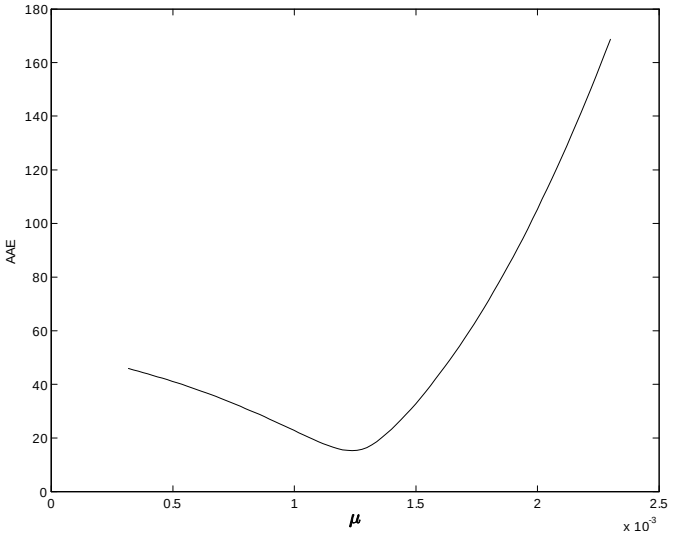

Figure 5.6: AAE as a function of μ.

The minimum AAE is achieved when $\mu = 0.0012$. At this value for μ we obtain the following distance measures:

$$AAE = 15.2577, \quad ARE = 3.0636, \quad RMSE = 21.0684. \tag{5.10}$$

As was the case with the β calibration, the calibration of μ reduces the distance measures considered only slightly.

Calibrating the value of μ took 31.154 seconds.

The generalised mean correcting martingale measure

We now consider the application of the generalised mean correcting martingale measure (GMCMM), see Section 4.4. As was the case with the β calibration and the μ calibration, one free parameter is available in the model. Again we define a grid of possible values for θ. Figure 5.7 shows the AAE as a function of θ.

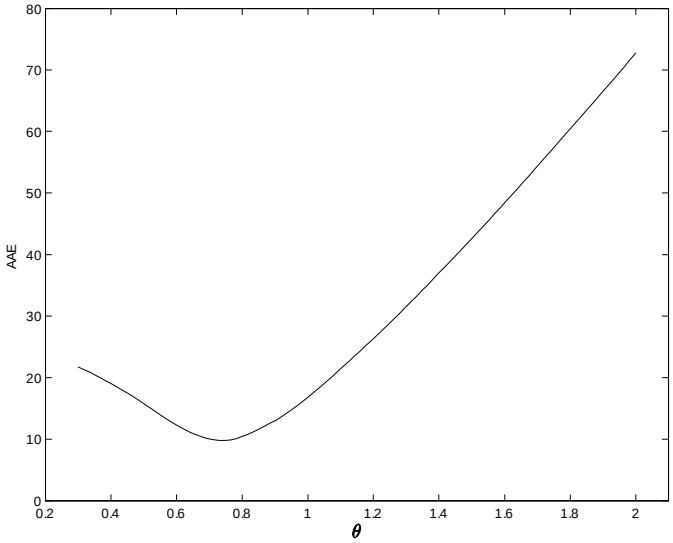

Figure 5.7: AAE as a function of θ.

The minimum AAE is achieved when $\theta = 0.7465$. The application of the GMCMM affects each of the parameters. Applying this change of measure results in the following parameter

set used for option pricing:

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (66.5443, -8.7103, 0.001, 0.008)\,.$$

The corresponding distance measures are

$$AAE = 9.771, \quad ARE = 0.7973, \quad RMSE = 16,9016. \tag{5.11}$$

The distance measures shown in (5.11) are substantially smaller that those shown in (5.9) and (5.10).

It took 28.891 seconds to apply this technique.

Remarks

In each of the three calibration procedures considered above a single free parameter is available, the value of which can be varied in order to lower the AAE. Neither the β calibration nor the μ calibration result in material reductions of the AAE. However, the GMCMM produces substantially smaller distance measures.

Note that the parameter set obtained using the GMCMM satisfies the martingaleness requirement imposed on the discounted price process, while the β calibration and the μ calibration do not satisfy this requirement.

5.2.4 Martingale restricted calibration

In the estimation procedures and the single parameter calibrations considered above the parameters used are obtained by fitting a $N \circ IG$ distribution to the log-returns of the price process and changing these parameters in some way. In the calibrations considered below we ignore the distribution of the log-returns and choose the parameters so that the AAE is as small as possible. However, the requirement that the discounted index price forms a martingale is enforced; we consider Esscher restricted calibration as well as MCMM restricted calibration.

Esscher restricted calibration

Esscher restricted calibration entails minimising the AAE by freely varying α, μ, and δ, while β is calculated using the Esscher transform. To achieve this minimisation we use Matlab's fminsearch.m. In order to employ the optimisation procedure we require starting values. To obtain these values we proceed in a way similar to that used to obtain starting values for parameter estimation. We define a grid of possible starting values for α, μ and δ. For each possible combination of these values we calculate β using the Esscher transform as well as the corresponding AAE. The parameter set resulting in the smallest AAE is used as starting values for the optimisation procedure. The possible starting values considered for α, μ, and δ are the same as those used in Section 3.3.3. In the remainder of this chapter we denote the starting value obtained for some parameter γ by γ_s. The starting values obtained using this procedure are

$$(\alpha_s, \beta_s, \mu_s, \delta_s) = (5, -0.4704, 0, 0.001)\,. \tag{5.12}$$

The distance measures calculated using the starting values are as follows:

$$AAE = 13.5186, \quad ARE = 6.3216, \quad RMSE = 19.2273. \tag{5.13}$$

The time taken to obtain these starting values is 92.851 seconds.

Using the starting values reported in (5.12) we minimise the AAE using the method described above. The parameters obtained are as follows:

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (14.839, -10.8948, 0.0008, 0.0008)\,.$$

The distance measures corresponding to this parameter set are

$$AAE = 7.7669, \quad ARE = 0.2778, \quad RMSE = 14.5842. \tag{5.14}$$

The calibration procedure took 142.506 seconds.

MCMM restricted calibration

In the MCMM restricted calibration the calibration procedure used is similar to that used in the Esscher restricted calibration; the values of α, β and δ are varied while μ is calculated using the MCMM. As before we require starting values for this optimisation procedure. Again, the possible starting values for α, β and δ are those reported in Section 3.3.3. We use the MCMM to calculate the value of μ for each of the parameter combinations and we calculate the corresponding AAE. The parameter set resulting in the smallest AAE is used as starting values. The starting values obtained in this manner are

$$(\alpha_s, \beta_s, \mu_s, \delta_s) = (20, 10 - 0.0006, 0.001)\,.$$

The corresponding distance measures are

$$AAE = 13.7018, \quad ARE = 2.2615, \quad RMSE = 21.9438. \tag{5.15}$$

Obtaining starting values for this calibration took 73.008 seconds.

Using these starting values the calibration procedure results in the following parameter set:

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (15.5674, 11.4427, 0.0009, 0.0009)\,.$$

The corresponding distance measures are

$$AAE = 7.7668, \quad ARE = 0.2829, \quad RMSE = 14.6017. \tag{5.16}$$

The calibration procedure took 138.637 seconds.

Remarks

The distance measures obtained using the two methods of calibration are similar. However, these distance measures are substantially lower than those obtained using the single parameter calibrations. This is to be expected because more free parameters are available, the values of which can be adjusted in order to obtain smaller distance measures.

5.2.5 Full calibration

Below we consider a calibration procedure in which neither local equivalence nor martingaleness is enforced. The algorithm used in this section allows all four of the parameters of the $N \circ IG$ distribution to vary freely in an attempt to minimise the AAE.

As before we require starting values for the calibration procedure and we calculate the AAE for each combination of a grid of possible starting values (specified in Section 3.3.3). The parameter set associated with the smallest AAE is chosen as the starting values. Using this procedure we arrive at the following starting values for the optimisation procedure:

$$(\alpha_s, \beta_s, \mu_s, \delta_s) = (50, 10, 0, 0.001)\,.$$

The distance measures obtained using the starting values are as follows:

$$AAE = 14.3137, \quad ARE = 0.4861, \quad RMSE = 26.8601. \tag{5.17}$$

Note that the AAE reported in (5.17) is similar to the AAEs reported in (5.13) and (5.15). However, calculating these starting values took 9444.506 seconds (more than two and a half hours), while calculating the starting values for the Esscher restricted calibration and the MCMM restricted calibration took less than two minutes each. In light of the time required and the distance measures obtained we opt to simply discard these starting values and we choose to use the starting values used for the Esscher restricted calibration. The same convention is used for the remaining datasets.

Using the starting values reported in (5.12) and Matlab's fminsearch.m we attempt to minimise the AAE by varying all four parameter values. The resulting parameters are

$$\left(\hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}\right) = (115.9247, 2.3349, -0.0001, 0.0088)\,.$$

The distance measures corresponding to this parameter set are

$$AAE = 8.2198, \quad ARE = 0.3288, \quad RMSE = 15.56.$$

This calibration procedure took 771.5 seconds to complete.

Remarks

In this calibration there are more free parameters than was the case when using martingale restricted calibration. In spite of this the AAE obtained here is slightly larger than the AAEs obtained using martingale restricted calibration. This can be explained by noting that the optimisation algorithm used converges to a local minimum. We conclude that, in this example, it is not possible to materially improve the fit of the model by discarding the martingaleness requirement imposed on the discounted price process.

The full calibration procedure used is time-consuming and the numerical results obtained using this method generally do not outperform those associated with martingale restricted calibration. As a result, we use martingale restricted calibration as our preferred calibration procedure in the remainder of the thesis.

5.2.6 Black-Scholes model results

Table 5.2 shows the parameters obtained using the various estimation and calibration procedures for the Black-Scholes model.

	μ	σ
Estimated parameters	0	0.0146
Esscher transform	-0.0001	0.0146
μ calibration	-0.0002	0.0146
GMCMM	-0.0001	0.0109
Esscher calibration (start)	-0.0001	0.012
Esscher calibration	-0.0001	0.011
Full calibration (start)	0.0001	0.008
Full calibration	0	0.009

Table 5.2: Parameters obtained using the Black-Scholes model.

As was the case with the previous model considered we are mainly interested in the various distance measures obtained using the values reported in Table 5.2. Table 5.3 shows the distance measures obtained using the Black-Scholes model together with the times taken to apply the various techniques.

	AAE	ARE	$RMSE$	Time
Esscher transform	16.269	3.3055	22.3107	0.234
μ calibration	15.0098	2.8985	20.8142	22.099
GMCMM	9.8906	0.7487	17.1587	23.478
Esscher calibration (start)	10.7389	1.2802	16.8804	1.997
Esscher calibration	9.8825	0.7812	17.0802	5.694
Full calibration (start)	8.9814	0.3196	17.4187	11.577
Full calibration	8.2804	0.3776	15.5019	13.542

Table 5.3: Empirical results obtained using the Black-Scholes model.

As before the AAE is reduced by the various calibration techniques. Under the Black-Scholes model the application of the GMCMM as well as the Esscher restricted calibration allow a single parameter to vary. Note that the distance measures associated with these calibration techniques are similar. The full calibration procedure is much less problematic when applied to the Black-Scholes model than when applied to the exponential $N \circ IG$ model. The reason for this is closely connected to the number of parameters in each of the two models.

The results pertaining to the remaining datasets can be found in Appendix D.

Chapter 6

Time changed option pricing models

In this chapter we consider time changed models used in financial modelling. These models differentiate between calender time and operational or business time. The premise is that the pace at which trading takes place, measured for example in the number of trades on a certain stock per hour, is not constant. In order to mimic this phenomenon models are used in which operational or business time is modelled as a non-decreasing stochastic process. This type of model is an alternative to stochastic volatility models.

Time changed financial models were first proposed in a paper entitled "A subordinated stochastic process model with finite variance for speculative prices" by Clark (1973:135-156). In the opinion of the author this important paper is often overlooked because of the way in which some of the content is presented. In Section 6.1 we provide an overview of some of the main results in this paper, including alternative proofs for several of these results.

After considering the contributions made in Clark's paper we define various time changed models used to obtain numerical results in Chapter 7. The time changed linear Brownian motion models used are defined in Section 6.2, while the time changed exponential Brownian motion models used are defined in Section 6.3.

6.1 An overview of Clark (1973:135-156)

Clark proposes the use of a subordinated (or time changed) stochastic process as a model for financial time series. He fits the proposed model to cotton futures prices and shows that this model fits these data better than the stable distributions proposed in Mandelbrot

(1963:394-419). Some interesting theoretical results are presented in the paper, notably a generalisation of the central limit theorem used to obtain the marginal distribution of the increments in the stock price under Clark's model.

Below we are mainly interested in the theoretical results presented. We refer the interested reader to Clark (1973:135-156) for the numerical results obtained.

6.1.1 Introduction and motivation

Interestingly, Clark does not model the log-returns of financial time series, instead he uses first differences. Let S_t denote the price at time t of a stock and let $\Delta S_t = S_t - S_{t-1}$. Empirical studies indicate that on short time intervals (for example daily intervals)

$$E\left[\Delta S_t\right] = 0, \tag{6.1}$$

and

$$E\left[\Delta S_t \Delta S_u\right] = 0, \tag{6.2}$$

for all $t \neq u$. Arbitrage free option pricing theory requires that the discounted value of the stock price form a martingale under the pricing measure used. Assuming a positive interest rate, this implies that the expected change in the stock price should be positive. However, the expected price change is small and can be ignored for the sake of this analysis. Alternatively the reader may think of S_t as the discounted stock price. A more recent reference on the empirical properties of financial returns is Cont (2001:223-236).

It is well known that changes in the stock price on a short time interval are not normally distributed; rather ΔS_t follows some leptokurtic distribution. This is evident when considering the kurtosis of ΔS_t which empirical studies find to be much larger than the kurtosis of the normal distribution. However, ΔS_t is the sum of a large number of small price changes. The non-normality of ΔS_t indicates that the necessary conditions for the central limit theorem are not satisfied by the small price changes that add up to ΔS_t. The aim of Clark's paper is to investigate which of these conditions are not satisfied.

A random walk model can be used to account for the properties of financial time series presented in (6.1) and (6.2). This model also has a theoretical justification provided by Bachelier, see Davis and Etheridge (2006:27-28). The basic argument is that if the changes in the stock price were indeed correlated speculators could use this information to buy and sell at the correct time in order to realise a profit. For example, consider a market in which ΔS_t and ΔS_{t+1} are negatively correlated. An alert speculator with access to this information and seeing that ΔS_t is positive would know that ΔS_{t+1} will likely be negative. At time t she would sell the stock in order to realise a profit. However, if many speculators use this trading strategy the demand for the stock will decrease at time t while the supply will increase. As a result, the price of the stock will increase, this increase in price will continue until the correlation between the increments is removed.

From the arguments above it seems reasonable to model the stock price using a random walk:

$$S_t = S_{t-1} + \varepsilon_t,$$

where

$$E[\varepsilon_t] = 0,$$

$$E[\varepsilon_s \varepsilon_t] = 0,$$

for all $t \neq s$, and where ε_t follows some leptokurtic distribution.

Mandelbrot (1963:394-419) provides a possible explanation for the non-normality of ΔS_t by arguing that the small price changes adding up to ΔS_t are independent, but have infinite variance. In this case the random walk model can be used with ε_t following a stable distribution (which is the limit distribution of the sum of independent random variables).

Clark presents a different hypothesis explaining the distributional properties of ΔS_t. He asserts that S_t evolves at different rates during identical time intervals. Clark's hypothesis is that the distribution of ΔS_t is not normal, but subordinate to a normal distribution.

This means that the number of small price changes summed together to obtain ΔS_t is a random variable. Under Clark's hypothesis ε_t should be modelled as the sum of a random number of normally distributed random variables. The central limit theorem does not guarantee the normality of ΔS_t in this case.

6.1.2 Generalisations of the central limit theorem

Below we consider two generalisations of the central limit theorem in which the number of random variables summed together are variable. Both of these generalisations are special cases of Anscombe's theorem, see Anscombe (1949:455-458). However, the special cases presented here are much simpler to prove than Anscombe's theorem.

Theorem 28 *Let $N_1, N_2, ...$ be a sequence of positive integral valued random variables such that*

$$\lim_{n\to\infty}\left(\frac{N_n}{n}\right) = 1,$$

where the equality holds in probability. Let $Y_1, Y_2, ...$ be a sequence of independent and identically distributed random variables with mean 0 and variance 1, independent of the sequence $N_1, N_2, ...$ If

$$X_{N_n} = \sum_{j=1}^{N_n} Y_j,$$

then $\frac{X_{N_n}}{\sqrt{n}}$ converges in distribution to the standard normal distribution as $n \to \infty$.

For a proof of this theorem see Feller (1971:265). This theorem shows that if N_n has small variation around n for large values of n, then the distribution of X_{N_n} approaches normality.

The second theorem considers the case where N_n has appreciable variance around n for large values of n. Below we denote the largest integer smaller than X by $\lfloor X \rfloor$.

Theorem 29 *Let $Y_1, Y_2, ...$ be a sequence of independent and identically distributed random variables with mean 0 and variance 1. Let $N_n = \lfloor Zn \rfloor$ for large n, where Z is a random*

variable with mean 1 *and variance* $\Gamma > 0$ *independent of* $Y_1, Y_2, ...$ *If*

$$X_{N_n} = \sum_{j=1}^{N_n} Y_j,$$

then the conditional distribution of $\frac{X_{N_n}}{\sqrt{n}}$ *given* Z *converges to a normal distribution with mean* 0 *and variance* Z*.*

Proof. Denote the characteristic function of a random variable X by ϕ_X and let

$$A_n(t) = \phi_{\frac{X_{N_n}}{\sqrt{n}}}(t).$$

In order to show that the distribution of $\frac{X_{N_n}}{\sqrt{n}}$ converges to that of a $N(0, Z)$ random variable, it suffices to show that

$$\lim_{n\to\infty} A_n(t) = \exp\left(-\frac{1}{2}Zt^2\right).$$

Consider

$$\begin{aligned}
A_n(t) &= \phi_{\frac{X_{N_n}}{\sqrt{n}}}(t) \\
&= E\left[\exp\left(\frac{it}{\sqrt{n}}\sum_{j=1}^{N_n} Y_j\right)\right] \\
&= E\left[E\left[\exp\left(\frac{it}{\sqrt{n}}\sum_{j=1}^{N_n} Y_j\right)\middle| N_n\right]\right] \\
&= E\left[\left(\phi_{Y_1}\left(\frac{t}{\sqrt{n}}\right)\right)^{N_n}\middle| N_n\right] \\
&= E\left[\exp\left(\frac{N_n}{n} n \log\left(\phi_{Y_1}\left(\frac{t}{\sqrt{n}}\right)\right)\right)\middle| N_n\right]. \qquad (6.3)
\end{aligned}$$

Taking limits, (6.3) becomes

$$\begin{aligned}
\lim_{n\to\infty} A_n(t) &= \lim_{n\to\infty} E\left[\exp\left(\frac{N_n}{n} n \log\left(\phi_{Y_1}\left(\frac{t}{\sqrt{n}}\right)\right)\right)\middle| N_n\right] \\
&= E\left[\exp\left(\lim_{n\to\infty}\left\{\frac{N_n}{n}\right\}\lim_{n\to\infty}\left\{n\log\left(\phi_{Y_1}\left(\frac{t}{\sqrt{n}}\right)\right)\right\}\right)\middle| N_n\right] \\
&= E\left[\exp\left(-\frac{1}{2}Zt^2\right)\middle| N_n\right] \\
&= \exp\left(-\frac{1}{2}Zt^2\right), \qquad (6.4)
\end{aligned}$$

where we used dominated convergence to obtain the second equality above; from the properties of characteristic functions

$$|A_n(t)| = \left|\phi_{\frac{X_{N_n}}{\sqrt{n}}}(t)\right| \leq 1. \qquad (6.5)$$

In the third equality of (6.4), we used

$$\lim_{n\to\infty}\left\{n\log\left(\phi_{Y_1}(t)\right)\right\} = -\frac{1}{2}t^2,$$

which follows from the central limit theorem.

This completes the proof. ■

6.1.3 Subordinated stochastic processes

Clark explains subordinated stochastic processes as follows. Instead of indexing a stochastic process $X_1, X_2, ...$ by an increasing sequence of integers the process can be indexed as follows; $X_{t_1}, X_{t_2}, ...$ where $t_1 \leq t_2 \leq ...$ and $t_1, t_2, ...$ are realisations of a stochastic process with non-negative increments. We can therefore form a new stochastic process X_{T_t}, where T_t is a non-decreasing stochastic process. This process is called subordinated to X_t and T_t

is called the directing process. The distribution of ΔX_{T_t} is subordinate to that of ΔX_t. We consider some distributional properties of ΔX_{T_t} below.

Theorem 30 *Let X_t and T_t be processes, independent of each other, each with stationary and independent increments. If*

$$\begin{aligned} E\left[\Delta X_t\right] &= 0, \\ Var\left[\Delta X_t\right] &= \sigma^2, \end{aligned} \tag{6.6}$$

for all $t \geq 0$, and

$$\begin{aligned} E\left[\Delta T_t\right] &= \alpha, \\ T_{t+1} - T_t &\geq 0, \end{aligned} \tag{6.7}$$

for all $t \geq 0$, then the subordinated stochastic process X_{T_t} has stationary and independent increments, and

$$\begin{aligned} E\left[\Delta X_{T_t}\right] &= E\left[X_{T_{t+1}} - X_{T_t}\right] = 0, \\ Var\left[\Delta X_{T_t}\right] &= Var\left[X_{T_{t+1}} - X_{T_t}\right] = \alpha\sigma^2. \end{aligned} \tag{6.8}$$

Proof. The stationarity and independence of the increments of X_{T_t} follow directly from the stationarity and the independence of the increments of X_t and T_t. Consider the mean

of ΔX_{T_t};

$$\begin{aligned} E\left[\Delta X_{T_t}\right] &= E\left[E\left[X_{T_{t+1}} - X_{T_t}|T_t, T_{t+1}\right]\right] \\ &= 0. \end{aligned}$$

Consider the variance of ΔX_{T_t};

$$\begin{aligned} Var\left[\Delta X_{T_t}\right] &= E\left[Var\left[X_{T_{t+1}} - X_{T_t}|T_t, T_{t+1}\right]\right] \\ &= \sigma^2 E\left[T_{t+1} - T_t\right] \\ &= \alpha\sigma^2. \end{aligned}$$

■

Note that the variance of ΔX_{T_t} is not affected by the variance of ΔT_t. Therefore, it is possible to obtain a family of distributions with mean 0 and identical variance by changing the variance of ΔT_t. Clark states and proves the corollary below to the previous theorem.

Corollary 31 *If X_t and T_t are defined as in Theorem 30 and ΔX_t is normally distributed, then the kurtosis of ΔX_{T_t} is an increasing function of the variance of ΔT_t.*

Proof.

$$\begin{aligned} kurt\left[\Delta X_{T_t}\right] &= \frac{E\left[\left(X_{T_{t+1}} - X_{T_t}\right)^4\right]}{E^2\left[\left(X_{T_{t+1}} - X_{T_t}\right)^2\right]} \\ &= \frac{1}{\alpha^2\sigma^4} E\left[E\left[\left(X_{T_{t+1}} - X_{T_t}\right)^4 |T_t, T_{t+1}\right]\right]. \end{aligned} \tag{6.9}$$

Since $X_{T_{t+1}} - X_{T_t}$ given T_t and T_{t+1} is a normal random variable,

$$\frac{E\left[\left(X_{T_{t+1}} - X_{T_t}\right)^4 | T_t, T_{t+1}\right]}{E^2\left[\left(X_{T_{t+1}} - X_{T_t}\right)^2 | T_t, T_{t+1}\right]} = 3. \tag{6.10}$$

From (6.9) and (6.10)

$$\begin{aligned} kurt\left[\Delta X_{T_t}\right] &= \frac{3}{\alpha^2\sigma^4} E\left[E^2\left[\left(X_{T_{t+1}} - X_{T_t}\right)^2 | T_t, T_{t+1}\right]\right] \\ &= \frac{3}{\alpha^2} E\left[\left(T_{t+1} - T_t\right)^2\right] \\ &= \frac{3}{\alpha^2}\left(E^2\left[\Delta T_t\right] + Var\left[\Delta T_t\right]\right) \\ &= 3\left(1 + \frac{Var\left[\Delta T_t\right]}{\alpha^2}\right). \end{aligned}$$

■

The kurtosis of financial returns (and log-returns) are greater than that of the normal distribution. The corollary shows that this characteristic can be modelled by introducing a directing process T_t and that the kurtosis of ΔX_{T_t} can be manipulated by adjusting the variance of ΔT_t.

6.1.4 Processes subordinated to Brownian motion

We now consider the density function of a stochastic process subordinated to a Brownian motion.

Theorem 32 *Let X_t and T_t be processes, independent of each other, each with stationary and independent increments such that $\Delta X_t \sim N\left(\mu, \sigma^2\right)$ and ΔT_t has density function g.*

The density function of ΔX_{T_t} *is*

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_0^\infty \exp\left(-\frac{(x-\mu y)^2}{2\sigma^2 y}\right) \frac{g(y)}{\sqrt{y}} dy.$$

Proof. If F denotes the distribution function of ΔX_{T_t}, then

$$\begin{aligned} F(x) &= P\left(X_{T_{t+1}} - X_{T_t} \leq x\right) \\ &= P\left(X_{\Delta T_t} \leq x\right) \\ &= E\left[P\left(X_{\Delta T_t} \leq x | \, \Delta T_t\right)\right] \\ &= E\left[\Phi\left(\frac{x-\mu\Delta T_t}{\sigma\sqrt{\Delta T_t}}\right)\right], \end{aligned} \tag{6.11}$$

where Φ denotes the standard normal distribution function. Using (6.11) the density function of ΔX_{T_t} can be obtained as follows:

$$\begin{aligned} f(x) &= \frac{\partial}{\partial x} E\left[\Phi\left(\frac{x-\mu\Delta T_t}{\sigma\sqrt{\Delta T_t}}\right)\right] \\ &= E\left[\frac{\partial}{\partial x}\Phi\left(\frac{x-\mu\Delta T_t}{\sigma\sqrt{\Delta T_t}}\right)\right] \\ &= E\left[\frac{1}{\sigma\sqrt{\Delta T_t}}\phi\left(\frac{x-\mu\Delta T_t}{\sigma\sqrt{\Delta T_t}}\right)\right] \\ &= \frac{1}{\sigma\sqrt{2\pi}} \int_0^\infty \exp\left(-\frac{(x-\mu y)^2}{2\sigma^2 y}\right) \frac{g(y)}{\sqrt{y}} dy. \end{aligned}$$

■

Later in the thesis we use the lognormal, Pareto and gamma processes as directing processes. The three corollaries below provide the density functions of the resulting subordinated processes.

Corollary 33 *If $\Delta T_t \sim \log N(\alpha, \beta)$, then the density function of ΔX_{T_t} is*

$$f(x) = \frac{1}{2\pi\sigma\beta} \int_0^\infty y^{-\frac{3}{2}} \exp\left(-\frac{(x-\mu y)^2}{2\sigma^2 y} - \frac{(\log(y)-\alpha)^2}{2\beta^2} \right) dy.$$

If ΔT_t follows a lognormal distribution then X_{T_t} is said to be a lognormal-normal process.

Corollary 34 *If $\Delta T_t \sim Pareto(g, h)$, then the density function of ΔX_{T_t} is*

$$f(x) = \frac{hg^h}{\sigma\sqrt{2\pi}} \int_0^\infty \frac{1}{\sqrt{y}(g+y)^{h+1}} \exp\left(-\frac{(x-\mu y)^2}{2\sigma^2 y} \right) dy.$$

If ΔT_t follows a Pareto distribution then X_{T_t} is said to be a Pareto-normal process.

Definition 35 *A random variable is said to follow a gamma distribution with parameters $a > 0$ and $b > 0$ if it has density*

$$f(x) = \frac{x^{a-1}}{b^a \Gamma(a)} \exp\left(-\frac{x}{b}\right),$$

for $x > 0$.

Corollary 36 *If $\Delta T_t \sim gamma(a, b)$, then the density function of ΔX_{T_t} is*

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma b^a \Gamma(a)} \int_0^\infty y^{a-\frac{3}{2}} \exp\left(-\frac{(x-\mu y)^2}{2\sigma^2 y} - \frac{y}{b} \right) dy.$$

If ΔT_t follows a gamma distribution then X_{T_t} is said to be a gamma-normal process.

6.2 Time changed linear Brownian motion models

We now define time changed linear Brownian motion models. Under these models

$$S_t = \begin{cases} (S_0 + L_{T_t})e^{rt}, & 0 \leq t < \tau, \\ \\ 0, & \tau \leq t < \infty, \end{cases}$$

where T_t is a subordinator and L_t is a Brownian motion independent of T_t. τ is the first hitting time of 0 of S_t; $\tau = \inf\{t : S_t = 0\}$. A stock has value because it imparts partial ownership of the company that issued the stock to the holder. Economically τ can be interpreted as the time of default of the company under consideration.

In the next two chapters we use three different subordinators; the $\log N(\alpha, \beta)$ process, the $Pareto(g, h)$ process and the $gamma(a, b)$ process. We refer to the resulting models as the linear lognormal-normal model, the linear Pareto-normal model and the linear gamma-normal model.

Note that $e^{-rt}S_t$ forms a martingale under these models if and only if the location parameter of the Brownian motion is 0, see (6.8).

6.3 Time changed exponential Brownian motion models

Under time changed exponential Brownian motion models

$$S_t = S_0 \exp(L_{T_t}),$$

where T_t is a subordinator and L_t is a Brownian motion independent of T_t. As was the case for the time changed linear Brownian motion models, we use three different subordinators as before; the $\log N(\alpha, \beta)$, $Pareto(g, h)$ and $gamma(a, b)$ processes. We use naming conventions similar to those used above; the resulting models are referred to as the exponential lognormal-normal model, the exponential Pareto-normal model and the exponential

gamma-normal model.

In Chapter 7 we require knowledge of the value of μ that ensures the martingaleness of $e^{-rt}S_t$. The process $e^{-rt}S_t$ forms a martingale if and only if

$$
\begin{aligned}
& e^{-rt}E\left[S_t\right] = S_0 \\
\iff \quad & E\left[\exp\left(\Delta L_{T_t}\right)\right] = e^r \\
\iff \quad & E\left[E\left[\exp\left(L_{\Delta T_t}\right)\middle|\, \Delta T_t\right]\right] = e^r \\
\iff \quad & E\left[\exp\left(\mu \Delta T_t + \frac{1}{2}\sigma^2 \Delta T_t\right)\right] = e^r \\
\iff \quad & \int_0^\infty \exp\left(\left(\mu + \frac{1}{2}\sigma^2\right) x\right) f_{\Delta T_t}\left(x\right) dx = e^r, \qquad (6.12)
\end{aligned}
$$

where $f_{\Delta T_t}$ is the density of ΔT_t. If time is modelled using a lognormal or Pareto process the following restriction is required for the integral above to be well defined:

$$\mu + \frac{1}{2}\sigma^2 < 0.$$

If time is modelled using a $gamma\left(a, b\right)$ process this restriction changes to

$$\mu + \frac{1}{2}\sigma^2 - \frac{1}{b} < 0.$$

Using (6.12) we can solve for the value of μ that ensures the martingaleness of $e^{-rt}S_t$ numerically.

Chapter 7

Calibrating models to barrier option prices

In Schoutens *et al.* (2004:66-78) the authors calibrate various models to the market prices of European call options. The authors then go on to examine the prices of exotic options calculated using these calibrated models. In spite of the fact that these models provide nearly indistinguishable European call option prices the resulting exotic option prices vary substantially.

The authors do not calibrate the models to exotic option prices. One of the possible reasons for this is that the market prices of these options are not as readily available as the prices of European options. It remains unclear whether or not the models can be used to replicate the prices of exotic options with the same level of precision as is the case when using European call options. In this chapter we explore this possibility by considering the results when calibrating various models to barrier option prices. However, since barrier option prices are not as freely available as European option prices we calculate the prices of various sets of barrier option prices under a known model and we use these in the place of observed market option prices. We demonstrate that it is possible to very accurately replicate the option prices calculated under this model using a calibration procedure applied to various other models. This demonstrates that the models are interchangeable; similar option prices can be achieved under the various models.

In Section 7.1 we introduce the barrier options used in this chapter. The models used to obtain the empirical results are discussed in Section 7.2. We use Monte Carlo simulation to estimate option prices in this chapter and the next; in Section 7.3 we provide a confirmatory analysis demonstrating that the simulation procedures used give the expected results. In

Sections 7.4 and 7.5 we explain the method used in order to calculate option prices and we construct a hypothetical financial market (under which option prices are calculated using a known model). The calibration procedure used is explained in Section 7.6. Finally, Section 7.7 analyses the empirical results and Section 7.8 concludes.

7.1 Barrier options

The payoff of a barrier option depends on whether or not the stock price reaches some predetermined barrier before the maturity of the option. In practice barrier options are usually monitored at discrete time intervals. Throughout the thesis, the barrier options considered are assumed to be monitored daily; this means that the stock price at the end of each business day is used in order to assess whether or not the barrier level has been reached.

A digital barrier call option pays a fixed amount (which we take equal to 1 unit of currency) at maturity if the stock price reaches a predefined barrier level $H > S_0$ before the maturity of the option. If this barrier is not reached the option expires worthless. The payoff of a digital barrier call option with barrier level H and time to maturity T is

$$h = \mathbb{I}\left(\max_{0\leq j\leq T} S_j \geq H\right),$$

and the price of this option (calculated with respect to the pricing measure Q) is

$$\pi = e^{-rT} E^Q\left[\mathbb{I}\left(\max_{0\leq j\leq T} S_j \geq H\right)\right].$$

A down-and-out barrier call option has the same payoff as a European call option if the stock price does not reach or drop below a predefined barrier level $H < S_0$ in the lifetime of the option. If the stock price reaches this level the option becomes worthless. The payoff of a down-and-out barrier call option with strike price K, barrier level H and

time to maturity T is

$$h = \mathbb{I}\left(\min_{0\leq j\leq T} S_j > H\right)(S_T - K)^+,$$

and the price of the option (calculated with respect to the pricing measure Q) is

$$\pi = e^{-rT}E^Q\left[\mathbb{I}\left(\min_{0\leq j\leq T} S_j > H\right)(S_T - K)^+\right].$$

Similarly, an up-and-out barrier call option has the same payoff as a European call option if the stock price does not reach or exceed some barrier level $H > S_0$ before maturity, otherwise the payoff of the option is 0. The payoff of an up-and-out barrier call option with strike price K, barrier level H and time to maturity T is

$$h = \mathbb{I}\left(\max_{0\leq j\leq T} S_j < H\right)(S_T - K)^+,$$

and the price of the option (calculated with respect to the pricing measure Q) is

$$\pi = e^{-rT}E^Q\left[\mathbb{I}\left(\max_{0\leq j\leq T} S_j < H\right)(S_T - K)^+\right].$$

7.2 Option pricing models

Below we discuss the various option pricing models used in this chapter and the next. The calibration procedures used in Schoutens *et al.* (2004:66-78) ensure that the discounted stock price forms a martingale using the MCMM. We follow the same convention below for all of the models considered. This means that, when calibrating these models we allow each of the parameters to vary freely, with the exception of the location parameter. This parameter is chosen so as to ensure that the discounted stock price process forms a martingale.

The majority of the models used in this chapter are discussed in detail in the preceding chapters. However, for ease of reference we include a brief discussion of the different models below.

The algorithms used to simulate price paths under the models discussed below are shown in Appendix E.

7.2.1 The Heston model

This model was introduced in Heston (1993:327-343). The dynamics of the stock price under the Heston model are as follows:

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu dt + \sigma_t dW_t, \\ d\sigma_t^2 &= \kappa\left(\eta - \sigma_t^2\right) dt + \theta\sigma_t d\tilde{W}_t, \end{aligned} \tag{7.1}$$

where W_t and $\tilde{W}_t$ are two correlated standard Brownian motions with $Cov\left(dW_t d\tilde{W}_t\right) = \rho dt$. In the calibration procedure used to obtain the numerical results shown later in this chapter the value of the location parameter μ is set to r; this ensures the martingaleness of $e^{-rt}S_t$.

Under the Heston model the volatility process σ_t^2 follows a Cox-Ingersoll-Ross process. As a result, if $2\kappa\eta \geq \theta^2$ the volatility of the process is strictly positive and the stationary distribution of σ_t^2 is $gamma\left(\frac{2\kappa\eta}{\theta^2}, \frac{\theta^2}{2\kappa}\right)$. In the numerical analysis to follow we discard parameter sets for which the requirement that $2\kappa\eta \geq \theta^2$ is not met, meaning that the calibration procedure is not allowed to converge to a parameter set such that $2\kappa\eta < \theta^2$.

7.2.2 The time changed exponential $N \circ IG$ model

Under this model the stock price

$$S_t = S_0 \exp\left(L_{T_t}\right),$$

where L_t is a $N \circ IG(\alpha, \beta, \mu, \delta)$ process and the operational time T_t follows a Cox-Ingersoll-Ross process;

$$dT_t = \kappa(\eta - T_t)\,dt + \lambda\sqrt{T_t}dW_t, \tag{7.2}$$

where W_t is a standard Brownian motion. Under this model the MCMM sets the location parameter to

$$\mu = r + \delta\left(\sqrt{\alpha^2 - (\beta+1)^2} - \sqrt{\alpha^2 - \beta^2}\right).$$

Similar to the volatility of the Heston model, T_t is strictly positive and the stationary distribution of T_t is $gamma\left(\frac{2\kappa\eta}{\lambda^2}, \frac{\lambda^2}{2\kappa}\right)$ if $2\kappa\eta \geq \lambda^2$. As before, we only consider parameter sets for which the requirement that $2\kappa\eta \geq \lambda^2$ is met.

7.2.3 Exponential Lévy models

The exponential Lévy models used are discussed in Chapter 4. Under these models

$$S_t = S_0 \exp(L_t),$$

where L_t is some Lévy process (a Brownian motion or a $N \circ IG$ process). The formulae for the location parameter in these models are given in (4.9) and (4.10).

7.2.4 Linear Lévy models

Under a linear Lévy model the dynamics of S_t are given by

$$S_t = \begin{cases} (S_0 + L_t)\,e^{rt}, & 0 \leq t < \tau, \\ \\ 0, & \tau \leq t < \infty, \end{cases}$$

where L_t is a Brownian motion or a $N \circ IG$ process and τ is the first hitting time of 0 of S_t; $\tau = \inf\{t : S_t = 0\}$. τ is the time at which the company that issued the stock defaults. The

application of the MCMM entails setting $\mu = 0$ in the case where L_t follows a Brownian motion. In the case where L_t follows a $N \circ IG(\alpha, \beta, \mu, \delta)$ process the application of the MCMM entails setting

$$\mu = \frac{-\delta\beta}{\sqrt{\alpha^2 - \beta^2}}.$$

7.2.5 Time changed linear Brownian motion models

Under the time changed linear Brownian motion models

$$S_t = \begin{cases} (S_0 + W_{T_t})e^{rt}, & 0 \leq t < \tau, \\ \\ 0, & \tau \leq t < \infty, \end{cases}$$

where T_t is a subordinator and W_t is a Brownian motion independent of T_t. As before τ is the first hitting time of 0 of S_t; $\tau = \inf\{t : S_t = 0\}$. We consider three different subordinators in the models used below; the $\log N(\alpha, \beta)$ process, the $Pareto(g, h)$ process and the $gamma(a, b)$ process. We refer to the resulting models as the linear lognormal-normal model, the linear Pareto-normal model and the linear gamma-normal model. Choosing the location parameter of W so that $e^{-rt}S_t$ forms a martingale entails setting $\mu = 0$ for all three models considered here. For more details on these models, see Chapter 6.

7.2.6 Time changed exponential Brownian motion models

The dynamics of S_t under the time changed exponential Brownian motion models are

$$S_t = S_0 \exp(W_{T_t}),$$

where T_t is a subordinator and W_t is a Brownian motion independent of T_t. Again we use the $\log N(\alpha, \beta)$, $Pareto(g, h)$ and $gamma(a, b)$ processes as subordinators. We use similar naming conventions for these models as for their linear counterparts. The location parameter ensuring the martingaleness of the discounted stock price can be determined

numerically, see (6.12).

7.3 Confirmatory analysis

Below we include two tests used to confirm that the algorithms used provide the expected results. This is done so as to remove any doubt as to the validity of the results shown and the conclusions drawn based on the empirical results obtained. The first test is applied only to the Heston and time changed exponential $N \circ IG$ models. This test gives strong indications that the stationary distributions of the volatility (in the case of the Heston model) and the operational time (in the case of the time changed exponential $N \circ IG$ model) follow the gamma laws with the parameters specified in Section 7.2. The second test is applied to all of the models used and gives strong indications that the discounted stock prices satisfy the martingaleness requirement. Both tests are informal in nature.

7.3.1 Stationary distribution of volatility and operational time

In the algorithms provided in Appendix E we randomly generate starting values for the volatility (under the Heston model) and the operational time (under the time changed exponential $N \circ IG$ model) from the gamma laws that constitute the stationary distributions of these processes. As a result, it is important to confirm that the stationary distributions of these processes conform to the specified gamma laws. Below we consider the Heston and time changed exponential $N \circ IG$ models in turn.

The stationary distribution of the volatility follows a $gamma\left(\frac{2\kappa\eta}{\theta^2}, \frac{\theta^2}{2\kappa}\right)$ distribution under the Heston model. To confirm that the distribution of the volatility process, as simulated using the algorithm provided in Appendix E, converges to this distribution we perform the following test. We simulate 10000 price paths, each with a length of 111 days (which is the maximum time to maturity of the options considered). Figure 7.1 shows, as a dashed line, a density estimate of the terminal volatility of each of the 10000 paths simulated from the Heston model with parameter set $(r, \kappa, \eta, \rho, \theta) = (0.1/252, 0.1, 0.0001, -0.7, 0.001)$. This parameter set is used in Section 7.5 in the construction of a

hypothetical market. The density function of the stationary distribution of the volatility is superimposed in the figure as a solid line.

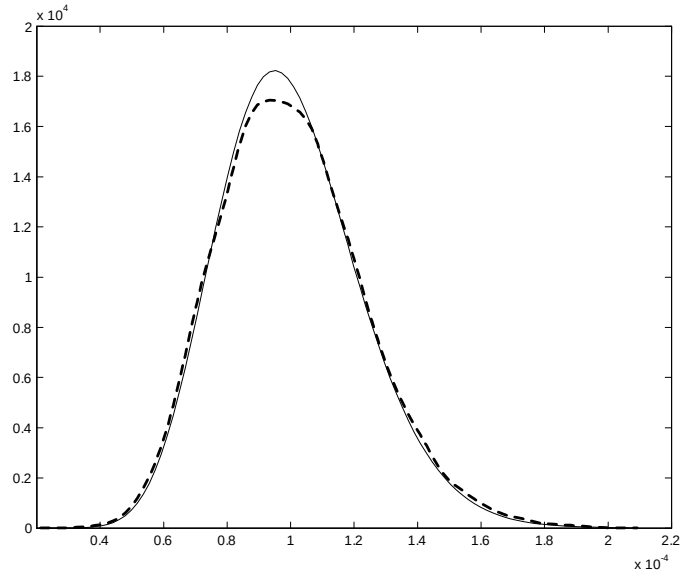

Figure 7.1: Density estimate of the terminal volatilities under the Heston model (dashed line) and the corresponding gamma distribution (solid line).

Figure 7.1 shows that the density estimate of the terminal volatility of the simulated price paths strongly resembles the density of the stationary distribution of the volatility process.

When considering the time changed exponential $N \circ IG$ model we simulate an initial value for the increment in operational time from its stationary distribution, which follows a $gamma\left(\frac{2\kappa\eta}{\lambda^2}, \frac{\lambda^2}{2\kappa}\right)$ random variable. Figure 7.2 shows a density estimate of the terminal value of the increment in operational time obtained from 10000 simulated price paths of length 111 days as a dashed line. The price paths are simulated from a time changed exponential $N \circ IG$ model with $(r, \kappa, \eta, \lambda, \alpha, \beta, \delta) = (0.1/252, 0.1, 1, 0.3, 50, -25, 0.01)$. No special significance is attached to this specific parameter set. The density function of the stationary distribution of the increments is shown in the figure using a solid line.

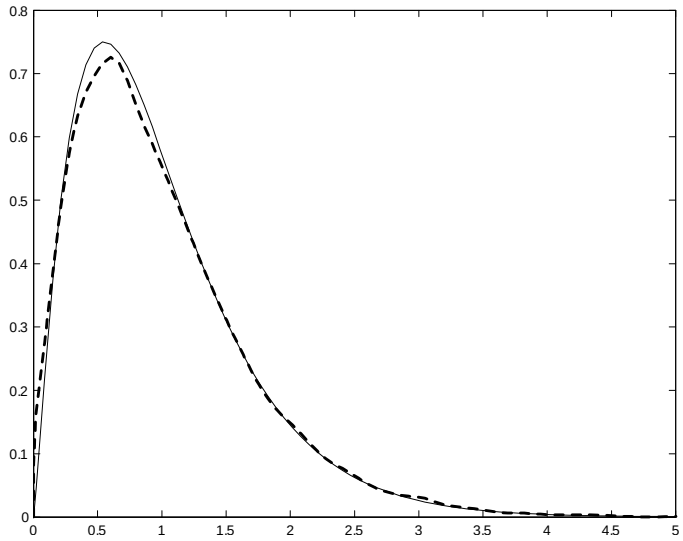

Figure 7.2: Density estimate of the terminal increment in the operational time under the time changed exponential $N \circ IG$ model (dashed line) and the corresponding gamma distribution (solid line).

Figures 7.1 and 7.2 give strong indications that the stationary distributions of the volatility process and the operational time process conform to the gamma laws specified in Section 7.2.

7.3.2 Martingaleness of the discounted stock price

We apply the MCMM to the models specified in Section 7.2 in order to ensure that the discounted stock price forms a martingale under each of the models. In order to show that the discounted stock price does indeed form a martingale under the Heston model we proceed as follows. We simulate 1000 price paths (using the same parameters as those specified in the previous test), each with a length of 111 days. The starting value of the stock price is taken to be 1. We discount the terminal value of each price path and calculate the average. The entire process is repeated 500 times, resulting in 500 average discounted prices. If the requirement of martingaleness is satisfied, we expect that these averages should be close to 1. The same procedure is applied to the remaining models. For the time changed exponential $N \circ IG$ model we use the parameter values used in Section 7.3.1.

For the remaining models we use the parameter sets obtained when calibrating the models to the digital barrier call options later in this chapter; these parameter sets are reported in Appendix H. Table 7.1 shows the mean and the standard deviation of the 500 average terminal stock prices obtained using each model. In the table below and throughout the remainder of the thesis we use TC as an abbreviation for time changed; we also use exp, lin and std dev as abbreviations for exponential, linear and standard deviation respectively.

	Average	Std dev
Heston	1.0001	0.0034
TC exp $N \circ IG$	1.0001	0.0059
Exp $N \circ IG$	1.0001	0.0033
Black-Scholes	0.9999	0.0031
Lin $N \circ IG$	1.0001	0.0034
Lin Brownian motion	0.9999	0.0032
Lin lognormal-normal	1.0000	0.0033
Lin Pareto-normal	0.9999	0.0036
Lin gamma-normal	1.0000	0.0034
Exp lognormal-normal	1.0000	0.0032
Exp Pareto-normal	0.9999	0.0035
Exp gamma-normal	0.9999	0.0034

Table 7.1: Means and standard deviations of average terminal stock prices.

The results in Table 7.1 support our claim that the discounted stock price forms a martingale under each of the models considered.

7.4 The calculation of option prices

In the empirical results shown in this chapter we calculate the prices of the three types of barrier options discussed in Section 7.1 under each of the models discussed in Section 7.2.

Closed form formulae do not exist for the prices of these options. As a result, we estimate the prices of these options using Monte Carlo simulation.

In order to calculate estimates for the option prices under a given model we proceed as follows. We simulate a price path for the stock. The length of this price path is equal to the maximum time to maturity of the options under consideration. We calculate the payoff of each of the options given this price path. The simulation of the price path is repeated 1000 times, resulting in 1000 realised payoffs for each option. The price of a given option is estimated by discounting the average of the payoffs associated with this option.

We opt to use 1000 simulations when calculating option price estimates throughout the remainder of the thesis in spite of the fact that these estimates can be calculated with greater accuracy if the number of simulations is increased. This is done in order to keep the computer time required for calculations within manageable limits.

Below we refer interchangeably to the estimates calculated using this method as option prices and option price estimates.

7.5 Construction of a hypothetical market

Observed barrier option prices are not as readily available as observed European option prices. As a result, we do not use observed option price data in this chapter, instead we opt for prices calculated in a hypothetical market. In this market there exists a risk free bond with a daily interest rate of $0.1/252$ as well as a single stock; the daily price of this stock follows a Heston model with parameter set $(r, \kappa, \eta, \rho, \theta) = (0.1/252, 0.1, 0.0001, -0.7, 0.001)$. The current price of the stock is 1. In this market the three types of barrier options defined in Section 7.1 are available for various barrier levels and times to maturity. The strike prices of all of the down-and-out as well as up-and-out barrier options are fixed at 1; note that the same convention is used in Schoutens *et al.* (2004:66-78). For each option type we consider 25 options with times to maturity ranging from 21 days to 111 days. For the market prices, barrier levels and times to maturity of the options see Appendix F.

In order to calculate the market prices of the options defined in this hypothetical Heston

market we calculate option price estimates using the method described in Section 7.4. These estimates are taken to be the market prices of the options. As a result, we consider the market prices of the options to be known constants, not random variables.

Below we calibrate various models to these market option prices. In order to later be able to assess the fit of the models we consider the variability of the option price estimates under the known, true model underlying the market. We repeat the procedure used to calculate the market option prices to obtain a second set of option price estimates. We then calculate the AAE between the fixed market option prices and the newly calculated option price estimates. This process is repeated 500 times to achieve 500 realisations of the AAE. Table 7.2 shows the average and the standard deviation of the AAEs obtained in this way for each option type considered. In the tables shown in the remainder of the thesis we abbreviate digital barrier by DB, down-and-out barrier by DOB and up-and-out barrier by UOB.

	Average	Std dev
DB	0.0107	0.0035
DOB	0.0017	0.0007
UOB	0.0012	0.0005

Table 7.2: Averages and standard deviations of the AAEs under the true model.

7.6 Model calibration

In this section we consider the calibration of the various models to the market prices of the barrier options available in the hypothetical market defined above. This is done in order to ascertain the accuracy with which it is possible to replicate the market option prices using the models. Below we calibrate each of the option pricing models used to the market prices of the three different option types available in the market separately.

The calibration procedures can be thought of as follows. Different analysts are inter-

ested in the workings of the market constructed in Section 7.5. Each analyst believes that the prices of the options in the market can be modelled accurately by his or her favourite model. The first analyst decides to calibrate the Heston model to the market option prices. The analyst is not aware that the Heston model is the true model underlying the market. He calibrates the model to each of the option types in turn, starting with the digital barrier options.

The analyst aims to find a parameter set that, when used with the Heston model, leads to the smallest possible AAE. He decides to use an optimisation algorithm in which the parameters of the model are varied in an attempt to minimise the associated AAE (we use Matlab's fminsearch.m to obtain the results below). However, the algorithm that he uses requires starting values for the parameters. Since this is the first time that he performs a calibration in this market he does not have a good idea of what parameter values might lead to a small AAE (and hence he is uncertain of which starting values to use). He uses past experience, obtained in other markets, to decide on a number of possible starting values for each parameter. Denote by γ_s the grid of possible starting values considered for a parameter γ. The possible starting values that he considers are as follows:

$$\begin{aligned}
\kappa_s &= (0.01, 0.05, 0.15, 0.2, 0.25)\,, \\
\eta_s &= (0.00001, 0.00005, 0.00015, 0.0002, 0.0003)\,, \\
\rho_s &= (-0.9, -0.8, -0.6, -0.4, 0)\,, \\
\theta_s &= (0.0001, 0.0005, 0.0015, 0.002, 0.003)\,.
\end{aligned}$$

Note that the possible starting values that the analyst considers do not contain the true parameters underlying the market; this is because the analyst does not have any knowledge

of the true parameter values. For each combination of parameters that can be constructed from the possible starting values the analyst calculates the option prices under the Heston model as well as the corresponding AAE. He uses the parameter set leading to the smallest AAE as starting values for the optimisation algorithm. He uses the same procedure to calibrate the parameters of the Heston model to the down-and-out and up-and-out barrier option prices.

Table 7.3 shows the parameters obtained when calibrating the Heston model to the three different sets of option prices. The true parameters underlying the market are included in this table to aid comparisons.

	κ	η	ρ	θ
True model	0.01	0.0001	-0.7	0.001
DB	0.0103	0.0001	-0.6433	0.0015
DOB	0.0097	0.0002	-0.8951	0.0015
UOB	0.0096	0.0002	-0.4138	0.0015

Table 7.3: Calibrated parameter sets for the Heston model.

Note that, while the values obtained for κ in the calibrated parameter sets are fairly close to the true value, this is not the case for the other parameters. However, for our purposes the value of the individual parameters are not of much importance; we are mainly interested in the option prices that are obtained when using the various parameter sets. As a result, we do not analyse the parameter values obtained by the calibration procedures. Instead we focus on the AAE associated with each calibration.

Various analysts calibrate the remaining models using the technique described above. The grids of possible starting values used are reported in Appendix G, while the parameter sets obtained by calibration are reported in Appendix H.

We are mainly interested in the AAE values achieved by the calibration of the various models. Recall that since the option price estimates are calculated using Monte Carlo

simulation the price estimates are random variables. Therefore, the AAE associated with a given calibration is itself is a random variable. In order to assess the fit of the calibrated Heston model to the market prices of the digital barrier options for example we proceed as follows. We estimate the prices of the digital barrier options under the calibrated Heston model using Monte Carlo simulation and we calculate the AAE between these price estimates and the market prices. This process is repeated 500 times, resulting in 500 realisations of the AAE. The same procedure is used for the remaining option types and models. Table 7.4 shows the average and the standard deviation (in brackets) of the resulting AAEs. We include the AAEs obtained using the true model in order to aid comparison.

	DB	DOB	UOB
True model	0.0107 (0.0035)	0.0017 (0.0007)	0.0012 (0.0005)
Heston	0.015 (0.004)	0.0026 (0.0012)	0.0015 (0.0004)
TC $N \circ IG$	0.0154 (0.0039)	0.0019 (0.0008)	0.0013 (0.0005)
Exp $N \circ IG$	0.0139 (0.0033)	0.0018 (0.0005)	0.0012 (0.0005)
Black-Scholes	0.0131 (0.0043)	0.0017 (0.0006)	0.0014 (0.0005)
Lin $N \circ IG$	0.0108 (0.0029)	0.0017 (0.0005)	0.0013 (0.0005)
Lin Brownian motion	0.0101 (0.003)	0.0027 (0.0008)	0.0014 (0.0005)
Lin lognormal-normal	0.0099 (0.0031)	0.0021 (0.0009)	0.0015 (0.0005)
Lin Pareto-normal	0.0108 (0.0029)	0.0019 (0.0009)	0.0013 (0.0005)
Lin gamma-normal	0.0103 (0.003)	0.002 (0.001)	0.0016 (0.0005)
Exp lognormal-normal	0.0117 (0.0032)	0.0022 (0.0009)	0.0013 (0.0005)
Exp Pareto-normal	0.0155 (0.0034)	0.002 (0.001)	0.0014 (0.0006)
Exp gamma-normal	0.018 (0.0047)	0.0019 (0.0008)	0.0015 (0.0005)

Table 7.4: Average and standard deviation of the AAEs under the calibrated models.

An analysis and interpretation of the results in Table 7.4 is given in the next section.

7.7 Analysis of calibration results

The aim of this chapter is to ascertain whether or not it is possible replicate barrier option prices calculated under one model using another model (the parameters of second model are obtained using the calibration procedure discussed in Section 7.6). In order to determine if one model is able to replicate the option prices calculated under another we introduce the concept of a perfect calibration and we classify each of the calibrations as either perfect or imperfect. We test the following hypotheses for each individual calibration:

$$\begin{aligned} H_0 &: \text{The calibration is perfect.} \\ H_A &: \text{The calibration is not perfect.} \end{aligned} \tag{7.3}$$

In order to explore this line of thought further it is necessary to define what we mean when we call a calibration perfect. Since the AAE associated with a given calibration is a random variable we cannot simply insist on the AAE being equal to 0 in order for the calibration to be called perfect. Instead we shall call a calibration perfect if the population mean of the AAE associated with the calibration does not exceed the population mean of the AAE under the true model.

Denote by AAE_C the AAE associated with a given calibration and denote by AAE_T the AAE under the true model. Table 7.4 reports the means and standard deviations of 500 realisations of the AAE_C for each of the calibrations as well as 500 realisations of the AAE_T. Denote by μ_C and μ_T the mean of the AAE_C and the AAE_T respectively. Restating the hypotheses in (7.3) we obtain

$$\begin{aligned} H_0 &: \mu_C = \mu_T. \\ H_A &: \mu_C > \mu_T. \end{aligned} \tag{7.4}$$

Note that we do not consider the case where $\mu_C < \mu_T$ since this would imply that the calibrated model fits the observed option prices better than the true model.

Perhaps the first idea that comes to mind when contemplating the question if the mean of the AAE_C exceeds that of the AAE_T is a formal, frequentist hypothesis test. However, because of the large sample size used, this test will reject the null hypothesis for very small differences between the sample means. As a result, we discard this approach and we rather consider two other formal tests below. The first utilises a mixed Bayesian-frequentist approach and the second uses a bootstrap methodology.

We use the calibration of the exponential lognormal-normal model to the up-and-out barrier option prices as an example to illustrate the calculations involved in both tests. Figure 7.3 shows a density estimate of the AAE_C associated with this calibration (solid line) as well as a density estimate of the AAE_T (dashed line).

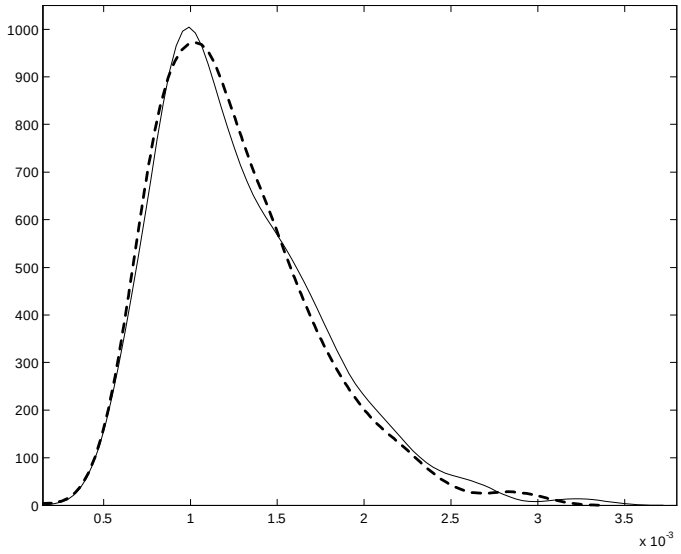

Figure 7.3: Density estimates of the true AAE for the up-and-out barrier options (dashed line) and the AAE of the exponential lognormal-normal model calibration (solid line).

7.7.1 Mixed Bayesian-frequentist hypothesis test

The discussion to follow is based on an alternative to frequentist significance tests explained in Cox and Hinkley (1974:394-396). Below we outline this method in general terms before

describing the specifics of the test used to answer questions relating to perfect calibrations.

Using Bayesian techniques to draw inference about the value of some parameter θ we assume that the parameter is a realisation of a random variable Θ with a known probability distribution. Consider the case where we have a sample $\mathbf{y}$ from a known distribution with unknown parameter θ and we want to test a null hypothesis (H_0) specifying $\theta = \theta_0$. In order to specify the distribution of Θ completely we assign a probability of p_0 to $\theta = \theta_0$ (p_0 is the a priori probability of the null hypothesis being true) and we assign a prior distribution $p_A(\theta)$ to Θ under the alternative hypothesis (H_A). This means that the prior distribution of Θ consists of an atom of probability p_0 at $\theta = \theta_0$ and a probability density function of $(1 - p_0) p_A(\theta)$ over the remaining values of θ.

In order to determine whether or not $\theta = \theta_0$ we use the posterior odds given $\mathbf{y}$; this is defined as the probability of H_0 given $\mathbf{y}$ divided by the probability of H_A given $\mathbf{y}$. The posterior odds is calculated as follows:

$$\frac{P(H_0|Y = \mathbf{y})}{P(H_A|Y = \mathbf{y})} = \frac{p_0 f_{Y|\Theta}(\mathbf{y}|\theta_0)}{(1 - p_0) \int_{\theta \neq \theta_0} p_A(\theta) f_{Y|\Theta}(\mathbf{y}|\theta) \, d\theta}, \tag{7.5}$$

where $f_{Y|\Theta}$ denotes the likelihood function associated with the data for a given value of θ. If the data support H_0 the value of the posterior odds will be large. Below we reject H_0 for posterior odds smaller than 1.

In order to apply the technique outlined above we proceed in a way similar to Example 10.12 in Cox and Hinkley (1974:395-396). However, our approach deviates from the methodology described in this example in that we opt for a mixed Bayesian-frequentist approach; this is explained below.

As an example we consider the AAEs resulting from the calibration of the exponential lognormal-normal model to the up-and-out barrier options. The mean and standard deviation of the AAEs associated with this calibration are 0.0013 and 0.0005 respectively. This seems fairly close to the mean and standard deviation of the AAEs associated with the true model which are 0.0012 and 0.0005 respectively.

The Bayesian technique discussed above is based on a single population. Therefore, it cannot be applied directly to test the hypotheses in (7.4). In order to test these hypotheses we define a variable $X = AAE_C - AAE_T$ and denote the mean of X by μ_X. In order to sample from X we subtract the realised values of AAE_T from the realised values of AAE_C, resulting in a sample of size 500. Now we can reformulate the hypotheses in (7.4) once more into a more suitable form;

$$H_0 \quad : \quad \mu_X = 0.$$

$$H_A \quad : \quad \mu_X > 0. \tag{7.6}$$

The hypotheses specified in (7.6) are in a form that can be tested using the method described in Cox and Hinkley (1974:395-396). In order to apply this method we have to specify the distribution of X; we consider the normal distribution as a candidate for the distribution of X. The Lilliefors test does not reject the assumption of normality at a significance level of 1% (the p-value is calculated to be 0.0618). Therefore we assume that X follows a normal distribution. Figure 7.4 shows a density estimate of X.

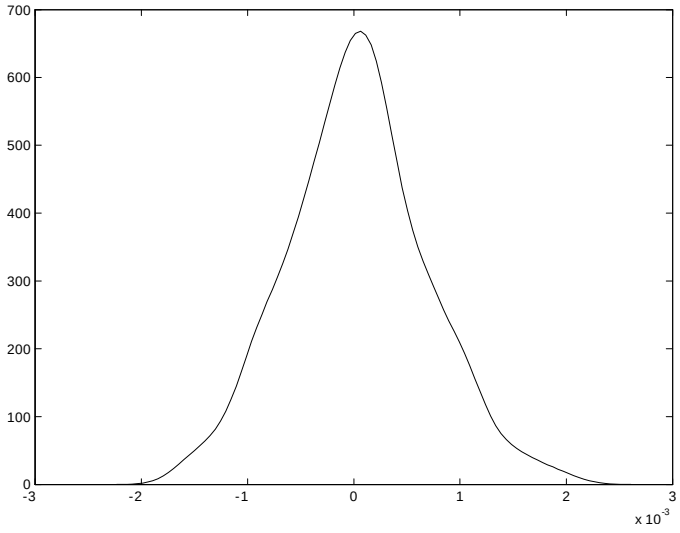

Figure 7.4: Density estimate of X.

Under the null hypothesis $X \sim N\left(0, \sigma_0^2\right)$ where σ_0^2 is a known constant. We opt for a mixed Bayesian-frequentist approach in that we estimate σ_0^2 from sample data. We set σ_0^2 equal to the variance of the realisations of X; this is calculated to be 4.2568×10^{-7}. Under the alternative hypothesis $X \sim N\left(\mu, \sigma_0^2\right)$ where μ is itself a random variable following some known (a priori) distribution. For the sake of simplicity we assume that $\mu \sim N\left(0, \nu\right)$, with ν some known constant. However, we are only interested in the one-sided alternative that $\mu > 0$. As a result, the density function of μ under the alternative hypothesis is

$$f_{H_A}\left(\mu\right) = \sqrt{\frac{2}{\pi\nu}} \exp\left(-\frac{\mu^2}{2\nu}\right),$$

for $\mu > 0$.

It remains to specify the value of ν. Again we opt for a mixed approach in that we estimate the value of ν from the data as follows. Calculate the average of X, this can be viewed as a realisation of μ. Still considering only up-and-out barrier options, define Y as the difference between the AAEs obtained using the Heston model and the AAEs obtained using the true model. In the same way as before we calculate realisations of Y as well as the average of Y. We can view this as a second realisation of μ. Continuing in this way we use the results of each of the remaining models to calculate a realisation of μ; we obtain twelve realisations of μ using the calibrations relating to the up-and-out-barrier options. The value of ν is set equal to the variance of this sample, which is calculated to be 9.8781×10^{-9}.

The last quantity required for the test that remains unspecified is the value of p_0, the prior probability of H_0 being true. Since we have no prior knowledge and little in the way of intuition regarding whether or not the null hypothesis holds, we arbitrarily set $p_0 = 0.5$.

If we denote by $\mathbf{x} = (x_1, x_2, ..., x_{500})$ the realised values of X, then the posterior odds

defined in (7.5) is calculated to be

$$\frac{P\left(H_0|X=\mathbf{x}\right)}{P\left(H_A|X=\mathbf{x}\right)}=\frac{p_0\sqrt{\pi\nu}\exp\left(-\frac{\sum x_j^2}{2\sigma_0^2}\right)}{\sqrt{2}\left(1-p_0\right)\int_0^\infty\exp\left(-\frac{\sum\left(x_j-\mu\right)^2}{2\sigma_0^2}-\frac{\mu^2}{2\nu}\right)d\mu}. \tag{7.7}$$

Using (7.7) we calculate the posterior odds to be 1.1345 in the example considered. This means that the test does not reject the hypothesis that $\mu_X = 0$ in (7.6). By extension this means that the test does not reject the hypothesis that the calibration of the exponential lognormal-normal model to the up-and-out barrier option prices is perfect.

The same procedure is used to calculate the posterior odds for each of the remaining calibrations. The results are shown in Table 7.5 below. In the table we indicate the calibrations for which the assumption of normality is rejected by the Lilliefors test at a 1% significance level using an asterisk.

	DB	DOB	UOB
Heston	0*	0	0
TC $N \circ IG$	0	0.0019*	0.0137
Exp $N \circ IG$	0	6.1606*	3.5960
Black-Scholes	0	17.0209*	0.0116
Lin $N \circ IG$	5.4726*	14.3854*	0.0764
Lin Brownian motion	51.5668*	0	0.0001
Lin lognormal-normal	57.8602*	0*	0*
Lin Pareto-normal	6.1884*	0.0595*	0.015*
Lin gamma-normal	36.8054*	0.0001*	0
Exp lognormal-normal	0.0001	0*	1.1345
Exp Pareto-normal	0	0.0005*	0
Exp gamma-normal	0	0.0205*	0

Table 7.5: Posterior odds associated with the various calibrations.

The table indicates that the posterior odds associated with many of the calibrations are greater than 1. This seemingly implies that these calibrations can be classified as perfect. However, in the majority of the calibrations resulting in posterior odds greater than 1 the Lilliefors test rejects the assumption of normality at a 1% significance level. This casts some doubt as to the accuracy of the test in these cases. In order to clarify the results we consider a non-parametric test below.

7.7.2 Bootstrap based hypothesis test

Allison (2008:26-29) shows how to test hypotheses of the kind shown in (7.4) using the bootstrap; we use this methodology below. Continuing with the same example as shown above we consider the calibration of the exponential lognormal-normal model to the up-and-out barrier option prices. For ease of notation let X represent the AAE associated with this calibration and let Y denote the AAE associated with the true model. Also let $x_1, ..., x_{500}$ and $y_1, ..., y_{500}$ represent the samples taken from X and Y respectively. Let $\overline{x}$ and $\overline{y}$ be the means of the two samples and let $s_x^2 = \frac{1}{500}\sum_{j=1}^{500}(x_j - \overline{x})^2$ and $s_y^2 = \frac{1}{500}\sum_{j=1}^{500}(y_j - \overline{y})^2$. Denote by μ_C the mean of X and by μ_T the mean of Y. Below we test the hypotheses in (7.4) using a 5% significance level.

The test statistic used is

$$T = \frac{\overline{x} - \overline{y}}{\sqrt{\frac{s_x^2}{500} + \frac{s_y^2}{500}}}.$$

In the example under consideration the realised value of T is calculated to be 1.1258. We use the bootstrap to estimate the distribution of T under H_0 as follows. Under the null hypothesis $\mu_C = \mu_T$. In order to mimic this equality in the bootstrap world we apply the following transformation to the sample data; $r_j = x_j - \overline{x}$ and $s_j = y_j - \overline{y}$ for $j = 1, ..., 500$. We then draw a simple random sample of size 500 with replacement from $r_1, ..., r_{500}$; denote the elements of this sample by $x_1^\star, ..., x_{500}^\star$. We also draw a simple random sample of size 500 with replacement from $s_1, ..., s_{500}$; denote the elements of this sample by $y_1^\star, ..., y_{500}^\star$.

Using notation similar to those defined above we calculate

$$T^\star = \frac{\overline{x^\star} - \overline{y^\star}}{\sqrt{\frac{s^2_{x^\star}}{500} + \frac{s^2_{y^\star}}{500}}}.$$

$T^\star$ is a realisation of the test statistic under the null hypothesis. We repeat this process 10000 times to obtain 10000 realisations of $T^\star$. Figure 7.5 shows a density estimate of the distribution of $T^\star$.

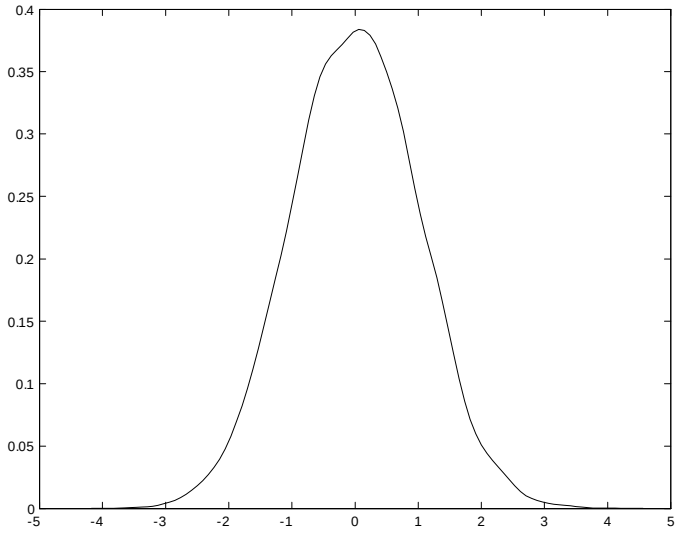

Figure 7.5: Density estimate of $T^\star$.

The p-value associated with the test is equal to the proportion of $T^\star$s that are larger than or equal to the realised value of T. For the calibration in question the p-value is calculated to be 0.1343. As a result, we do not reject H_0 at a 5% level of significance. For this reason we consider the calibration of the exponential lognormal-normal model to the up-and-out barrier option prices to be perfect.

The same procedure is applied to test the hypotheses that the remaining calibrations are perfect. Table 7.6 shows the resulting p-values associated with each of the calibrations. In the table we use an asterisk to indicate that the hypothesis of a perfect calibration is rejected.

	DB	DOB	UOB
Heston	0*	0*	0*
TC $N \circ IG$	0*	0*	0.0006*
Exp $N \circ IG$	0*	0.3913	0.4992
Black-Scholes	0*	0.9098	0.0001*
Lin $N \circ IG$	0.1874	0.8228	0.0041*
Lin Brownian motion	0.9983	0*	0*
Lin lognormal-normal	0.9998	0*	0*
Lin Pareto-normal	0.2141	0.0022*	0.0002*
Lin gamma-normal	0.9618	0*	0*
Exp lognormal-normal	0*	0*	0.1343
Exp Pareto-normal	0*	0*	0*
Exp gamma-normal	0*	0.0007*	0*

Table 7.6: Bootstrap p-values associated with the various calibrations.

When comparing the results shown in Table 7.5 to those reported in Table 7.6 we see that the results are identical; all calibrations classified as perfect by the mixed Bayesian-frequentist hypothesis test are also classified as perfect by the bootstrap based hypothesis test and vice versa.

7.7.3 A comment on the results of the imperfect calibrations

Both tests classify the same ten calibrations as perfect. One might easily imagine that the remaining calibrations do not fit the observed option prices well. This is not the case. As an example, consider the linear Pareto-normal model calibration to the down-and-out barrier option prices. Both of the formal tests decisively reject the hypothesis that this calibration is perfect. Figure 7.6 shows a density estimate of the AAE associated with this calibration (solid line) as well as a density estimate of the true AAE associated with the

down-and-out barrier option prices (dashed line).

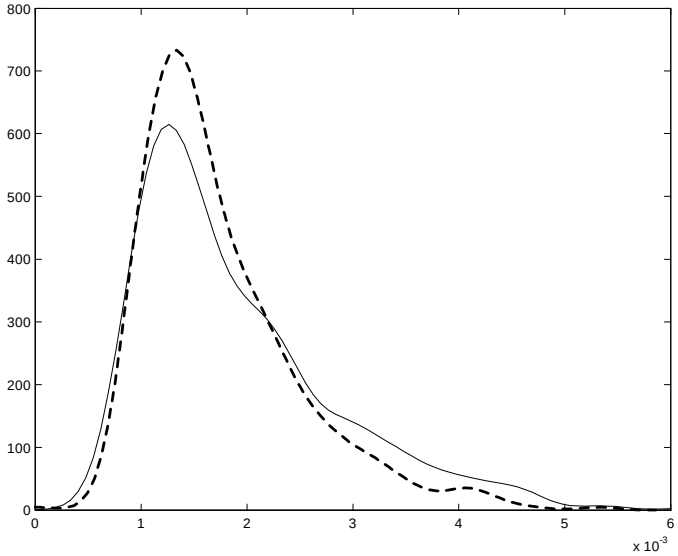

Figure 7.6: Density estimates of the true AAE for the down-and-out barrier options (dashed line) and the AAE of the linear Pareto-normal model calibration (solid line).

Figure 7.6 is included to demonstrate that a calibration not classified as perfect can have a small associated AAE. Therefore, a calibration that is not classified as perfect can still fit the market option prices well.

7.8 Conclusions

In this chapter we consider the interchangeability of barrier option pricing models. This entails exploring the possibility of accurately replicating the prices calculated under one model using another model (which is calibrated to the option prices under consideration). We consider three different types of barrier options; digital barrier call options, down-and-out barrier call options and up-and-out barrier call options. Since the prices of the barrier options are not as readily available as their European counterparts we construct a hypothetical market. In this market the stock price follows a Heston model and we calculate the market prices of the options using this model.

We use the average absolute error (AAE) as a distance measure between the market

option prices and the prices calculated using the models considered. The calibration procedure used attempts to minimise the AAE using martingale restricted calibration and the mean correcting martingale measure (MCMM). Closed form expressions for the option prices considered do not exist, so we estimate the prices of the options using a Monte Carlo procedure. Since the option price estimates are random variables, the AAE associated with a given model (even the true model underlying the market) is also a random variable. In order to aid our comparison between the models considered we introduce the concept of a perfect calibration. We define a calibration to be perfect if the mean of the AAE associated with the calibration does not exceed the mean of the AAE under the true model.

The models are calibrated in turn to each option type. We employ two formal hypothesis testing procedures to decide whether or not the calibrations are perfect. Both tests classify the same ten calibrations as perfect. Some of the calibrations are not classified as perfect by the two formal hypothesis tests, but still exhibit small AAE values and fit the observed option prices well.

Since various models provide perfect calibrations to option prices calculated using the Heston model we conclude that it is possible to mimic the behaviour of one of the models by using another. Therefore, we conclude that the models are interchangeable.

It is interesting to observe that although the true model underlying the market is a Heston model, the Heston model is one of the few models for which none of the calibrations are deemed perfect. Surprisingly, it seems that prior knowledge that the market prices are calculated using the Heston model does not result in more accurate calibrations. The author realises that the preceding statement may cause doubt as to the accuracy of the algorithms used in order to obtain the results. A confirmatory analysis of these algorithms is included in Section 7.3 in order to address this issue. A possible explanation of the absence of perfect calibrations associated with the Heston model lies therein that the optimisation algorithm used converges to a local minimum and not a global minimum. However, the AAEs associated with the calibrated results are small. This indicates that the calibrations relating to the Heston model might not be perfect, but the option prices calculated using

the calibrated models fit the corresponding market prices well.

The optimisation problems considered are generally non-convex. As a result, the only way to ensure that the global minimum is achieved is by calculating the AAE for each possible parameter set. Since this is not possible in practice we are generally unable to attain global minima.

The means of the AAEs associated with the majority of the calibrations are small. We conclude that it is indeed possible to accurately (even perfectly) calibrate various models to the prices of barrier options. As a result, we conclude that barrier option pricing models are interchangeable.

Chapter 8

Calibrations and estimations relating to barrier options

The main aim of this thesis is to investigate the effects of calibrating Lévy option pricing models to market option prices. In this chapter we endeavour to uncover some of these effects, especially relating to the AAE associated with a calibration. The discussion below is limited to barrier options.

As was the case in the previous chapter we construct a hypothetical market containing three types of barrier options; digital barrier call options, down-and-out barrier call options and up-and-out barrier call options. All options are considered to be monitored daily. In this chapter we return to the exponential Lévy models considered earlier in the thesis; the exponential $N \circ IG$ model and the Black-Scholes model.

We calibrate each of these models to the prices of the various option types, using the techniques described in Chapter 7, in turn. An alternative to calibrating the parameters of the models is estimating these parameters by fitting the relevant distribution to the log-returns of the stock price. Below we refer to this technique as model estimation (as opposed to model calibration). We are interested in determining whether or not it is possible to substantially reduce the AAE associated with a given model estimation by using a calibration procedure.

In the previous chapter we defined a calibration to be perfect if the mean of the AAE associated with the calibration does not exceed the mean of the AAE under the true model. Here we use the same criterion to classify a calibration as either perfect or not.

Similarly, we define an estimation to be perfect if the mean of the AAE associated with the estimation does not exceed the mean of the AAE under the true model. Below we identify perfect calibrations and estimations in order to aid our comparison of the results associated with the two methods. As before, this is done in order to demonstrate that barrier option pricing models are interchangeable.

In Section 8.1 we construct a hypothetical exponential $N \circ IG$ market containing the three types of barrier options used. We discuss the calibration and estimation procedures used and their results in Section 8.2. Similar analyses as those performed in the previous chapter are included in Section 8.3 in order to ascertain whether or not the resulting model calibrations and estimations are perfect. Finally Section 8.4 concludes.

8.1 Construction of a hypothetical market

As is the case in the previous chapter, we do not use option price data observed in some financial market. Instead we calculate option prices in a hypothetical market. In this market there exists a risk free bond with a daily interest rate of $0.1/252$ as well as a single stock with a current price of 1. The price of the stock follows an exponential $N \circ IG$ model. In order to obtain a realistic parameter set we fit a $N \circ IG$ distribution to the log-returns of the Standard and Poor 500 index for the year ending on 18 April 2002 (this specific dataset is chosen as a matter of convenience). The resulting parameters are $(\alpha, \beta, \mu, \delta) = (74.8181, 9.8668, -0.0040, 0.0297)$.

The three types of barrier options considered are available for various barrier levels and times to maturity (the barrier levels, times to maturity and strike prices used are the same as in the previous chapter). As before we use a strike price of 1 for all down-and-out barrier options as well as all up-and-out barrier options. We consider 25 options of each type. Appendix I shows the market prices, barrier levels and times to maturity of the options used.

In order to calculate the market prices of the options defined above we calculate option price estimates using the method described in Section 7.4. As before, we consider these

estimates to be the fixed market prices of the options.

Below we consider the variability of the AAE under the true model underlying the market. This is done in order to establish a benchmark that can later be used to classify a given calibration or estimation as being perfect or imperfect. We use the sampling procedure described in Section 7.5 to obtain realisations of the AAE based on 1000 Monte Carlo simulations. Table 8.1 shows the averages and standard deviations of the realised AAEs for the various option types.

	Average	Std dev
DB	0.0161	0.0047
DOB	0.0026	0.0011
UOB	0.0011	0.0002

Table 8.1: Averages and standard deviations of the AAEs under the true model.

8.2 Model calibration and estimation

The calibration of the two models considered to the market prices of the options is done using the procedure described in Section 7.6. As before, we calibrate each model separately to each option type, resulting in a total of six calibrations. The possible starting values considered for the exponential $N \circ IG$ model are the same as those considered in the previous chapter. The starting values considered for σ in the Black-Scholes model are $(0.015, 0.0175, 0.02, 0.0225, 0.025)$.

We now turn our attention to the estimation procedure used to fit the relevant distributions to the log-returns of the stock. All estimations are based on a dataset consisting of 252 log-returns (corresponding to one year's stock price history). Consider first the estimation procedure relating to the exponential $N \circ IG$ model. Parameter estimation for the $N \circ IG$ distribution is discussed in Section 3.3. Below we calculate parameter estimates using the IMLE, see Section 3.3.1. However, this estimator is tailored to this specific ap-

plication as follows. In the estimations below we use an optimisation algorithm that varies the parameters of the distribution in an attempt to maximise the likelihood function; we use Matlab's fminsearch.m for this purpose. As is the case for the calibration procedure used, we do not allow the location parameter μ to vary freely. This parameter is chosen so as to ensure the martingaleness of the discounted stock price using the MCMM. In effect the algorithm maximises the likelihood function subject to the constraint that the price process forms a martingale. A similar estimation method is used in order to obtain the estimated Black-Scholes model; the value of σ is varied in an attempt to maximise the likelihood function while μ is calculated using the MCMM.

The reason for imposing the martingaleness constraint on the estimation algorithm is twofold. Firstly, it makes direct comparisons between the results of the calibrations and that of the estimations possible since the same restriction in placed on both methods. Secondly, since the stock price is guaranteed to form a martingale under the estimated model we do not need to perform a measure change. Various measure changes and their effects are discussed in detail in Chapter 4. However, in this chapter we are not concerned with the effects of measure changes, we want to isolate the effects of the calibration and estimation procedures. As a result, we opt for an estimation technique that does not require any adjustment of the estimated parameter set using a change of measure.

As is the case for the calibration algorithm, we require starting values for the optimisation algorithm. In the case of the Brownian motion model we simply use the standard deviation of the log-returns as the starting value for the estimate of σ. For the exponential $N \circ IG$ model we specify the same possible starting values that are used for the calibration procedure. For each combination of these values we calculate the likelihood function. The parameter set resulting in the largest value of the likelihood function is used as starting values.

Below we consider the results of the calibration and estimation procedures. Table 8.2 shows the parameter sets obtained for the calibrations and estimations relating to the exponential $N \circ IG$ model. Since the model is calibrated to each set of option prices

separately we report three different parameter sets for the calibrations compared to the single parameter set reported for the estimation. The true parameters underlying the market are included in Table 8.2 to aid comparisons.

	α	β	μ	δ
True model	74.8181	9.8668	−0.0040	0.0297
DB calibration	300.6594	−91.7852	0.0332	0.1039
DOB calibration	150.9598	−4.9909	0.0017	0.0516
UOB calibration	149.0349	−46.6797	0.0170	0.0519
Estimation	81.1759	0.7942	−0.0004	0.0337

Table 8.2: Parameter sets for the $N \circ IG$ model.

As was the case in the previous chapter, we do not concentrate on the values of the parameters themselves but rather on the AAEs associated with each of the parameter sets.

Table 8.3 shows the parameters obtained by means of calibration and estimation for the Black-Scholes model.

	μ	σ
DB calibration	−0.0001	0.0195
DOB calibration	−0.0002	0.0253
UOB calibration	−0.0001	0.0199
Estimation	−0.0001	0.0203

Table 8.3: Parameter sets for the Black-Scholes model.

Next we consider the AAEs achieved by the various calibrations and estimations. As before, we obtain 500 realisations of the AAE for each of the calibrations and estimations. Table 8.4 shows the averages and standard deviations (in brackets) of the realised

AAEs. The average and standard deviation of the AAEs obtained using the true model are included in order to aid comparison.

	DB	DOB	UOB
True model	0.016 (0.0048)	0.0027 (0.0012)	0.0012 (0.0003)
Exp $N \circ IG$ model	0.015 (0.0044)	0.0028 (0.0012)	0.0011 (0.0002)
Black-Scholes model	0.0161 (0.0054)	0.0048 (0.0022)	0.0011 (0.0002)

Table 8.4: Averages and standard deviations
of the AAEs under the calibrated models.

Some of the results shown in Table 8.4 hint at the presence of perfect calibrations. Consider for example the average of the AAE associated with the calibration of the $N \circ IG$ model to the digital barrier options. The sample mean of the AAEs associated with this calibration is lower than that of the AAEs under the true model.

Table 8.5 shows the averages and standard deviations of the realised AAEs associated with the various estimations.

	DB	DOB	UOB
True model	0.016 (0.0048)	0.0027 (0.0012)	0.0012 (0.0003)
Exp $N \circ IG$ model	0.0161 (0.0050)	0.0027 (0.0011)	0.0011 (0.0002)
Black-Scholes model	0.0158 (0.0053)	0.0027 (0.0012)	0.0012 (0.0003)

Table 8.5: Averages and standard deviations
of the AAEs under the estimated models.

As was the case for the calibration results, the table contains small average AAE values. The next section provides a detailed analysis of the results shown in Tables 8.4 and 8.5.

8.3 Analysis of calibration and estimation results

In order to aid the comparison of the results we classify each calibration and estimation as perfect or imperfect. Recall that a calibration or estimation is classified as perfect if

the mean of the associated AAE does not exceed that of the AAE under the true model. Denote by γ the mean of the AAE associated with a given calibration or estimation and let μ_T denote the mean of the AAE under the true model. The hypotheses tested to classify a calibration or estimation as perfect or imperfect are

$$H_0 \quad : \quad \gamma = \mu_T.$$

$$H_A \quad : \quad \gamma > \mu_T. \tag{8.1}$$

If H_0 is not rejected for a given calibration or estimation it is deemed perfect.

Below we use a mixed Bayesian-frequentist hypothesis test as well as a bootstrap based hypothesis test to classify each of the calibrations and estimations as perfect or imperfect. Both of these methods are discussed briefly below; for a more detailed explanation the reader is referred to Section 7.7.

After the formal tests we include an informal graphical test comparing the density estimates of the AAE calculated under the calibrations and estimations to that of the AAE under the true model. We include this graphical test only for the results pertaining to the digital barrier options.

8.3.1 Mixed Bayesian-frequentist hypothesis test

Consider the AAE associated with some calibration or estimation. Define X to be this AAE minus the AAE under the true model. Recall that we have calculated 500 realisations of the AAE associated with each calibration and estimation as well as the true model. As before we obtain a sample from X using these realisations. Let μ_X denote the mean of X.

We can reformulate the hypotheses in (8.1) as follows:

$$H_0 \quad : \quad \mu_X = 0.$$

$$H_A \quad : \quad \mu_X > 0. \tag{8.2}$$

Below we use mixed Bayesian-frequentist techniques to calculate the posterior odds for the hypotheses in (8.2). As before, we reject H_0 if the posterior odds is calculated to be less than 1.

In order to apply this test we assume that X follows a normal distribution with known variance σ_0^2. We test the assumption of normality using the Lilliefors test. A significance level of 1% is used for this test throughout. We use a mixed Bayesian-frequentist approach in that we estimate σ_0^2 by the sample variance of X. Under the null hypothesis $X \sim N\left(0, \sigma_0^2\right)$. Under the alternative hypothesis $X \sim N\left(\mu, \sigma_0^2\right)$ where $\mu \sim N\left(0, \nu\right)$. However, we consider only the case where $\mu > 0$, therefore the density function of μ under the alternative hypothesis is

$$f_{H_A}\left(\mu\right) = \sqrt{\frac{2}{\pi\nu}} \exp\left(-\frac{\mu^2}{2\nu}\right), \tag{8.3}$$

for $\mu > 0$.

We use a mixed Bayesian-frequentist approach in that we estimate the value of ν in (8.3) as follows. Consider, for example, the case where X is obtained from the calibration of the exponential $N \circ IG$ model to the digital barrier option prices. We calculate the sample mean of X; this is regarded as a realisation of μ. Let X_2 be the difference between the AAE associated with the Black-Scholes model calibration and the AAE under the true model. The sample mean of X_2 is also regarded as a realisation of μ. Continuing in the same fashion, we calculate X_3 and X_4 based on the sample means of the AAEs associated with the two model estimations. We estimate ν by the variance of the sample means of X, X_2, X_3 and X_4. The same approach is used to obtain suitable estimates for ν for the remaining option types.

As was the case in Section 7.7.1, we arbitrarily set the prior probability of H_0 being true to 0.5.

Table 8.6 shows the posterior odds calculated for each of the calibrations and estimations. The presence of an asterisk in Table 8.6 indicates that the Lilliefors test rejects the assumption of normality at a 1% significance level.

	DB	DOB	UOB
Exp $N \circ IG$ calibration	8.583	4.0137	1.9692
Black-Scholes calibration	1.5158	0*	3.2247
Exp $N \circ IG$ estimation	2.0115	16.6147*	2.6307
Black-Scholes estimation	2.7241*	15.7908*	1.3863

Table 8.6: Posterior odds associated with the calibrations and estimations.

From Table 8.6 we conclude that five of the six models are calibrated perfectly to the option prices considered. The table also indicates that all of the estimations are perfect. However, for three of the estimations, the hypothesis of normality is rejected. This leaves some doubt as to the reliability of the test used.

8.3.2 Bootstrap based hypothesis test

Below we use the bootstrap to test the hypotheses in (8.1) in order to classify calibrations and estimations as either perfect or imperfect. We use the same methodology to test the hypotheses here as was used in Section 7.7.2. As before we use a significance level of 5%.

Table 8.7 shows the p-values calculated for the hypotheses in (8.1) for each of the calibrations and estimations. In the table we use an asterisk to indicate that the hypothesis of a perfect calibration is rejected.

	DB	DOB	UOB
Exp $N \circ IG$ calibration	0.9997	0.1206	0.7708
Black-Scholes calibration	0.3895	0^*	0.9846
Exp $N \circ IG$ estimation	0.491	0.5248	0.9123
Black-Scholes estimation	0.7436	0.5811	0.5364

Table 8.7: Bootstrap p-values associated
with the calibrations and estimations.

The results in Table 8.7 support the results in Table 8.6. We conclude that the majority of the calibrations are perfect. The only calibration that is not classified as perfect is the calibration of the $N \circ IG$ model to the down-and-out barrier options. The bootstrap based test concurs with the mixed Bayesian-frequentist test in that it classifies all of the estimations as perfect.

8.3.3 Graphical tests

Below we compare density estimates of the AAEs associated with the calibrations and estimations. For the sake of brevity we show only the results relating to the digital barrier options. In each case we include a density estimate of the AAE under the true model in order to aid comparisons.

Denote by AAE_T, AAE_C and AAE_E the AAEs associated with the true model, the calibrated model and the estimated model respectively. Figure 8.1 shows density estimates of the AAE_C (using a circled line) and that of the AAE_E (using a dashed line) associated with the exponential $N \circ IG$ model. A density estimate of the AAE under the true model is shown using a solid line.

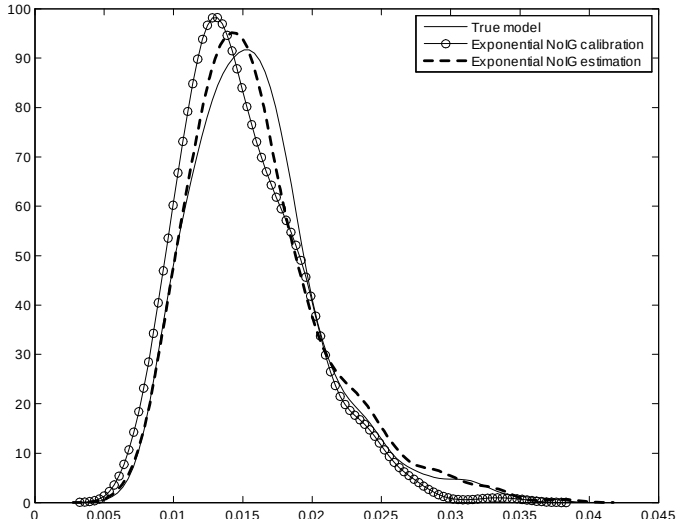

Figure 8.1: Density estimates of the AAEs associated with the exponential $N \circ IG$ model and the digital barrier option prices.

The density estimates of AAE_C and AAE_E are similar to that of the AAE_T. In fact, Figure 8.1 seems to indicate that the mean of AAE_C is smaller than the mean of the AAE_T.

Figure 8.2 shows density estimates of the various AAEs associated the Black-Scholes model. As before the density estimate of the AAE_C is shown using a circled line, that of the AAE_E using a dashed line and that of the AAE_T using a solid line.

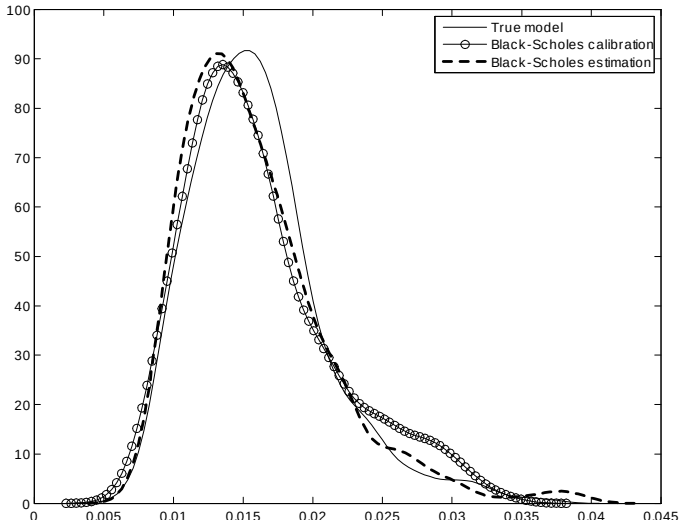

Figure 8.2: Density estimates of the AAEs associated with the Black-Scholes model and the digital barrier option prices.

The density functions in Figure 8.2 suggest that there is little difference between the means of the AAEs associated with the calibration and the estimation. Both the AAE_C and the AAE_E seem to have slightly smaller means than that of the AAE_T.

8.4 Conclusions

This chapter compares the AAEs obtained using model calibration with the AAEs obtained using model estimation. The estimation algorithm estimates the parameters of a model by fitting the relevant distribution to the log-returns of the stock using a maximum likelihood estimator. The models considered are the exponential $N \circ IG$ model and the Black-Scholes model. As in the previous chapter we consider three types of exotic options; digital barrier call options, down-and-out barrier call options and up-and-out barrier call options. The prices of these options are obtained from a hypothetical financial market in which the stock price follows an exponential $N \circ IG$ model.

In order to aid the comparisons between the AAEs associated with the calibrations and estimations we classify a given calibration or estimation as perfect if the mean of the AAE associated with the calibration or estimation does not exceed the mean of the AAE under the true model. We use two formal tests in this classification exercise; a mixed Bayesian-frequentist hypothesis test and a bootstrap based hypothesis test.

A total of six calibrations are performed and five of these are classified as perfect. The only calibration that is not deemed perfect is that of the Black-Scholes model to the down-and-out barrier option prices. Similarly six estimations are performed. Surprisingly, all of these estimations are deemed perfect.

That more estimations than calibrations are deemed perfect is surprising since the aim of a calibration algorithm is to minimise the AAE. However, the optimisation algorithm used in the calibration process converges to a local minimum and not a global minimum. Therefore, it is possible for estimations to outperform the corresponding calibrations. As was mentioned in the previous chapter, the non-convex nature of the calibration problem prevents us from achieving the global minimum.

The AAEs associated with the calibrations and estimations are generally small. We conclude that it is possible to replicate barrier option prices very accurately not only by using calibration procedures but also by using estimation procedures. As a result, barrier option pricing models are interchangeable.

Appendix A

Proof that (4.7) and (4.8) hold for the $N \circ IG$ process

Theorem 37 *The following two properties hold for a $N \circ IG$ process with Lévy measure $\nu(dx)$:*

$$\nu((-\infty, a]) > 0, \tag{A.1}$$

for all $a \in \mathbb{R}$, and

$$\nu \text{ has no atom and } \int_{[-1,0)} |x| \nu(dx) = \int_{(0,1]} x\nu(dx) = \infty, \tag{A.2}$$

Proof. Proof that (A.1) holds for the $N \circ IG(\alpha, \beta, \mu, \delta)$ process:

For a $N \circ IG(\alpha, \beta, \mu, \delta)$ process $\alpha > 0$, $\beta < |\alpha|$, $\delta > 0$ and

$$\nu(dx) = \frac{\delta\alpha}{\pi} \frac{e^{\beta x} K_1(\alpha|x|)}{|x|} dx,$$

where

$$K_1(z) = \int_0^\infty \exp\left(-\frac{1}{2} z\left(u + u^{-1}\right)\right) du.$$

As as result,

$$\nu((-\infty, a]) = \frac{\delta\alpha}{\pi} \int_{-\infty}^{a} \frac{e^{\beta x} K_1(\alpha|x|)}{|x|} dx > 0,$$

for all $a \in \mathbb{R}$, since $K_1(z) > 0$ for all $z \in \mathbb{R}$.

Proof that (A.2) holds for the $N \circ IG$ process:

From its definition, v has no atom. It remains to show that

$$\int_{[-1,0)} |x| \, \nu(dx) = \int_{(0,1]} x\nu(dx) = \infty. \tag{A.3}$$

Consider the first integral in (A.3):

$$\begin{aligned}
\int_{[-1,0)} |x| \, \nu(dx) &= \frac{\delta\alpha}{\pi} \int_{-1}^{0} e^{\beta x} K_1(-\alpha x) \, dx \\
&= \frac{\delta\alpha}{2\pi} \int_{-1}^{0} \int_{0}^{\infty} \exp\left(\left[\beta + \frac{\alpha}{2} \left(u + u^{-1} \right) \right] x \right) du dx \\
&= \frac{\delta\alpha}{2\pi} \int_{0}^{\infty} \int_{-1}^{0} \exp\left(\left[\beta + \frac{\alpha}{2} \left(u + u^{-1} \right) \right] x \right) dx du \\
&= \frac{\delta\alpha}{2\pi} \int_{0}^{\infty} \frac{1 - \exp\left(- \left[\beta + \frac{\alpha}{2} \left(u + u^{-1} \right) \right] \right)}{\beta + \frac{\alpha}{2} \left(u + u^{-1} \right)} du.
\end{aligned} \tag{A.4}$$

In order to simplify this integral note that, for $u \geq 0$,

$$\begin{aligned}
& u + u^{-1} \geq 2 \\
\Longrightarrow \quad & \frac{\alpha}{2} \left(u + u^{-1} \right) \geq \alpha \\
\Longrightarrow \quad & - \left(\beta + \frac{\alpha}{2} \left(u + u^{-1} \right) \right) \leq - (\alpha + \beta) \\
\Longrightarrow \quad & 1 - \exp\left(- \left(\beta + \frac{\alpha}{2} \left(u + u^{-1} \right) \right) \right) \geq 1 - e^{-(\alpha+\beta)}.
\end{aligned} \tag{A.5}$$

Using (A.4) and (A.5),

$$\begin{aligned}\int_{[-1,0)} |x|\,\nu(dx) &\geq \frac{\delta\alpha}{2\pi}\left(1-e^{-(\alpha+\beta)}\right)\int_0^\infty \frac{1}{\beta+\frac{\alpha}{2}\left(u+u^{-1}\right)}du \\ &\geq \frac{\delta\alpha}{2\pi}\left(1-e^{-(\alpha+\beta)}\right)\int_1^\infty \frac{1}{\beta+\frac{\alpha}{2}\left(u+u^{-1}\right)}du, \qquad \text{(A.6)}\end{aligned}$$

where the last inequality holds since the integrand in (A.6) is positive. Note that, for all $u \geq 1$,

$$u+u^{-1} \leq 2u \implies \frac{1}{\beta+\frac{\alpha}{2}\left(u+u^{-1}\right)} \geq \frac{1}{\beta+\alpha u}. \qquad \text{(A.7)}$$

Using (A.6) and (A.7), we obtain

$$\begin{aligned}\int_{[-1,0)} |x|\,\nu(dx) &\geq \frac{\delta\alpha}{2\pi}\left(1-e^{-(\alpha+\beta)}\right)\int_1^\infty \frac{1}{\beta+\alpha u}du \\ &= \frac{\delta}{2\pi}\left(1-e^{-(\alpha+\beta)}\right)\int_{\alpha+\beta}^\infty \frac{1}{y}dy \\ &= \infty. \qquad \text{(A.8)}\end{aligned}$$

Consider the second integral in (A.3):

$$\begin{aligned}
\int_{(0,1]} x\nu\left(dx\right) &= \frac{\delta\alpha}{\pi}\int_0^1 e^{\beta x}K_1\left(\alpha x\right)dx \\
&= \frac{\delta\alpha}{2\pi}\int_0^1\int_0^\infty \exp\left(\left[\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right]x\right)dudx \\
&= \frac{\delta\alpha}{2\pi}\int_0^\infty\int_0^1 \exp\left(\left[\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right]x\right)dxdu \\
&= \frac{\delta\alpha}{2\pi}\int_0^\infty \frac{\exp\left(\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right)-1}{\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)}du \\
&= \frac{\delta\alpha}{2\pi}\int_0^\infty \frac{1-\exp\left(\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right)}{\frac{\alpha}{2}\left(u+u^{-1}\right)-\beta}du. \qquad \text{(A.9)}
\end{aligned}$$

For all $u \geq 0$,

$$\begin{aligned}
u+u^{-1}\geq 2 \quad &\Longrightarrow \quad \exp\left(\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right)\leq e^{\beta-\alpha} \\
&\Longrightarrow \quad 1-\exp\left(\beta-\frac{\alpha}{2}\left(u+u^{-1}\right)\right)\geq 1-e^{\beta-\alpha}. \qquad \text{(A.10)}
\end{aligned}$$

Combining (A.9) and (A.10) we obtain

$$\begin{aligned}
\int_{(0,1]} x\nu\left(dx\right) &\geq \frac{\delta\alpha}{2\pi}\left(1-e^{\beta-\alpha}\right)\int_0^\infty \frac{1}{\frac{\alpha}{2}\left(u+u^{-1}\right)-\beta}du \\
&\geq \frac{\delta\alpha}{2\pi}\left(1-e^{\beta-\alpha}\right)\int_1^\infty \frac{1}{\frac{\alpha}{2}\left(u+u^{-1}\right)-\beta}du, \qquad \text{(A.11)}
\end{aligned}$$

where the last inequality holds because the integrand is positive.

For all $u \geq 1$,

$$u + u^{-1} \leq 2u \implies \frac{1}{\frac{\alpha}{2}\left(u + u^{-1}\right) - \beta} \geq \frac{1}{\alpha u - \beta}. \tag{A.12}$$

Combining the results in (A.11) and (A.12), we obtain

$$\begin{aligned} \int_{(0,1]} x\nu\left(dx\right) &\geq \frac{\delta\alpha}{2\pi}\left(1 - e^{\beta-\alpha}\right)\int_{1}^{\infty}\frac{1}{\alpha u - \beta}du \\ &= \frac{\delta}{2\pi}\left(1 - e^{\beta-\alpha}\right)\int_{\alpha-\beta}^{\infty}\frac{1}{y}dy \\ &= \infty. \end{aligned} \tag{A.13}$$

Property (A.2) is shown to hold by combining the results in (A.8) and (A.13). ■

Appendix B

Derivation of the Lévy characteristics under the GMCMM

Let L_t be a Lévy process with triplet of Lévy characteristics given by $\left(\gamma, \sigma^2, \nu\left(dx\right)\right)$ under probability measure P. Under the probability measure P^θ, where P^θ is the GMCMM introduced in Section 4.4, the Lévy process changes to

$$Y_t = \theta L_t + rt - \log\left(E^P\left[\exp\left(\theta L_t\right)\right]\right), \tag{B.1}$$

for $\theta > 0$.

Below we derive the triplet of Lévy characteristics of θL_t. From the Lévy-Khintchine formula we know that

$$\log\left(E^P\left[\exp\left(aL_t\right)\right]\right) = t\left\{\gamma a + \frac{\sigma^2 a^2}{2} + \int_{-\infty}^{\infty}\left(e^{ax} - 1 - ax\mathbb{I}\left(|x| < 1\right)\right)\nu\left(dx\right)\right\}. \tag{B.2}$$

Setting $a = u\theta$ in (B.2) we obtain

$$\begin{aligned}
\log\left(E^P\left[\exp\left(u\theta L_t\right)\right]\right) &= t\left\{\theta\gamma u + \frac{\theta^2\sigma^2u^2}{2} + \int_{-\infty}^{\infty}\left(e^{\theta ux} - 1 - \theta ux\mathbb{I}\left(|x|<1\right)\right)\nu\left(dx\right)\right\} \\
&= t\left\{\theta\gamma u + \frac{\theta^2\sigma^2u^2}{2} + \int_{-\infty}^{\infty}\left(e^{\theta ux} - 1 - \theta ux\mathbb{I}\left(|\theta x|<1\right)\right)\nu\left(dx\right)\right. \\
&\quad \left. + u\int_{-\infty}^{\infty}\theta x\left(\mathbb{I}\left(|\theta x|<1\right) - \mathbb{I}\left(|\theta x|<\theta\right)\right)\nu\left(dx\right)\right\} \\
&= t\left\{u\left[\theta\gamma + \int_{-\infty}^{\infty} y\left(\mathbb{I}\left(|y|<1\right) - \mathbb{I}\left(|y|<\theta\right)\right)\nu\left(\frac{dy}{\theta}\right)\right]\right. \\
&\quad \left. + \frac{\theta^2\sigma^2u^2}{2} + \int_{-\infty}^{\infty}\left(e^{uy} - 1 - uy\mathbb{I}\left(|y|<1\right)\right)\nu\left(\frac{dy}{\theta}\right)\right\}, \qquad \text{(B.3)}
\end{aligned}$$

where we use the substitution $y = \theta x$ to obtain the last equality. The triplet of Lévy characteristics of θL_t can be obtained from (B.3); this triplet is given by

$$\left(\theta\gamma + \int_{-\infty}^{\infty} x\left(\mathbb{I}\left(|x|<1\right) - \mathbb{I}\left(|x|<\theta\right)\right)\nu\left(\frac{dx}{\theta}\right), \theta^2\sigma^2, \nu\left(\frac{dx}{\theta}\right)\right). \qquad \text{(B.4)}$$

The Brownian motion part and the jump measure of Y_t defined in (B.1) equals that of θL_t. Setting $a = \theta$ in (B.2) we obtain

$$\begin{aligned}
\log\left(E^P\left[\exp\left(\theta L_t\right)\right]\right) &= t\left\{\theta\gamma + \frac{\theta^2\sigma^2}{2} + \int_{-\infty}^{\infty}\left(e^{\theta x} - 1 - \theta x\mathbb{I}\left(|x|<1\right)\right)\nu\left(dx\right)\right\} \\
&= t\left\{\theta\gamma + \frac{\theta^2\sigma^2}{2} + \int_{-\infty}^{\infty}\left(e^{y} - 1 - y\mathbb{I}\left(|y|<\theta\right)\right)\nu\left(\frac{dy}{\theta}\right)\right\}. \qquad \text{(B.5)}
\end{aligned}$$

The drift of Y_t can be calculated using (B.4) and (B.5) as follows:

$$\begin{aligned}
\gamma_{GMCMM} &= \theta\gamma + \int_{-\infty}^{\infty} x\left(\mathbb{I}\left(|x|<1\right) - \mathbb{I}\left(|x|<\theta\right)\right)\nu\left(\frac{dx}{\theta}\right) + r - \frac{\log\left(E^P\left[\exp\left(\theta L_t\right)\right]\right)}{t} \\
&= r - \frac{\theta^2\sigma^2}{2} + \int_{-\infty}^{\infty} x\left(\mathbb{I}\left(|x|<1\right) - \mathbb{I}\left(|x|<\theta\right)\right)\nu\left(\frac{dx}{\theta}\right) \\
&\quad - \int_{-\infty}^{\infty}\left(e^x - 1 - x\mathbb{I}\left(|x|<\theta\right)\right)\nu\left(\frac{dx}{\theta}\right) \\
&= r - \frac{\theta^2\sigma^2}{2} - \int_{-\infty}^{\infty}\left(e^x - 1 - x\mathbb{I}\left(|x|<1\right)\right)\nu\left(\frac{dx}{\theta}\right). \qquad \text{(B.6)}
\end{aligned}$$

Combining (B.4) and (B.6) the triplet of Lévy characteristics is calculated to be

$$\left(r - \frac{\theta^2\sigma^2}{2} - \int_{-\infty}^{\infty}\left(e^x - 1 - x\mathbb{I}\left(|x|<1\right)\right)\nu\left(\frac{dx}{\theta}\right), \theta^2\sigma^2, \nu\left(\frac{dx}{\theta}\right)\right).$$

Appendix C

Bounds for European call option prices under the GMCMM

Let

$$S_t = S_0 \exp\left(L_t\right),$$

under P where L_t is a Lévy process with the properties shown in (4.7) and (4.8). Let P^θ be the probability measure obtained using the GMCMM for arbitrary $\theta > 0$.

Theorem 38 *The price of a European call option, with strike price $K > 0$ and time to maturity $T > 0$, calculated with respect to P^θ is contained in the interval*

$$\left(\left(S_0 - e^{-rT}K\right)^+, S_0\right). \tag{C.1}$$

Proof. Denote by $\pi\left(Q\right)$ the price of the option calculated with respect to the probability measure Q;

$$\begin{aligned}
\pi\left(P^\theta\right) &= E^{P^\theta}\left[e^{-rT}\left(S_T - K\right)^+\right] \\
&\leq E^{P^\theta}\left[e^{-rT}S_T\right] \\
&= S_0, \qquad \text{(C.2)}
\end{aligned}$$

where the last equality above holds since $e^{-rt}S_t$ is a P^θ-martingale. However, in order to prove the theorem we need to show that the inequality in (C.2) is strict; i.e. we need to show that $\pi\left(P^\theta\right) < S_0$. Let F_T denote the distribution function of S_T under P^θ. Consider

$$\begin{aligned}
\pi\left(P^\theta\right) - S_0 &= E^{P^\theta}\left[e^{-rT}\left(S_T - K\right)^+\right] - E^{P^\theta}\left[e^{-rT}S_T\right] \\
&= e^{-rT}\left\{\int_K^\infty (x-K)\,F_T\,(dx) - \int_0^\infty x F_T\,(dx)\right\} \\
&= e^{-rT}\left\{-KP\left(S_T \geq K\right) - \int_0^K x F_T\,(dx)\right\} \\
&< 0. \qquad \text{(C.3)}
\end{aligned}$$

For the terms in the second last step to equal 0 we require that $S_T = K$ almost surely. However, properties (4.7) and (4.8) excludes this possibility. As a result, the inequality in (C.3) is strict.

We now turn our attention to the lower bound shown in (C.1). Let

$$g\left(x\right) = \left(x - e^{-rT}K\right)^+. \qquad \text{(C.4)}$$

Since g is a convex function and $e^{-rt}S_t$ forms a P^θ-martingale, $g\left(e^{-rt}S_t\right)$ forms a sub-martingale. Therefore,

$$\begin{aligned}
\pi\left(P^\theta\right) &= E^{P^\theta}\left[g\left(e^{-rT}S_T\right)\right] \\
&\geq g\left(E^{P^\theta}\left[e^{-rT}S_T\right]\right) \\
&= \left(S_0 - e^{-rT}K\right)^+, \qquad \text{(C.5)}
\end{aligned}$$

where the inequality in the second step above is obtained using Jensen's inequality. In the last step we use the P^θ-martingaleness of $e^{-rt}S_t$.

It remains to show that

$$\pi\left(P^\theta\right) \neq \left(S_0 - e^{-rT}K\right)^+. \tag{C.6}$$

Assume that (C.6) does not hold. It then follows that

$$\begin{aligned}
\pi\left(P^\theta\right) &= \left(S_0 - e^{-rT}K\right)^+ \\
&= \left(E^{P^\theta}\left[e^{-rT}S_T\right] - e^{-rT}K\right)^+ \\
&= g\left(E^{P^\theta}\left[e^{-rT}S_T\right]\right). \qquad \text{(C.7)}
\end{aligned}$$

However,

$$\pi\left(P^\theta\right) = E^{P^\theta}\left[g\left(e^{-rT}S_T\right)\right]. \tag{C.8}$$

Combining (C.7) and (C.8) we have that

$$E^{P^\theta}\left[g\left(e^{-rT}S_T\right)\right] = g\left(E^{P^\theta}\left[e^{-rT}S_T\right]\right). \tag{C.9}$$

Since g is a convex function, (C.9) implies that g must be linear on the support of the random variable $e^{-rT}S_T$. It follows from the properties in (4.7) and (4.8) that the support of L under the measure P is $(-\infty, \infty)$. From (4.11) we know that the support of L under the measure P^θ must then also be $(-\infty, \infty)$ because this change of measure induces a translation and scale change in the values attainable by L. As a result, the support of $e^{-rT}S_T$ is $(0, \infty)$. However, since $K > 0$, we know that g cannot be linear; $g(x) = 0$ for $x \leq e^{-rT}K$ and $g(x) = x - e^{-rT}K$ for $x > e^{-rT}K$. As a result, (C.9) leads to a contradiction, meaning that (C.7) cannot hold and $\pi\left(P^\theta\right)$ cannot equal $\left(S_0 - e^{-rT}K\right)^+$.

Together with (C.5) this means that

$$\pi\left(P^{\theta}\right) > \left(S_0 - e^{-rT}K\right)^{+}.$$

Combining this result with (C.3) we have that

$$\pi\left(P^{\theta}\right) \in \left(\left(S_0 - e^{-rT}K\right)^{+}, S_0\right).$$

■

Appendix D

Additional numerical calibration results for European options

The numerical results discussed in Section 5.2 are shown below. We consider six datasets consisting of European option prices; call and put options are considered for the S&P 500 index, the PowerShares index and Google shares. We present the results pertaining to each of these datasets below. For each set of call option datasets we include 4 tables. The first table reports the parameters obtained when using the exponential $N \circ IG$ model, while the second reports three distance measures (the AAE, the ARE and the $RMSE$) calculated using the parameters reported in the preceding table as well as the time taken (in seconds) to calculate each of the results. The third and fourth tables provide the same information for the Black-Scholes model as the first two tables provide for the exponential $N \circ IG$ model. Similarly, four tables are included for each of the put option datasets as well.

The results shown below are obtained using the procedure detailed in Section 5.2. We do not include an interpretation for each of the tables presented below. However, some general remarks are in order. As before we are mainly interested in the distance measures associated with each of the calibration methods considered. Consider first the results obtained using the exponential $N \circ IG$ model. When using this model the β calibration, the μ calibration and the application of the GMCMM each provide a single free parameter that can be manipulated to reduce the AAE. The AAE associated with these calibration procedures are generally smaller than the AAE calculated using either the Esscher transform or the MCMM. The Esscher and MCMM restricted calibrations each provide

three free parameters. Generally these methods provide a smaller AAE than is the case for the calibration procedures relying on the manipulation of a single parameter. The full calibration procedure allows the manipulation of four parameters. However, because of the numerical problems discussed previously, this method performs better than the Esscher and MCMM restricted calibrations in some cases and worse in others.

Consider the results associated with the Black-Scholes model. The μ calibration procedure, the application of the GMCMM and the Esscher restricted calibration each introduce a single free parameter. As a result, the AAEs associated with these procedures are smaller than the AAE associated with the prices calculated using the Esscher transform. In general the GMCMM achieves a smaller AAE value than the μ calibration and the Esscher restricted calibration achieves a smaller AAE still. In the case of the Black-Scholes model the full calibration outperforms the Esscher calibration for the results below.

For ease of reference we include the results for the calibrations relating to the call options on the S&P 500 index below. These results are discussed in detail in Section 5.2.

D.1 S&P 500 options

Table D.1 shows the parameters obtained for the exponential $N \circ IG$ model, while Table D.2 shows the distance measures and times taken for this model relating to the S&P 500 call options. Tables D.3 and D.4 report the same information relating to the Black-Scholes model. Tables D.5 to D.8 provide the same information for the put options than Tables D.1 to D.4 provide for the call options.

	α	β	μ	δ
Estimated parameters	49.673	-6.5019	0.0014	0.0107
Esscher transform	49.673	-7.0377	0.0014	0.0107
MCMM	49.673	-6.5019	0.0013	0.0107
β calibration	49.673	-7.3576	0.0014	0.0107
μ calibration	49.673	-6.5019	0.0012	0.0107
GMCMM	66.5443	-8.7103	0.001	0.008
Esscher calibration (start)	5	-0.4704	0	0.001
Esscher calibration	14.839	-10.8948	0.0008	0.0008
MCMM calibration (start)	20	10	-0.0006	0.001
MCMM calibration	15.5674	11.4427	0.0009	0.0009
Full calibration (start)	50	10	0	0.001
Full calibration	115.9247	2.3349	-0.0001	0.0088

Table D.1: Parameters obtained using the exponential $N \circ IG$ model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	16.9058	3.5015	23.1673	0.359
MCMM	16.8282	3.4935	23.0692	0.28
β calibration	15.3201	3.0706	21.0684	30.514
μ calibration	15.2577	3.0636	21.0076	31.154
GMCMM	9.771	0.7973	16.9016	28.891
Esscher calibration (start)	13.5186	6.3216	19.2273	92.851
Esscher calibration	7.7669	0.2778	14.5842	142.506
MCMM calibration (start)	13.7018	2.2615	21.9438	73.008
MCMM calibration	7.7668	0.2829	14.6017	138.637
Full calibration (start)	14.3147	0.4861	26.8601	9444.506
Full calibration	8.2198	0.3288	15.56	771.5

Table D.2: Empirical results obtained using the exponential $N \circ IG$ model (call options).

	μ	σ
Estimated parameters	0	0.0146
Esscher transform	-0.0001	0.0146
μ calibration	-0.0002	0.0146
GMCMM	-0.0001	0.0109
Esscher calibration (start)	-0.0001	0.012
Esscher calibration	-0.0001	0.011
Full calibration (start)	0.0001	0.008
Full calibration	0	0.009

Table D.3: Parameters obtained using the Black-Scholes model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	16.269	3.3055	22.3107	0.234
μ calibration	15.0098	2.8985	20.8142	22.099
GMCMM	9.8906	0.7487	17.1587	23.478
Esscher calibration (start)	10.7389	1.2802	16.8804	1.997
Esscher calibration	9.8825	0.7812	17.0802	5.694
Full calibration (start)	8.9814	0.3196	17.4187	11.577
Full calibration	8.2804	0.3776	15.5019	13.542

Table D.4: Empirical results obtained using the Black-Scholes model (call options).

	α	β	μ	δ
Estimated parameters	49.673	-6.5019	0.0014	0.0107
Esscher transform	49.673	-7.0377	0.0014	0.0107
MCMM	49.673	-6.5019	0.0013	0.0107
β calibration	49.673	-7.1443	0.0014	0.0107
μ calibration	49.673	-6.5019	0.0013	0.0107
GMCMM	49.4731	-6.4758	0.0013	0.0107
Esscher calibration (start)	5	-0.4704	0	0.001
Esscher calibration	2.0818	-1.1286	0.0002	0.0005
MCMM calibration (start)	5	0	-0.0001	0.001
MCMM calibration	2.0817	-1.1285	0.0002	0.0005
Full calibration	4.3908	-3.1477	0.0004	0.0005

Table D.5: Parameters obtained using the exponential $N \circ IG$ model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	8.7449	0.5632	19.9823	0.515
MCMM	8.751	0.5641	20.0079	0.452
β calibration	8.64	0.5623	19.2176	37.752
μ calibration	8.64	0.5623	19.0901	37.377
GMCMM	8.7474	0.5635	19.9644	46.41
Esscher calibration (start)	8.4424	0.6787	20.5668	127.717
Esscher calibration	5.8711	2.4774	16.8178	180.523
MCMM calibration (start)	8.6914	0.6349	20.622	107.796
MCMM calibration	5.8711	2.4774	16.8178	318.038
Full calibration	5.7299	2.3981	16.0385	256.121

Table D.6: Empirical results obtained using the exponential $N \circ IG$ model (put options).

	μ	δ
Estimated parameters	0	0.0146
Esscher transform	-0.0001	0.0146
μ calibration	-0.0001	0.0146
GMCMM	-0.0001	0.0151
Esscher calibration (start)	-0.0001	0.016
Esscher calibration	-0.0001	0.0151
Full calibration	-0.0001	0.0143

Table D.7: Parameters obtained using the Black-Scholes model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	8.8983	0.5942	20.2902	0.359
μ calibration	8.7223	0.5925	19.2548	28.803
GMCMM	8.8534	0.5877	19.8794	36.131
Esscher calibration (start)	9.1583	0.5849	19.7197	3.151
Esscher calibration	8.8514	0.5882	19.9087	5.694
Full calibration	8.7075	0.5961	19.26	11.28

Table D.8: Empirical results obtained using the Black-Scholes model (put options).

D.2 PowerShares options

Table D.9 reports the parameters obtained using the exponential $N \circ IG$ model where calibration is done with respect to the call option prices. Table D.10 shows the distance measures associated with the parameter sets shown in Table D.9 as well as the times required to obtain the results. Tables D.11 and D.12 provide the same results for the Black-Scholes model. Tables D.13 to D.16 show the same results as Tables D.9 to D.12, with the exception that the results in Tables D.13 to D.16 relate to the put option prices.

	α	β	μ	δ
Estimated parameters	69.9467	-10.165	0.0026	0.0151
Esscher transform	69.9467	-12.2708	0.0026	0.0151
MCMM	69.9467	-10.165	0.0021	0.0151
β calibration	69.9467	-12.2336	0.0026	0.0151
μ calibration	69.9467	-10.165	0.0021	0.0151
GMCMM	77.6314	-11.2817	0.0019	0.0136
Esscher calibration (start)	5	-0.4704	0	0.001
Esscher calibration	6.5525	-4.9182	0.0006	0.0007
MCMM calibration (start)	100	-95	0.0029	0.001
MCMM calibration	101.1419	-94.777	0.0026	0.001
Full calibration	98.6023	9.9995	-0.0009	0.0097

Table D.9: Parameters obtained using the exponential $N \circ IG$ model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	0.9325	1.0475	1.3425	0.296
MCMM	0.9271	1.0254	1.3414	0.172
β calibration	0.914	1.0608	1.3288	721.532
μ calibration	0.9191	1.0421	1.3235	737.07
GMCMM	0.8726	0.6358	1.3573	18.72
Esscher calibration (start)	0.8856	1.3443	1.3362	67.736
Esscher calibration	0.6702	0.2045	1.1319	61.854
MCMM calibration (start)	0.8742	0.4033	1.2554	52.775
MCMM calibration	0.7151	0.2519	1.2156	68.687
Full calibration	0.6638	0.2207	1.1033	349.753

Table D.10: Empirical results obtained using the exponential $N \circ IG$ model (call options).

	μ	δ
Estimated parameters	0.0004	0.0148
Esscher transform	−0.0001	0.0148
μ calibration	−0.0001	0.0148
GMCMM	−0.0001	0.0133
Esscher calibration (start)	−0.0001	0.012
Esscher calibration	−0.0001	0.0133
Full calibration	0.0001	0.0101

Table D.11: Parameters obtained using the Black-Scholes model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	0.9298	1.0406	1.3451	0.156
μ calibration	0.9179	1.0986	1.2977	14.354
GMCMM	0.8794	0.6482	1.364	15.46
Esscher calibration (start)	0.9151	0.3807	1.4487	1.31
Esscher calibration	0.8793	0.6322	1.3669	2.091
Full calibration	0.662	0.2169	1.1022	8.455

Table D.12: Empirical results obtained using the Black-Scholes model (call options).

	α	β	μ	δ
Estimated parameters	69.9467	-10.165	0.0026	0.0151
Esscher transform	69.9467	-12.2708	0.0026	0.0151
MCMM	69.9467	-10.165	0.0021	0.0151
β calibration	69.9467	-12.2336	0.0026	0.0151
μ calibration	69.9467	-10.165	0.0021	0.0151
GMCMM	69.6652	-10.1241	0.0021	0.0151
Esscher calibration (start)	5	-0.4704	0	0.001
Esscher calibration	3.3401	-1.6477	0.0003	0.0008
MCMM calibration (start)	100	-95	0.0029	0.001
MCMM calibration	98.7958	-94.9771	0.0019	0.0006
Full calibration	2.042	-0.7659	0.0002	0.0007

Table D.13: Parameters obtained using the exponential $N \circ IG$ model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	0.3794	0.4126	0.5561	0.281
MCMM	0.3805	0.4157	0.5604	0.171
β calibration	0.3784	0.4135	0.5621	723.358
μ calibration	0.381	0.4169	0.5695	738.474
GMCMM	0.3804	0.4144	0.5584	19.064
Esscher calibration (start)	0.4144	0.3743	0.6456	65.941
Esscher calibration	0.1944	0.5101	0.3321	77.111
MCMM calibration (start)	0.36	0.5253	0.5187	48.782
MCMM calibration	0.2393	0.4696	0.4019	62.384
Full calibration	0.2037	0.5633	0.3459	49.078

Table D.14: Empirical results obtained using the exponential $N \circ IG$ model (put options).

	μ	σ
Estimated parameters	0.0004	0.0148
Esscher transform	−0.0001	0.0148
μ calibration	−0.0001	0.0148
GMCMM	−0.0001	0.0149
Esscher calibration (start)	−0.0001	0.016
Esscher calibration	−0.0001	0.0149
Full calibration	−0.0001	0.0159

Table D.15: Parameters obtained using the Black-Scholes model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	0.3869	0.4274	0.569	0.156
μ calibration	0.3877	0.4287	0.5789	13.137
GMCMM	0.3867	0.426	0.5668	15.366
Esscher calibration (start)	0.4161	0.4092	0.5865	1.357
Esscher calibration	0.3866	0.4242	0.5642	2.324
Full calibration	0.3818	0.4102	0.5838	4.852

Table D.16: Empirical results obtained using the Black-Scholes model (put options).

D.3 Google options

In Tables D.17 to D.24 we report the same results for the Google options as are reported in Tables D.9 to D.16 for the PowerShares options.

	α	β	μ	δ
Estimated parameters	47.5159	−5.7369	0.0024	0.0154
Esscher transform	47.5159	−7.6894	0.0024	0.0154
MCMM	47.5159	−5.7369	0.0017	0.0154
β calibration	47.5159	−7.5729	0.0024	0.0154
μ calibration	47.5159	−5.7369	0.0017	0.0154
GMCMM	49.8841	−6.0229	0.0016	0.0146
Esscher calibration (start)	300	−0.4821	0	0.1
Esscher calibration	420.5982	−416.003	0.0036	0.0006
MCMM calibration (start)	300	0	−0.0002	0.1
MCMM calibration	26.4058	−0.0614	−0.0001	0.0079
Full calibration	349.2106	−0.512	0.0001	0.078

Table D.17: Parameters obtained using the exponential $N \circ IG$ model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	5.5328	0.3626	10.028	0.421
MCMM	5.5016	0.3543	10.05	0.328
β calibration	5.2903	0.3855	9.0931	33.197
μ calibration	5.1899	0.3752	9.0885	32.776
GMCMM	5.3546	0.3005	10.2441	34.554
Esscher calibration (start)	5.6222	0.3781	10.0954	110.276
Esscher calibration	4.1743	0.358	8.818	249.819
MCMM calibration (start)	5.6224	0.3782	10.0955	90.199
MCMM calibration	5.4149	0.2879	10.2897	177.949
Full calibration	3.4901	0.2557	7.4707	1038.042

Table D.18: Empirical results obtained using the exponential $N \circ IG$ model (call options).

	μ	σ
Estimated parameters	0.0005	0.0188
Esscher transform	-0.0002	0.0188
μ calibration	-0.0002	0.0188
GMCMM	-0.0001	0.0172
Esscher calibration (start)	-0.0001	0.016
Esscher calibration	-0.0001	0.0173
Full calibration	0	0.0149

Table D.19: Parameters obtained using the Black-Scholes model (call options).

	AAE	ARE	$RMSE$	Time
Esscher transform	5.8862	0.4223	10.1412	0.265
μ calibration	5.81	0.4307	9.8341	25.18
GMCMM	5.423	0.3098	10.3024	28.44
Esscher calibration (start)	5.919	0.272	11.0131	2.544
Esscher calibration	5.4229	0.3123	10.2861	4.555
Full calibration	3.4823	0.2559	7.4582	15.725

Table D.20: Empirical results obtained using the Black-Scholes model (call options).

	α	β	μ	δ
Estimated parameters	47.5159	−5.7369	0.0024	0.0154
Esscher transform	47.5159	−7.6894	0.0024	0.0154
MCMM	47.5159	−5.7369	0.0017	0.0154
β calibration	47.5159	−7.4331	0.0024	0.0154
μ calibration	47.5159	−5.7369	0.0018	0.0154
GMCMM	53.7609	−6.4909	0.0015	0.0136
Esscher calibration (start)	300	−0.4821	0	0.1
Esscher calibration	295.5133	−289.6395	0.0033	0.0007
MCMM calibration (start)	100	−95	0.0029	0.001
MCMM calibration	100.5458	−94.7995	0.003	0.0011
Full calibration	315.8269	−0.4942	0.0001	0.0916

Table D.21: Parameters obtained using the exponential $N \circ IG$ model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	4.6524	0.4028	8.0946	0.344
MCMM	4.5178	0.4028	7.9176	0.296
β calibration	3.5274	0.3934	6.4467	26.77
μ calibration	3.5076	0.396	6.4161	26.582
GMCMM	3.5414	0.4304	6.3511	29.297
Esscher calibration (start)	4.6734	0.4156	8.1135	99.933
Esscher calibration	2.6435	0.5881	5.6684	177.466
MCMM calibration (start)	2.7473	0.7026	5.7542	77.486
MCMM calibration	2.6314	0.587	5.6624	81.557
Full calibration	3.4621	0.4259	6.2595	980.493

Table D.22: Empirical results obtained using the exponential $N \circ IG$ model (put options).

	μ	σ
Estimated parameters	0.0005	0.0188
Esscher transform	-0.0002	0.0188
μ calibration	-0.0001	0.0188
GMCMM	-0.0001	0.0159
Esscher calibration (start)	-0.0001	0.016
Esscher calibration	-0.0001	0.0159
Full calibration	-0.0001	0.017

Table D.23: Parameters obtained using the Black-Scholes model (put options).

	AAE	ARE	$RMSE$	Time
Esscher transform	5.2114	0.4214	8.8121	0.234
μ calibration	3.6821	0.4049	6.6444	21.107
GMCMM	3.5981	0.4462	6.3821	23.837
Esscher calibration (start)	3.6	0.4448	6.3887	2.122
Esscher calibration	3.5971	0.4486	6.3757	3.51
Full calibration	3.46	0.4261	6.2614	6.973

Table D.24: Empirical results obtained using the Black-Scholes model (put options).

Appendix E

Algorithms for the simulation of price paths

Let n represent the number of days that we wish to simulate the stock price for and let N be the number of simulated price paths that we require. Below we provide the algorithms used to simulate price paths for the various models used in Chapters 7 and 8.

E.1 The Heston model

1. Generate σ^2 from a $gamma\left(\frac{2\kappa\eta}{\theta^2}, \frac{\theta^2}{2\kappa}\right)$ distribution.
2. Generate W_1 from a $N(0,1)$ distribution.
3. Calculate $S(1,k) = 1 + \mu + \sigma W_1$.
4. Generate W_2 and W_3 independently from a $N(0,1)$ distribution.
5. Calculate $W_4 = \rho W_2 + \sqrt{1-\rho^2} W_3$.
6. Calculate $S(j,k) = (1 + r + \sigma W_2) S(j-1,k)$.
7. Calculate $\sigma^2 = \max\left(\sigma^2 + \kappa\left(\eta - \sigma^2\right) + \sigma\theta W_4, 0\right)$.
8. Repeat steps 4 to 7 for $j = 2$ to n.
9. Repeat steps 1 to 8 for $k = 1$ to N.

E.2 The time changed exponential $N \circ IG$ model

1. Generate y from a $gamma\left(\frac{2\kappa\eta}{\lambda^2}, \frac{\lambda^2}{2\kappa}\right)$ distribution.

2. Generate W from a $N(0,1)$ distribution.

3. Calculate $y = \max\left(y + \kappa(\eta - y) + \lambda\sqrt{y}W, 0.001\right)$.

4. Generate $L(j,k)$ from a $N \circ IG(\alpha, \beta, \mu y, \delta y)$ distribution.

5. Repeat steps 2 to 4 for $j = 1$ to n.

6. Repeat steps 1 to 5 for $k = 1$ to N.

7. Calculate $S(1,k) = \exp(L(1,k))$.

8. Calculate $S(j,k) = S(j-1,k)\exp(L(j-1,k))$.

9. Repeat step 8 for $j = 2$ to n.

10. Repeat steps 7 to 9 for $k = 1$ to N.

E.3 Exponential Lévy models

1. Generate an n by N matrix Δ where each element of Δ follows a $N \circ IG(\alpha, \beta, \mu, \delta)$ distribution.

2. Calculate $\Delta^{\star}(j,k) = \sum_{l=1}^{j} \Delta(l,k)$.

3. Calculate $S(j,k) = \exp(\Delta^{\star}(j,k))$.

4. Repeat steps 2 and 3 for $j = 1$ to n.

5. Repeat steps 2 to 4 for $k = 1$ to N.

The algorithm above is used to generate price paths for the exponential $N \circ IG$ model. In order to generate price paths for the Black-Scholes model replace $N \circ IG(\alpha, \beta, \mu, \delta)$ in step 1 with $N\left(\mu, \sigma^2\right)$.

E.4 Linear Lévy models

1. Generate an n by N matrix Δ where each element of Δ follows a $N \circ IG(\alpha, \beta, \mu, \delta)$ distribution.

2. Calculate $\Delta^{\star}(j,k) = \sum_{l=1}^{j} \Delta(l,k)$.

3. Repeat step 2 for $j = 1$ to n.

4. Repeat steps 2 and 3 for $k = 1$ to N.

5. For fixed k, if $\min(\Delta^{\star}(1,k), \Delta^{\star}(2,k), ..., \Delta^{\star}(n,k)) \leq 0$, denote the minimum value of j for which $\Delta^{\star}(j,k) \leq 0$ by τ.

6. Set $\Delta^{\star}(\tau,k) = \Delta^{\star}(\tau+1,k) = ... = \Delta^{\star}(n,k) = 0$.

7. Repeat steps 5 and 6 for $k = 1$ to N.

8. Let R denote an n by N matrix where each element equals r.

9. Calculate $R^{\star}(j,k) = \sum_{l=1}^{j} R(l,k)$.

10. Calculate $S(j,k) = \Delta^{\star}(j,k) \exp(R^{\star}(j,k))$.

11. Repeat steps 9 and 10 for $j = 1$ to n.

12. Repeat steps 9 to 11 for $k = 1$ to N.

The algorithm above is used to generate price paths for the linear $N \circ IG$ model. In order to generate price paths for the linear Brownian motion model replace $N \circ IG(\alpha, \beta, \mu, \delta)$ in step 1 with $N(\mu, \sigma^2)$.

E.5 Time changed linear Brownian motion models

1. Generate an n by N matrix T where each element of T follows a $\log N(\alpha, \beta)$ distribution.

2. Generate an n by N matrix Δ where each element of Δ follows a $N(0,1)$ distribution.

3. Calculate $\Delta_1(j,k) = \mu T(j,k) + \sigma\sqrt{T(j,k)}\Delta(j,k)$.

4. Calculate $\Delta^\star(j,k) = \sum_{l=1}^{j} \Delta_1(l,k)$.

5. Repeat steps 3 and 4 for $j = 1$ to n.

6. Repeat steps 3 to 5 for $k = 1$ to N.

7. For fixed k, if $\min(\Delta^\star(1,k), \Delta^\star(2,k), ..., \Delta^\star(n,k)) \leq 0$, denote the minimum value of j for which $\Delta^\star(j,k) \leq 0$ by τ.

8. Set $\Delta^\star(\tau,k) = \Delta^\star(\tau+1,k) = ... = \Delta^\star(n,k) = 0$.

9. Repeat steps 7 and 8 for $k = 1$ to N.

10. Let R denote an n by N matrix where each element equals r.

11. Calculate $R^\star(j,k) = \sum_{l=1}^{j} R(l,k)$.

12. Calculate $S(j,k) = \Delta^\star(j,k)\exp(R^\star(j,k))$.

13. Repeat steps 11 and 12 for $j = 1$ to n.

14. Repeat steps 11 to 13 for $k = 1$ to N.

The algorithm above is used to generate price paths for the linear lognormal-normal model. In order to generate price paths for the linear Pareto-normal model or the linear gamma-normal model replace $\log N(\alpha,\beta)$ in step 1 with $Pareto(g,h)$ or $gamma(a,b)$.

E.6 Time changed exponential Brownian motion models

1. Generate an n by N matrix T where each element of T follows a $\log N(\alpha,\beta)$ distribution.

2. Generate an n by N matrix Δ where each element of Δ follows a $N(0,1)$ distribution.

3. Calculate $\Delta_1(j,k) = \mu T(j,k) + \sigma\sqrt{T(j,k)}\Delta(j,k)$.

4. Calculate $\Delta^{\star}(j,k) = \sum_{l=1}^{j} \Delta_1(l,k)$.

5. Calculate $S(j,k) = \exp(\Delta^{\star}(j,k))$.

6. Repeat steps 3 to 5 for $j = 1$ to n.

7. Repeat steps 3 to 6 for $k = 1$ to N.

The algorithm above is used to generate price paths for the exponential lognormal-normal model. In order to generate price paths for the exponential Pareto-normal model or the exponential gamma-normal model replace $\log N(\alpha,\beta)$ in step 1 with $Pareto(g,h)$ or $gamma(a,b)$.

Appendix F

Options available in the hypothetical Heston market

Table F.1 shows the market prices of the digital barrier call options (calculated using the Heston model) together with the corresponding times to maturity and barrier levels.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.05	0.2717	0.4988	0.6006	0.6985
1.075	0.1061	0.3054	0.4302	0.5694
1.1	0.0278	0.1709	0.2745	0.4402
1.125		0.0805	0.1723	0.3253
1.15		0.0353	0.1061	0.2392
1.175			0.0594	0.1799
1.2			0.0331	0.1196
1.225				0.0708
1.25				0.0469
1.275				0.0287

Table F.1: Digital barrier call option prices.

Table F.2 shows the market prices of the down-and-out barrier call options together with the corresponding barrier levels as well as the times to maturity. All of the down-and-out barrier options have a strike price of 1.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
0.99	0.0219	0.0337	0.0460	0.0640
0.98	0.0211	0.0314	0.0451	0.0634
0.97	0.0195	0.0286	0.0431	0.0617
0.96	0.0157	0.0254	0.0394	0.0600
0.95		0.0197	0.0350	0.0569
0.94			0.0299	0.0500
0.93			0.0225	0.0425
0.92				0.0355
0.91				0.0265

Table F.2: Down-and-out barrier call option prices.

Table F.3 shows the market prices of the up-and-out barrier call options as well as the corresponding times to maturity and barrier levels. The strike price of each up-and-out barrier call option equals 1.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.1	0.0192	0.0169	0.0143	
1.125	0.0213	0.0246	0.0221	0.0156
1.15	0.0217	0.0300	0.0293	0.0235
1.175	0.0224	0.0327	0.0357	0.0308
1.2		0.0345	0.0401	0.0394
1.225		0.0353	0.0435	0.0479
1.25			0.0450	0.0528
1.275				0.0572
1.3				0.0597

Table F.3: Up-and-out barrier call option prices.

Appendix G

Starting values for calibration procedures

Tables G.1 to G.11 below list the possible starting values considered when calibrating the option pricing models in Chapter 7.

κ	0.1	0.25	
η	0.85	1	
λ	0.03	0.05	0.07
α	50	90	150
β	-25	-5	0
δ	0.005	0.01	0.02

Table G.1: Possible starting values for the time changed exponential $N \circ IG$ model.

α	10	30	50	150	300
β	-95	-45	-5	0	10
δ	0.001	0.005	0.01	0.05	0.1

Table G.2: Possible starting values for the exponential $N \circ IG$ model.

σ	0.005	0.0075	0.01	0.0125	0.015

Table G.3: Possible starting values for the Black-Scholes model.

α	10	30	50	150	300
β	-95	-45	-5	0	10
δ	0.001	0.005	0.01	0.05	0.1

Table G.4: Possible starting values for the linear $N \circ IG$ model.

σ	0.005	0.0075	0.01	0.0125	0.015

Table G.5: Possible starting values for the linear Brownian motion model.

α	0.05	0.075	0.1	0.125	0.15
β	0.005	0.0075	0.01	0.0125	0.015
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.6: Possible starting values for the linear lognormal-normal model.

g	1	2	3	4	5
h	3	3.5	4	4.5	5
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.7: Possible starting values for the linear Pareto-normal model.

a	0.5	1	1.5	2	2.5
b	0.5	1	1.5	2	2.5
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.8: Possible starting values for the linear gamma-normal model.

α	0.05	0.075	0.1	0.125	0.15
β	0.005	0.0075	0.01	0.0125	0.015
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.9: Possible starting values for the exponential lognormal-normal model.

g	1	2	3	4	5
h	3	3.5	4	4.5	5
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.10: Possible starting values for the exponential Pareto-normal model.

a	0.5	1	1.5	2	2.5
b	0.5	1	1.5	2	2.5
σ	0.005	0.0075	0.01	0.0125	0.015

Table G.11: Possible starting values for the exponential gamma-normal model.

Appendix H

Parameter sets obtained using calibration

Tables H.1 to H.11 show the parameter sets obtained when calibrating the option pricing models to the market option prices in Chapter 7.

	κ	η	λ	α	β	δ
DB	0.1015	0.9562	0.0507	50.7288	-5.0728	0.0051
DOB	0.1009	0.8421	0.0506	150.9124	0	0.0206
UOB	0.1018	0.9884	0.0504	51.0908	-4.9679	0.0051

Table H.1: Calibrated parameter sets for the time changed exponential $N \circ IG$ model.

	α	β	μ	δ
DB	49.9992	-5.2514	0.0009	0.005
DOB	143.625	-47.0222	0.0038	0.01
UOB	52.5	-5	0.0008	0.005

Table H.2: Calibrated parameter sets for the exponential $N \circ IG$ model.

	μ	σ
DB	0.0004	0.0094
DOB	0.0004	0.0092
UOB	0.0004	0.0094

Table H.3: Calibrated parameter sets for the Black-Scholes model.

	α	β	μ	δ
DB	48.7427	-5.1645	0.0005	0.005
DOB	147.265	-46.0764	0.0034	0.0103
UOB	48.4612	0.0002	0	0.0052

Table H.4: Calibrated parameter sets for the linear $N \circ IG$ model.

	μ	σ
DB	0	0.0097
DOB	0	0.0079
UOB	0	0.01

Table H.5: Calibrated parameter sets for the linear Brownian motion model.

	α	β	μ	σ
DB	0.053	0.0143	0	0.0095
DOB	0.0503	0.0075	0	0.0104
UOB	0.0753	0.0052	0	0.0099

Table H.6: Calibrated parameter sets for the linear lognormal-normal model.

	g	h	μ	σ
DB	4.1132	3.4468	0	0.0078
DOB	2.9752	4.0966	0	0.0101
UOB	2.0892	3.0584	0	0.0101

Table H.7: Calibrated parameter sets for the linear Pareto-normal model.

	a	b	μ	σ
DB	0.965	1.0502	0	0.0098
DOB	0.5	2	0	0.1
UOB	1.041	0.9672	0	0.0105

Table H.8: Calibrated parameter sets for the linear gamma-normal model.

	α	β	μ	σ
DB	0.0512	0.0153	0.0003	0.0095
DOB	0.0497	0.0075	0.0003	0.0105
UOB	0.0767	0.0103	0.0003	0.0089

Table H.9: Calibrated parameter sets for the exponential lognormal-normal model.

	g	h	μ	σ
DB	3.9444	3.4417	0.0002	0.0079
DOB	3.0462	4.1264	0.0004	0.0103
UOB	2.1115	4.376	0.0006	0.0121

Table H.10: Calibrated parameter sets for the exponential Pareto-normal model.

	a	b	μ	σ
DB	0.9916	1.0306	0.0003	0.0101
DOB	1.9605	2.1262	0.0001	0.0051
UOB	0.965	0.5303	0.0007	0.0127

Table H.11: Calibrated parameter sets for the exponential gamma-normal model.

Appendix I

Options available in the hypothetical exponential $N \circ IG$ market

Table I.1 shows the market prices of the digital barrier call options (calculated using the exponential $N \circ IG$ model) as well as the corresponding times to maturity and barrier levels.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.05	0.5157	0.6345	0.6983	0.7550
1.075	0.3820	0.5321	0.6031	0.6800
1.1	0.2504	0.4247	0.5118	0.5911
1.125		0.3222	0.4206	0.5122
1.15		0.2556	0.3472	0.4461
1.175			0.2857	0.3799
1.2			0.2192	0.3158
1.225				0.2655
1.25				0.2378
1.275				0.1984

Table I.1: Digital barrier call option prices.

Table I.2 shows the market prices of the down-and-out barrier call options together with the corresponding barrier levels and times to maturity. All of the down-and-out barrier options have a strike price of 1.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
0.99	0.0186	0.0192	0.0206	0.0211
0.98	0.0244	0.0260	0.0273	0.0275
0.97	0.0300	0.0343	0.038	0.0394
0.96	0.0333	0.0396	0.0448	0.0473
0.95		0.0428	0.0491	0.0528
0.94			0.0528	0.0563
0.93			0.0565	0.0638
0.92				0.0672
0.91				0.0713

Table I.2: Down-and-out barrier call option prices.

Table I.3 shows the market prices, times to maturity and barrier levels of the up-and-out barrier call options. All of the up-and-out barrier options have a strike price of 1.

Barrier	$T = 21$	$T = 46$	$T = 68$	$T = 111$
1.1	0.0106	0.0043	0.0023	
1.125	0.0160	0.0087	0.0051	0.0032
1.15	0.0212	0.0129	0.0091	0.0060
1.175	0.0280	0.0186	0.0136	0.0102
1.2		0.0242	0.0203	0.0144
1.225		0.0294	0.0256	0.0175
1.25			0.0304	0.0205
1.275				0.0258
1.3				0.0305

Table I.3: Up-and-out barrier call option prices.

References

Allison, J.S. 2008. Bootstrap based hypothesis testing. Potchefstroom: NWU. (Thesis - PhD).

Anscombe, F.J. 1949. Large sample theory of sequential estimation. *Biometrika,* 36:455-458.

Attari, M. 2004. Option pricing using Fourier transforms: a numerically efficient simplification. Working paper. Boston: Charles River Associates, Inc. Available at SSRN: http://ssrn.com/abstract=520042.

Barndorff-Nielsen, O.E. 1997. Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics,* 24:1-13.

Barndorff-Nielsen, O.E. 1998. Processes of normal inverse Gaussian type. *Finance and Stochastics,* 2:41-68.

Barndorff-Nielsen, O.E. & Shiryaev, A. 2010. Change of time and change of measure. New Jersey: World Scientific.

Björk, T. 2011. An introduction to point processes from a martingale point of view. Unpublished. Available at http://www.math.kth.se/matstat/fofu/PointProc3.pdf.

Black, F. & Scholes, M. 1973. The pricing of options and corporate liabilities. *The Journal of Political Economy,* 3:637-654.

Carr, P. & Madan, D.B. 1999. Option valuation using the fast Fourier transform. *Journal of Computational Finance,* 4:61-73.

Clark, P.K. 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica,* 41(1):135-156.

Cont, R. 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance,* 1:223-236.

Cont, R. & Tankov, P. 2004. Financial modelling with jump processes. Boca Raton: Chapman & Hall.

Cox, D.R. & Hinkley, D.V. 1974. Theoretical statistics. London: Chapman and Hall.

Davis, M. & Etheridge, A. 2006. Louis Bachelier's theory of speculation: the origins of modern finance. Princeton: Princeton University Press.

Delbaen, F. & Schachermayer, W. 2005. The mathematics of arbitrage. New York: Springer.

Eberlein, E. & Jacod, J. 1997. On the range of option prices. *Finance and Stochastics,* 1:131-140.

Feller, W. 1971. An introduction to probability theory and its applications, volume 2. London: John Wiley & Sons, Inc.

Feuerverger, A. & McDunnough, P. 1981. On some Fourier methods for inference. *Journal of the American Statistical Association,* 76:379-387.

Fouque, J.P., Papanicolaou, G. & Sircar, K.R. 2001. Derivatives in financial markets with stochastic volatility. Cambridge: Cambridge University Press.

Gerber, H.U. & Shiu, E.S.W. 1994. Option pricing by Esscher transforms (with discussions). *Transactions of the Society of Actuaries,* 46:99-191.

Harrison, J.M. & Pliska, S.R. 1981. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications,* 11:215-260.

Heathcote, C.R. 1977. The integrated squared error estimation of parameters. *Biometrika,* 64:255-264.

Heston, S.L. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies,* 6:327-343.

Hull, J.C. 2009. Options, futures, and other derivatives. 7th ed. London: Pearson Prentice Hall.

Kemp, M.H.D. 2009. Market consistency: model calibration in imperfect markets. Chichester: John Wiley & Sons, Inc.

Mandelbrot, B. 1963. The variation of certain speculative prices. *Journal of Business,* 36:394-419.

Sato, K. 2005. Lévy processes and infinitely divisible distributions. Cambridge: Cambridge University Press.

Schoutens, W. 2003. Lévy processes in finance: pricing financial derivatives. Chichester: John Wiley & Sons, Inc.

Schoutens, W., Simons, E. & Tistaert, J. 2004. A perfect calibration! Now what? *Wilmott magazine,* 8:66-78.

Shiryaev, A.N. 2003. Essentials of stochastic finance: facts, models and theory. New Jersey: World Scientific.

Spiegel, M.R. & Liu, J. 1999. Mathematical handbook of formulas and tables. 5th ed. New York: McGraw-Hill.

Thorin, O. 1977a. On the infinite divisibility of the Pareto distribution. *Scandinavian Actuarial Journal,* 31-40.

Thorin, O. 1977b. On the infinite divisibility of the lognormal distribution. *Scandinavian Actuarial Journal,* 121-148.

Van der Vaart, A.W. 1998. Asymptotic statistics. Cambridge: Cambridge University Press.

Venter, J.H., De Jongh, P.J. & Griebenow, G. 2005. NIG-GARCH models based on open, close, high and low prices. *South African Statistical Journal,* 39:79-101.

Yahoo!finance. 2012. Google Inc. (GOOG). Available at: http://finance.yahoo.com/q/pr?s=GOOG. Date of access: 12 May 2012.

Yahoo!finance. 2012. PowerShares QQQ Trust, Series 1 (QQQ). Available at: http://finance.yahoo.com/q/pr?s=QQQ. Date of access: 12 May 2012.

Yahoo!finance. 2012. S&P 500 (^GSPC). Available at: http://finance.yahoo.com/q/pr?s=^GSPC. Date of access: 12 May 2012.

Yao, L., Yang, G. & Yang, X. 2011. A note on the mean correcting martingale measure for geometric Lévy processes. *Applied Mathematics Letters,* 24:593-597.

Yu, J. 2004. Empirical characteristic function estimation and its applications. *Econometric Reviews,* 23:93-123.